# COMPUTING THE SEMANTICS OF VERB ALTERNATIONS

## A CONSTRUCTIONAL APPROACH

by

Esther Seyffarth

A dissertation submitted to the
Faculty of Arts and Humanities at
Heinrich Heine University
for the degree of

Doctor of Philosophy

Supervisor:
Laura Kallmeyer

Originally submitted May 2nd, 2023
Revised version May 8th, 2025

D61

Esther Seyffarth
*Computing the Semantics of Verb Alternations:*
A constructional approach

# ABSTRACT

This thesis explores the syntactic and semantic properties of verbs, their arguments, and the argument structure constructions they appear in, and investigates in what ways these properties matter for a successful computational treatment of verb alternations. Verbs that participate in a verb alternation are able to instantiate an alternation-specific set of constructions in which certain syntactic argument positions are linked to certain semantic roles (Levin 1993).

A construction is defined by the combination of its syntactic form and semantic meaning (Goldberg 1995: 1). As the syntactic form of a construction is not necessarily unique, ambiguity is possible: constructions can share a syntactic form and differ in their meaning. This can pose issues for computational systems when they encounter sentences involving alternating verbs. A reliable distinction of constructions is a prerequisite for a successful interpretation of such sentences.

This thesis approaches these issues from two angles. First, it presents a metagrammar for a lexicalized tree adjoining grammar (Joshi & Schabes 1997) with semantic representations in the form of typed, recursive frames (Kallmeyer & Osswald 2013). The metagrammar models the syntax and semantics of a selection of English verb alternations, as well as a number of argument structure constructions whose syntactic form is identical to that of constructions in the alternations. Second, the thesis presents a set of classifiers to determine whether a given verb participates in a specific alternation or not, based on corpus attestations of the verb in question and of known alternating and non-alternating verbs. The classification features are designed to approximate the overlap in each verb's selectional preferences on arguments in the relevant syntactic positions in different constructions.

Based on the findings reported in this thesis, an optimal approach to a computational treatment of alternating verbs would be one that combines the advantages of the different approaches presented here. Handcrafted models like the metagrammar implemented in this thesis do not scale well, but are promising with respect to a reliable handling of ambiguous sentences. Instead of classifying verbs based on features approximating selectional restriction overlap, it could be beneficial to predict exactly the properties that are modeled in the metagrammar. This would involve learning a type hierarchy of a suitable granularity, including incompatibility and subtype relationships, as well as learning type requirements imposed on specific arguments both by argument structure constructions and by verbs. The implementation of a system that learns these properties and the corresponding type hierarchy is a promising direction for future work.

# ACKNOWLEDGEMENTS

The publication of this PhD thesis marks the end of an amazingly inspiring and enriching phase of my academic life. I will always value my memories of the time I spent at Heinrich Heine University, the connections I made both within and outside HHU's computational linguistics department, and of course, the many things I learned throughout my undergraduate and graduate studies.

The most important person accompanying me on this academic journey was Laura Kallmeyer, my *Doktormutter*, who tirelessly provided guidance, perspective, patience and feedback over the years. I am deeply grateful for her leadership and her support. I'm aware that some PhD advisors tend to either micro-manage or be too absent; but Laura was always available to give advice when I needed it, while also giving me space to try out my own ideas and make my own mistakes.

I have many others to thank for supporting my research and teaching me along the way, but all my thanks somehow come back to Laura again for the research environment and community she nurtured in Düsseldorf. The DFG-funded Collaborative Research Center 991, which was directed by Laura, and which I joined as a PhD student late in the funding period, was a stimulating research world to find myself in. Everyone I met there was smarter than me, as well as kind, curious and cooperative. Orienting myself as a fresh PhD student was made much easier and more pleasant by the hard work of Eva Gentes, who not only took care to make every member of the CRC's graduate school feel welcome, but also made sure I knew exactly what was expected of me and what I could expect from the graduate program and the people involved in it. Thanks in particular also to Wiebke Petersen for giving me my first opportunity to collaborate in a research project and publish a conference paper, and for taking on the role of the second grader for this thesis. Thanks to Rainer Osswald for all the paper and book recommendations that he suggested to me each time we talked – they were all relevant, and always improved my theoretical understanding of what was going on.

Being a member of the computational linguistics department at Heinrich Heine University, headed by Laura Kallmeyer, was an amazing experience that I benefited from immensely, particularly in regard to all the PhD students who became my friends, the postdocs and other department members who were always happy to make their expertise, computing power, or libraries available to me to help with specific aspects of my work, and the numerous visitors from near and far who gave so many interesting and relevant talks. Thanks to Tatiana Bladier, Benjamin Burkhardt, Rafael Ehren, Leda Berio and Kurt Erbach for their friendship during our shared time in graduate school – you have all been academic and personal role models to me in numerous ways. I am grateful for all the coffee chats, Mensa meals, summer parties, *Stammtische*, and the family-friendly atmosphere and working conditions in the department. I am grateful for the numerous motivated and passionate students I encountered whenever I was teaching classes. And of course, I am grateful for the funding that allowed me to visit summer schools and conferences to broaden my horizon beyond what I already knew, and to receive feedback and ideas for my ongoing research. I also want to thank in particular Tim Marton, Olga Böse, Janine Willems, and Lena Schwarz, whom I was always able to count on when I needed documents printed, rooms booked, equipment ordered, a bureaucratic process explained to me for the third time, a document faxed across campus for me, or any other sort of administrative help.

Out of all the computational linguistics department colleagues just mentioned, I am particularly grateful to Simon Petitjean for not only providing and maintaining the software infrastructure for the XMG implementations presented in this thesis, but also helping me with numerous tips, answers to confused questions, extremely helpful feedback on an earlier draft of this work and on many previous versions of the metagrammars I developed, and of course, a multitude of bug fixes for the tools used.

Parts of the data processing for the alternation identification experiments reported in Part IV were conducted using HHU's high-performance computing cluster, HILBERT.

Further thanks go to everyone whose contributions to LaTeX packages, templates, or settings helped me typeset this document, most notably André Miede and Ivo Pletikosić for the `classicthesis` package, Felix Kopecky for the `langsci-avm` package, and Timm Lichte for his TAG-friendly and frame-friendly Ti*k*Z and `forest` settings.

An extremely valuable contribution to me being able to write this thesis came from my family and friends outside of work, who supported me by directing my focus to all the things beyond work that are good and fun and interesting and important in the world. Thank you, Henrieke, Annika, Roman, Rebekka, Heike, Kai, Kathrin, Jay, M, Luke, Ronja, Vanessa, Anja, and everyone else. Finally, many thanks to all the neighborhood toddlers and parents whose friendship and community gave me so much joy and happiness particularly during the last months of writing this thesis.

Writing this PhD thesis made me feel both empowered and humbled: I am proud of my work, but also acutely aware of many things that could still be done to improve it. But everyone knows that a good thesis is a finished thesis, so I apologize for leaving unattended for now the long list of possible additions, tweaks, and details that could be studied further.

The present document is a slightly adapted version of the thesis I defended in 2023. I thank Laura Kallmeyer, Wiebke Petersen, and two anonymous Language Science Press reviewers for their helpful comments.

# CONTENTS

LIST OF FIGURES

# LISTINGS

# Part I

# INTRODUCTION

This part of the thesis motivates the research being presented, introduces the scope of the work, and gives an overview of the thesis structure.

# 1 INTRODUCTION

Verbs play an important role in language, and consequently, in natural language processing. They determine a large part of the syntax and semantics in a sentence, and this is reflected in various frameworks in computational linguistics. For instance, in dependency grammar, verbs are the most likely root element of dependency trees (Tesnière 2015: 7), in Head-Driven Phrase Structure Grammar, the verb in a sentence is the head of the sentence's representation (Pollard & Sag 1994: 34), and in Tree Adjoining Grammar, verbs anchor initial trees which can then be extended by substitution or adjunction to represent a full sentence (Abeillé et al. 1990: 11, Joshi & Schabes 1997: 29).

On the syntactic side, verbs can appear with a number of arguments in certain positions in the sentence, such as subjects, direct objects, or prepositional objects. On the semantic side, the meaning of a verb can involve a certain set of participants filling specific semantic roles or thematic roles. For instance, a verb denoting an activity can involve an ACTOR, which can either be explicitly referred to or be left unexpressed. Mapping the semantic roles of a verb to its syntactic arguments is not trivial, since there is no requirement for a one-to-one correspondence between syntactic argument positions and semantic roles (Van Valin 2005: 64). For instance, a verb may typically express its ACTOR in the subject position, but in a passivized sentence, the subject position instead contains the UNDERGOER.

A phenomenon that has received some attention in linguistics since the second half of the twentieth century are so-called verb alternations, also known as argument alternations, diathesis alternations, or valence alternations (for an overview, see Levin 2015). Each verb alternation is characterized by a set of argument structure constructions, also called alternants, that allow participating verbs to express their semantic roles (or a subset of those roles) in different syntactic positions. An example is the dative alternation, which is illustrated in (1).[1]

(1) a. Tashtego gave Queequeg his bucket.
    b. Tashtego gave his bucket to Queequeg.

The verb-specific semantic roles that are associated with the verb *give* are a "giver", a "thing that is being given" and a "person who is being given something". The latter two are commonly referred to as THEME and RECIPIENT[2]. In (1a), these two arguments are expressed in a double object construction, while (1b) expresses the THEME, the *bucket*, in the direct object position and the human RECIPIENT, *Queequeg*, as a prepositional object. In other words, two different syntactic realizations of the verb's semantic roles are available.

The two alternants shown in (1) seem to be nearly-synonymous paraphrases: a *giving* event is being described in which the giver, *Tashtego*, successfully transfers an object, the *bucket*, to a recipient, *Queequeg*. However, other instances of this alternation have been

1 The example sentences given throughout this thesis involve the human characters Tashtego, Queequeg, Daggoo, Ahab, Ishmael, and Starbuck from Herman Melville's *Moby Dick*.

2 In this thesis, names of semantic roles are typeset in small caps.

argued to show differences in the semantics of the constructions, as well as differences in the availability of one or the other variant depending on verb and argument choice; for an overview, see e.g. Goldberg (2002) and Levin (2006, 2015). Claims have been made that the double object variant as in (1a) denotes a caused possession, while the prepositional variant as in (1b) denotes a caused motion (Harley 2002: 40, Green 1974: 99,110, Goldsmith 1980: 424), or that the double object variant denotes a successful transfer, while the prepositional variant merely denotes an attempted transfer that is not necessarily completed (Green 1974: 110,117, Goldberg 1992: 49, Goldberg 1995: 147, Krifka 1999: 260, Harley 2002: 42, Rappaport Hovav & Levin 2008: 150).

Certain examples indicate that the distinction between successful and unsuccessful transfer is impacted by the chosen verb: *throw* does not entail successful transfer in either variant, but *give* is inherently a caused-possession verb and entails successful transfer in both variants (Oehrle 1976: 24, Rappaport Hovav & Levin 2008: 138). For verbs that license different entailments and implications in the two variants, that behavior has also been ascribed to verbal polysemy, such that one construction selects one sense of the verb and the other selects another sense (Speas 1990: 83, Harley 2002: 65, Krifka 2004: 29, Pinker 2013: 82). This contrasts with the idea that the semantic differences are instead mostly due to the different argument structure constructions, with the verb having some basic meaning and being incorporated into the constructions in different ways to yield distinct meanings (Goldberg 1992: 45, Goldberg 1995: 141).

Additionally, the constructions involved in the dative alternation have been found to impose different conditions on the THEME and RECIPIENT arguments. For instance, the RECIPIENT has to be construable as an animate entity in the case of the double object construction, while the prepositional variant can also take a non-animate entity in this role (Green 1974: 103, Goldsmith 1980: 425, 430, Harley 2002: 37). Harley (2002: 37) argues that the prepositional object in the prepositional variant is thematically a location, not necessarily a possessor, and that this difference is why certain types of entities are allowed in this position but not the corresponding position in the other alternant. Some relate the semantic differences to the idea that the first NP argument in the double object construction is more "affected" than the corresponding argument in the prepositional variant (Green 1974: 101, Oehrle 1976: 126, 129, Jackendoff 1990: 195, Goldberg 1995: 146, Pinker 2013: 98).

Beyond all these considerations, other factors also seem to be at play when it comes to whether certain arguments are allowed in each position of the variants involved in the dative alternation. These factors include information structural properties (Bresnan et al. 2007: 11, Bresnan & Nikitina 2010: 172-175), constituent weight (Quirk et al. 1985: 1375), givenness (Ransom 1979: 215, Smyth, Prideaux & Hogan 1979: 27, Gundel 1988: 210, Collins 1995: 41, Thompson 1995: 158, Arnold et al. 2000: 30, Wasow 2002: 2, Gries 2003: 12, Snyder 2003: 15), and prosodic considerations (Anttila, Adams & Speriosu 2010).

The dative alternation is among the most well-studied verb alternations in the English language (Levin 2015: 67). Some of the discussions concerning the behavior of verbs in the context of this alternation may also be applicable to other (English) alternations, while others seem to be specific to this particular alternation. Levin (2015: 66) notes that for some alternations, there is no obvious truth-conditional difference between their variants; for other alternations, like the causative-inchoative alternation, the meaning characterizing one argument structure construction clearly differs from the meaning associated with the other construction. In that alternation, alternants are not simply paraphrases of each other, and each of them allows a different subset of semantic roles to be expressed. Examples for the two constructions associated with the causative-inchoative alternation are given in (2).

(2) a. The door opened.
    b. Bildad opened the door.

An overview of a large number of English verb alternations and the verbs and verb classes participating in them can be found in Levin (1993). In her perspective, verb alternations require verbs to possess certain facets of meaning that allow them to appear in the relevant constructions. For more on this and other theoretical views on verb alternations, see Section 2.2 of this thesis.

The minimal number of argument structure constructions involved in an alternation is two; Levin also lists alternations involving three constructions, such as the Search alternations (Levin 1993: 70):

> Verbs of searching seem to have available three alternate ways of expressing their arguments: 'NPl V NP2 *in* NP3', 'NPl V NP3 *for* NP2', 'NPl V *for* NP2 *in* NP3'. Different verbs of searching display different subsets of these patterns, giving rise to a variety of alternations in the expression of their arguments. (...)
>
> (224) a. Ida hunted the woods for deer.
>       b. Ida hunted for deer in the woods.
>       c. Ida hunted deer in the woods.

As illustrated above in the context of the dative alternation, verb alternations pose a number of challenges for both linguistics and natural language processing. The correct representation of an event or state of affairs denoted by a sentence can only be found if the constructions instantiated by that sentence are identified correctly. Some of the constructions that are available to verbs that participate in certain alternations are unavailable to other verbs. To further complicate things, different constructions can share their surface forms but have different semantics; in other words, a sentence's syntactic structure alone is not a sufficient indicator for determining whether or not it instantiates an (alternation-specific) construction. For example, the causative alternant of the causative-inchoative alternation has the syntactic form of a transitive sentence; not all transitive sentences are instances of the causative construction. Alternation participation may correlate with certain meaning facets of verbs as argued by Levin (1993), but such latent properties are also challenging to acquire automatically.

This thesis approaches the topic of verb alternations from a perspective that is grounded in Construction Grammar (Goldberg 1995, 2013) and frame semantics (Kallmeyer & Osswald 2013, Lichte & Petitjean 2015, Petersen 2015, Osswald & Van Valin 2014). Part II provides an introduction to these frameworks, motivates their usefulness for the purposes of this thesis, and summarizes the two theoretical perspectives on verb alternations that are most relevant to the work presented here.

Part III proposes a strategy for modeling the alternation behavior of verbs under the assumption that the relevant meaning facets are known and encoded in the model. The proposed approach to this task relies on meaning facets as main factors controlling alternation participation, as argued for by Levin (1993). The model also implements the constructionist perspective in which meaning is contributed not only by the lexicon entries for verbs and their arguments, but also by the constructions that are instantiated by a given sentence (Goldberg 1995: 2, Goldberg 2013: 19). A verb can only instantiate a given argument structure construction if the meanings contributed by the verb and the construction are compatible with each other. The different meaning contributions are represented as semantic frames in the sense proposed by Kallmeyer & Osswald (2013), Lichte & Petitjean (2015), Osswald & Van Valin (2014), Petersen (2015), while the relevant syntactic structures are modeled via elementary trees in a lexicalized tree adjoining grammar (Vijay-Shanker 1987, Joshi & Schabes 1997, Abeillé & Rambow 2000).

The model takes the form of a metagrammar describing a lexicalized tree adjoining grammar that can be compiled to yield a grammar with which input sentences instantiating the relevant constructions can be parsed. The focus of the implementation is on the challenge of distinguishing constructions whose syntactic structure is identical, while their semantics differ. Alternation-specific constructions are only available to verbs participating in the relevant alternation.

In this task, alternation participation is encoded in the form of constraints indicating construction availability, based on a semantic type hierarchy in which all semantic types appearing in the grammar are described with respect to subtype relationships, mutual incompatibility, and type requirements on frame attributes. The implementation shows that it is essential to correctly predict the availability of every candidate construction in order to produce the correct semantic representation for various sentences sharing the same syntactic form.

The assumption that all relevant meaning facets and semantic type relationships are known and easily accessible is of course a simplification. This is why Part IV of the thesis presents a series of approaches to learning the distinction between alternating and non-alternating verbs automatically, based on dependency-parsed text data from an English web corpus (ENCOW, Schäfer & Bildhauer 2012, Schäfer 2015) and lists of alternating verbs from VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007). The classifiers developed in that thesis part rely on a number of features whose purpose is to approximate whether or not each verb has a tendency to instantiate the constructions associated with a specific alternation.

Using only dependency-parsed corpus data means that the classification task requires distinguishing between constructions that share the same syntactic structure. Continuing with the focus of the previous thesis part, this part also concerns itself with contrasting verbs that can appear in the same syntactic environments as alternating verbs, but which do not themselves participate in the alternation. The classifiers are thus applied exclusively to verbs that are observed at least once in the syntactic environments associated with each alternant of an alternation.

The classifiers are trained and evaluated on three English verb alternations: the causative-inchoative alternation, the instrument subject alternation, and the together reciprocal alternation (intransitive). The classification features either rely exclusively on the observable behavior of each alternation candidate verb, or additionally refer to pre-trained resources to approximate semantic properties of the observed instances of each verb. The classification results show that the constructions associated with different alternations pose different challenges. For instance, constructions that involve prepositional objects are subject to the high degree of polysemy of English prepositions; the classifier features do not resolve the resulting ambiguity well. This makes it difficult for the classifiers to distinguish between alternation-specific constructions and other constructions that are unrelated to the given alternation, but happen to share the same syntactic structure and preposition.

Additionally, alternations in a language can differ in their frequency. VerbNet lists over 500 participating verbs for the instrument subject alternation, but only 50 verbs for the together reciprocal alternation (intransitive). Not all verbs listed in VerbNet for an alternation are guaranteed to be attested in a corpus such as ENCOW. The classification results show that the classifiers for the together reciprocal alternation (intransitive) suffer more from sparsity than the classifiers for alternations with more participating verbs. At the same time, this alternation benefits most strongly from a feature set which augments observed corpus data with hypothetical additional instances of the relevant constructions.

One of the goals of the alternating verb identification task is to be able to extend a known set of alternating verbs based on corpus data, for instance in the context of a language for which VerbNet-like resources are either not available or not comprehensive. A useful asset for this goal is the set of "false positives" identified by the classifiers, that is, verbs that are initially assumed to not participate in an alternation, but which are found to resemble alternating verbs based on their feature values. Part IV of the thesis closes with a description of a strategy using manual annotation of false positives to identify good new alternation candidates.

Certain aspects of the work presented in this thesis have been published in an earlier form. A previous version of the metagrammar model covering the syntax and semantics of the induced action alternation and the caused-motion construction (described in Chapter 6) appeared as Seyffarth (2019b). A previous version of the classifiers for the causative-inchoative alternation (described in Chapter 10) appeared as Seyffarth (2019a). Finally, a previous version of the classifiers for the causative-inchoative alternation and the instrument subject alternation, including a manual annotation step to identify good new alternation candidates (described in Chapter 10 and specifically Section 10.7), appeared as Seyffarth & Kallmeyer (2020).

# Part II

# BACKGROUND AND DEFINITIONS

This part of the thesis provides the theoretical background for the subsequent parts and describes the perspective that will be taken with regard to the phenomena under investigation. Construction grammar, the projectionist view on verb alternations by Levin (1993), and frame semantics in the sense of Kallmeyer & Osswald (2013), Lichte & Petitjean (2015), Osswald & Van Valin (2014), Petersen (2015) will be introduced.

# 2 VERB ALTERNATIONS

## 2.1 INTRODUCTION

A variety of different theoretical approaches to verb alternations exist in linguistic research. The goal of this thesis is not to examine the cognitive or psycholinguistic properties of verb alternations; instead, the focus is on aspects of verb alternations that are relevant in computational applications, where the different behavior of alternating verbs needs to be taken into account.

The following section will introduce some existing theoretical perspectives on verb alternations. Section 2.3 will then describe which aspects of two of the most prominent theories will feature in this thesis. Chapter 3 will give some background on Construction Grammar and elaborates on the way constructions will be used in this thesis. Chapter 4 describes how semantic frames will be used in this thesis to denote semantic contributions from constructions and verb lemmas.

## 2.2 THEORETICAL APPROACHES TO ALTERNATIONS

Over time, several theoretical approaches have been proposed by various linguists to account for verb alternations with respect to the cognitive processes underlying them, the psycholinguistic or pragmatic factors determining when a specific alternant is used, the semantic and syntactic differences between alternants in a particular alternation, or the commonalities between different alternations within and across individual languages. This section will briefly describe the projectionist perspective, as exemplified by Levin (1993), and the constructionist perspective, as exemplified by Goldberg (2002). A more thorough overview of different accounts of verb alternations can be found in Levin (2015).

### 2.2.1 *The projectionist perspective*

Levin (1993) compiles a large list of English verb alternations and verbs that can participate in them, grouped into semantically-coherent verb classes. Her work is motivated by the idea that verb meaning to a large extent determines verb behavior; thus, she concludes, observing verb behavior is a useful tool for researching certain components of verb meaning.

Levin observes that verbs which behave similarly tend to have shared meaning components, and illustrates this with a set of four example verbs that overlap in certain properties and license different diathesis alternations. First, she demonstrates that each of the four verbs exhibits a different behavioral pattern that allows certain alternations and disallows others. The relevant example sentences are reproduced below in (3) (adapted from Levin 1993: 6–7). Levin discusses the availability of the conative construction, the body-part possessor ascension construction, and the middle construc-

Table 1: Alternations licensed by each of the verbs *touch*, *hit*, *cut*, *break* (from Levin 1993: 7).

| | touch | hit | cut | break |
|---|---|---|---|---|
| Conative | no | yes | yes | no |
| Body-Part Possessor Ascension | yes | yes | yes | no |
| Middle | no | no | yes | yes |

tion as a proxy for the availability of the alternations that are characterized by these constructions, respectively.

(3) a. Conative:
   i. * Terry *touched* at the cat.
   ii. Carla *hit* at the door.
   iii. Margaret *cut* at the bread.
   iv. * Janet *broke* at the vase.
  b. Body-Part Possessor Ascension:
   i. touch:
    1. Terry *touched* Bill's shoulder.
    2. Terry *touched* Bill on the shoulder.
   ii. hit:
    1. Carla *hit* Bill's back.
    2. Carla *hit* Bill on the back.
   iii. cut:
    1. Margaret *cut* Bill's arm.
    2. Margaret *cut* Bill on the arm.
   iv. break:
    1. Janet *broke* Bill's finger.
    2. *Janet *broke* Bill on the finger.
  c. Middle:
   i. * Cats *touch* easily.
   ii. * Door *frames* hit easily.
   iii. The bread *cuts* easily.
   iv. Crystal vases *break* easily.

Based on the given judgments for each combination of an alternation variant and a verb, Levin gives an alternation participation overview which is reproduced here as Table 1.

To account for the different patterns of these four verbs with respect to the listed alternations, Levin looks for shared meaning components and meaning components in which the verbs differ, based on prior work by Hale & Keyser (1986: 614), Fillmore (1970b: 130), and Hale & Keyser (1987). An overview of her analysis is given in Table 2 on page 13.

Table 2: Meaning components inherent to each of the verbs *touch*, *hit*, *cut*, *break*, according to Levin (1993: 7–9).

| Meaning component | touch | hit | cut | break |
|---|---|---|---|---|
| motion | no | yes | yes | no |
| contact | yes | yes | yes | no |
| change of state | no | no | yes | yes |

The distribution of the meaning facets shown in Table 2 correlates with the distribution of alternation participation for the four verbs as shown in Table 1. Levin also notes that this observation does not only hold for the individual verbs discussed, but that each of the verbs shares the relevant meaning components and behavior pattern with a set of similar verbs, forming verb classes that are characterized by their shared syntactic and semantic properties.

Levin observes that out of the examples shown, only *break*-like verbs can appear in the causative-inchoative alternation. She explains this with the fact that this alternation is sensitive to pure change of state verbs, and verbs like *cut* cannot participate in it because they also involve notions of motion and contact, in addition to the meaning facet of a change of state. She references an analysis by Guerssel et al. (1985) which explains the difference by the fact that change of state verbs like *break* have an event structure involving only a single participant (an entity undergoing a change of state), while verbs like *cut* (which inherently involves an instrument and requires an agent to use that instrument to bring about a change of state in some entity) have an event structure involving two arguments. This is why *cut* and other verbs that share its relevant meaning components can never appear in the inchoative construction, in which only one participant is expressed.

From these patterns, Levin concludes that facets of meaning like motion, contact, or change of state should be part of the lexical representation of verbs, since they are what allows speakers to decide whether a diathesis alternation is available for a given verb or not. If the meaning facets associated with a verb are compatible with an alternation, the verb participates in that alternation. Levin adds that speakers also take into account individual verb meaning and general principles that determine behavior from meaning. Instead of specifying the alternation behavior of each verb in the lexicon explicitly, these regular relationships between verb meaning and verb behavior can be "factored out" of the lexicon, leaving only idiosyncratic information in the lexicon entry.

Understanding the relevant meaning facets for alternations also allows speakers to generalize their knowledge to verbs they are unfamiliar with, either because they are not in everyday use anymore, or because they are newly-formed. Levin (1993: 4) cites a study by Hale & Keyser (1987: 2), who argue that speakers confronted with the archaic whaling term *gally* will form different hypotheses about its meaning based on a single usage example. Based on their hypothesis regarding the verb's meaning, speakers will allow and disallow different alternations for the verb: if *gally* is assumed to mean something similar to *see*, the middle alternation would not be licensed, but under the assumption that it means something similar to *frighten*, the middle alternation would be licensed by it. Levin takes this as evidence for the idea that verb behavior is influenced by meaning (facets), and also provides an example of a new-at-the-time verb, *modem*, which can appear in the argument structure constructions that are also licensed by the semantically-similar verbs *radio* or *wire*.

Regarding the cross-linguistic validity of these ideas, Levin (1993: 10) cites evidence from the Australian language Warlpiri which also shows the conative alternation. The class of verbs that license this alternation in Warlpiri is semantically similar to the class of English verbs that participate in it (Laughren 1988: 230). However, in later work, Levin (2015: 66) cites evidence from Alexiadou, Anagnostopoulou & Schäfer (2006: 199) showing that verbs like *destroy* participate in the causative alternation in Greek, but not in English, and from Kim (1999: 25) showing that the sets of verbs participating in the locative alternation exhibit systematic differences across languages.

Majewska et al. (2018) also observe that alternation behavior can differ between languages and language families. They are concerned with translating VerbNet classes and their members from English to other languages, and find that the differences in alternation behavior require a certain degree of language-specific tuning for each target language. Meaning facets like causativity are marked morphologically in some languages, while this is not the case in English. Italian and Polish do not have a resultative construction. These are examples of factors that impact to what extent an alternation is possible in each language, and which conditions must be met by a verb in order for it to participate in an alternation.

Frense & Bennett (1996) compare the behavior of English and German verbs with respect to their participation in the conative, middle, and locative alternations. They find that there is some overlap in the verb classes licensing each of these alternations, but there are also differences: certain German verbs license the conative alternation while their English counterparts cannot participate in this alternation, and for locative alternations, each of the two languages has several verbs that can participate in the alternation while their counterparts in the other language do not participate in it. Frense & Bennett find that the middle alternation is more sensitive to the syntactic valency and lexical aspect of verbs, rather than their semantic classes. Thus, while there are some similarities in these closely-related languages, the verb classes determining possible alternations are found to be subject to cross-linguistic variation.

Levin (1993) considers subcategorization requirements of lexical items to be largely derivable from the meaning of words, with the following explanation:

> Those facets of syntactic constructions that cannot be made to follow from general principles of grammar are considered to be projections of the lexical properties of the words in these constructions. (Levin 1993: 12)

Thus, the properties of a verb that allow it to participate in a certain alternation are part of that verb's lexical entry, and are projected into the constructions involved in the alternation.

In order to illustrate the way the observed alternation behavior of verbs can provide insights into verb meaning, Levin (1993: 15) gives an example in which a verb of sound emission, *whistle*, is used with an extended meaning as a verb of directed motion: *the bullet whistled through the window*. She notes that this verb is not basically of the type *directed motion*, but the directional prepositional phrase in the sentence licenses that "sense" of the verb. The following section will present an alternative view on sentences like this.

### 2.2.2 *The constructionist perspective*

In Levin's projectionist perspective on alternations, observing verbs in the constructions of an alternation provides insight into the verbs' meaning components, and each construction in an alternation can appear with a certain sense of verbs participating in the alternation. An alternative perspective is taken by Goldberg (1995, 2002, 2013), who argues for studying argument structure constructions on their own terms, instead of relating the constructions involved in an alternation to each other in an attempt to learn something about their meanings and about the meaning of verbs participating in the alternation. Goldberg's constructionist perspective is based on Construction Grammar (Goldberg 1995), which will be discussed in more detail in Chapter 3 of this thesis. With her *surface generalization hypothesis*, Goldberg (2002: 329) stresses that attempting to derive one variant of an alternation from the other makes it difficult to capture certain so-called surface generalizations, which reflect the ways in which one argument structure construction behaves independently of the verb with which it is instantiated.

The ditransitive construction is one of the alternants in the dative alternation (Levin 1993: 45). Goldberg (2002: 330) points out that while some theories (e.g., Baker 1985: 239, Larson 1988: 350) assume distinct derivations of the ditransitive from different paraphrases, the different instances of the construction share properties with each other and differ systematically from their paraphrases. The ditransitives she discusses and their rough paraphrases are given in (4) (adapted from Goldberg 2002: 330). Note that according to the generative theories referenced by Goldberg, the ditransitive instantiated with *bought* is presumably derived from a prepositional form with *for*, while the ditransitive instantiated with *sent* is presumably derived from a prepositional form with *to*.

(4) a. Mina *bought* Mel a book.
b. Mina *bought* a book for Mel.
c. Mina *sent* Mel a book.
d. Mina *sent* a book to Mel.

The similarities between these ditransitives include the ability to distantly instantiate the theme, as in (5), the inability to distantly instantiate the recipient argument, as in (6), the inability to separate the two NP arguments, as in (7), the inability to express the theme argument via the third person singular pronoun *it*, as in (8), and the requirement for the recipient argument to be construed to be animate, as illustrated in (9) (all examples adapted from Goldberg 2002: 330–331).

(5) a. What did Mina *buy* Mel?
b. What did Mina *send* Mel?

(6) a. ?? Who did Mina *buy* a book?
b. ?? Who did Mina *send* a book?

(7) a. * Mina *bought* Mel yesterday a book.
b. * Mina *sent* Mel yesterday a book.

(8) a. ?? Mina *bought* Mel it.

b. ?? Mina *sent* Mel it.

(9) a. ?? Mina *bought* that place a box.
b. ?? Mina *sent* that place a box.

Regarding the paraphrases, their shared properties include the ability to distantly instantiate the recipient argument, as in (10), the ability to express the theme argument as *it*, as in (11), and the ability to appear with a recipient argument that is not animate, as in (12).

(10) a. Who did Mina *buy* a book for?
b. Who did Mina *send* a book to?

(11) a. Mina *bought* it for Mel.
b. Mina *sent* it to Mel.

(12) a. Mina *bought* a box for that place.
b. Mina *sent* a box to that place.

In addition to these shared properties between the ditransitives on the one hand and between the paraphrases on the other hand, Goldberg (2002: 332) points out that the different instances of the ditransitive share their information-theoretical constraints and a meaning that evokes the notion of "giving" in some way, unlike the paraphrases. Depending on the verb, ditransitives can denote related types of events such as an expected transfer (if certain satisfaction conditions imposed by the verb are met), or the absence of a transfer, or a transfer towards the agent instead of a transfer from the agent to the recipient.

Goldberg (2002: 333) considers these constructional meanings to form a natural class of closely related concepts. These observations lead her to claim that surface generalizations which are based on the shared properties of the instances of the ditransitive construction are more robust than generalizations that are based on the paraphrase relationship between a prepositional variant and its ditransitive counterpart. She considers the "paraphrases" to be instances of separate constructions: a caused-motion construction for the prepositional variant with *to*, and a transitive construction extended by a benefactive adjunct construction for the prepositional variant with *for*.

In the constructionist perspective, in order to determine the meaning of a sentence involving a particular argument structure construction that is instantiated with a particular verb, the meaning contributed by the construction must be taken into account together with the meaning contributed by the verb and its arguments (see also Chapter 3 of this thesis). Thus, the meaning of a sentence is not directly "read off" from the surface form of a construction, but comes about based on the interaction between argument structure constructions and verbs. This explains constructional ambiguity, which can be resolved by identifying the verb classes that are involved in a specific case, as well as the shared meaning between variants associated with an alternation, which is due to the shared verb in both alternants. Subtle meaning differences between paraphrase pairs are seen as a result of the verb's participant roles instantiating different sets of argument roles provided by the alternative constructions (for more on these roles and the way they are integrated, see Chapter 3).

As a consequence, Levin's example sentence *The bullet whistled through the window*, which she explains by ascribing a motion sense to the verb *whistle* in this particular context, would be analyzed differently by Goldberg. The directed-motion meaning of the sentence is contributed by the construction involving a directional prepositional phrase, and the verb is not required to contribute a motion sense lexically, but must merely be compatible with the construction (Goldberg 2002: 339). Chapters 3 and 4 of this thesis will go into more detail concerning the process of integrating verbs into constructions, and what that means practically in the context of modeling or predicting alternation behavior. As Goldberg (2002: 348) notes, an actual observed expression or construct typically involves the combination of multiple different constructions; this idea will become relevant in Part III of this thesis.

The constructionist perspective on alternations taken by Goldberg (2002) contrasts with the projectionist perspective taken by Levin (1993) mainly in their different assumptions regarding the relationship between the two variants associated with an alternation (and the theoretical relevance of that relationship), and in the presumed source of meaning differences and similarities between alternants.

Levin (2015) gives an overview of these and other theoretical accounts of verb alternations over the course of several decades. Earlier approaches, such as Emonds (1972), Hall (1965), Fillmore (1965), are known as transformational accounts. In transformational accounts, one alternant is assumed to be created based on the other via some transformation. Under this perspective, alternating verbs are considered to have one single meaning, and transformations do not change meaning. There is either one basic alternant and one alternant that is created via transformation, or both alternants are assumed to be derived from some more basic variant. Later theories, such as Rappaport & Levin (1988), Speas (1990: 83), Harley (2002: 65), Krifka (2004: 29), posit verbal polysemy as the reason for alternation behavior. In these approaches, differences in meaning between alternants are ascribed to distinct but related meanings being associated with the verbs participating in the alternation. Each sense then licenses one of the alternants and imposes different semantic conditions on the syntactic arguments appearing in the relevant alternant.

Romain (2018: 74) additionally lists valency approaches, such as those proposed by Hampe & Schönefeld (2007), Faulhaber (2011), Herbst (2011), which are centered on the verb and focus more on form than on meaning. These accounts reject the position taken by Levin (1993) and others that meaning facets of verbs can be seen as the sole indicators of alternation participation or non-participation, and instead favor an item-based view, in which speakers have to learn for each item, i.e., each verb, whether it is compatible with any specific construction.

The focus of this thesis is not on the theoretical explanation for alternation behavior of verbs, but on predicting and modeling that behavior for selected phenomena. Therefore, a further discussion of the merits of different theoretical accounts of diathesis alternations is outside the scope of the thesis. The next section will describe how the following parts of the thesis will take different aspects of these theories into account.

## 2.3 aspects of alternation theories in practice

Parts III and IV of this thesis are concerned with practical approaches to modeling or predicting different aspects of alternation behavior. Each of these implementations relies on ideas from the theoretical accounts of verb alternations discussed in the previous section. This section provides an outlook on the ways the different theories will be taken into account in each thesis part.

Part III of this thesis focuses on modeling the behavior of verbs participating in certain verb alternations, contrasting them with verbs that do not participate in those alternations but can appear in syntactically similar constructions. This is achieved with a metagrammar that is compiled into a lexicalized tree adjoining grammar which can be used to parse input sentences and jointly derive the syntactic tree structure and a semantic representation for each sentence.

The metagrammar is built in a modular fashion, with classes at different levels of a hierarchy describing the syntactic form of basic tree fragments, grammatical functions, diathesis alternatives, constructions, and tree families. The classes defining constructions contain semantic representations in the form of frames; the values for the various attributes in these frames are determined by the lexical frames for the arguments in the relevant syntactic slots. The semantic content of alternating variants (and other, unrelated constructions) is specified as part of the description of the relevant constructions. The semantic contribution of the verb in a sentence that is being parsed is also taken into account and combined with the relevant construction (see also Chapter 4 of this thesis). In this aspect, the metagrammar implements the constructionist idea that verbs have a basic meaning, and subtle meaning differences between alternating variants are made possible by the semantic contributions from the constructions themselves.

TAG as a framework is particularly well-suited for modeling insights from Construction Grammar, since it is surface-oriented instead of assuming structure-destroying operations, it can be used to define constructions at varying levels of complexity and abstraction, and it allows partial descriptions to be reused in multiple places via inheritance and thus reflects an inheritance network of constructions (Lichte & Kallmeyer 2017: 209). For more on this, see Section 5.5. Furthermore, the metagrammar can be extended to cover other phenomena or additional variants of the constructions under discussion, such as passivized forms, nominalizations or sentential subjects. While these extensions are beyond the scope of this thesis, TAG provides all the tools necessary for such extensions (see e.g. Kroch & Joshi 1985 and Abeillé et al. 1990, and Kallmeyer et al. 2016 for the interplay of such syntactic forms with semantics).

In the metagrammar developed in this thesis, alternation participation is determined via the interplay of each verb's lexical semantics, the lexical semantics of the observed arguments, the semantic requirements of the constructions characterizing the alternation, and semantic type and attribute constraints in the frame type hierarchy. This partially reflects the argument by Levin (1993) that alternation availability is guided by meaning facets that are inherent to a verb's lexical semantics.

Part IV of this thesis focuses on the question to what extent verbs that participate in certain alternations can be identified automatically based on dependency-annotated corpus data. By selecting alternation participation as the relevant level of prediction, this task puts the focus on "pairs of paraphrases", rather than on individual constructions as favored by Goldberg (2002). The approach to alternation identification pursued in this thesis relies on examining pairs of alternating variants in order to determine whether verbs that appear in both syntactic forms participate in the relevant alternation or not.

Levin (1993: 4) connects alternation behavior to meaning facets that are inherent to individual verbs: if the meaning of a verb includes a relevant set of meaning facets, it means both variants of the corresponding alternation are available for that verb. In other words, alternation participation is seen as derivable from a set of lexical properties of verbs, and verbs can be labeled with respect to alternation participation in the form of a binary classification.

The implementation of the classification task in this thesis relies on gold data sourced from VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007). That resource takes the verb classes compiled by Levin (1993) as its point of departure and systematically adds refinements and extensions. In this aspect, the alternation identification setup developed in this thesis relies on Levin's theoretical approach.

An alternative approach based on the constructionist view could cast the task as a construction classification task instead of predicting alternation identification: given a certain verb and its attestations in a corpus, can that verb instantiate a particular construction or not? Constructions are characterized by their syntactic form and their semantic interpretation; only the syntactic form can be "read off" from dependency-annotated corpus data. The challenge then lies in distinguishing constructions from each other which share their syntactic form. Implementing this task would require a resource for gold labels with respect to each verb and construction under investigation. Examining constructions in isolation, without referring to their counterparts in the context of specific alternations, would make it more challenging to hypothesize about the semantic role being expressed in a given syntactic slot in a sentence. For more on how alternating variants are leveraged for the purpose of classifying verbs with respect to alternation participation, see Part IV of this thesis.

Overall, the different theoretical accounts of verb alternations are relevant for different aspects of the practical contributions of this thesis. The goal of the implementations developed here is to highlight aspects of alternation behavior that can pose problems for natural language processing tools, and to suggest approaches to overcoming those challenges.

So far, the variants involved in an alternation have been referred to multiple times as constructions, but no definition of constructions has been provided yet. Chapter 3 will give a brief introduction to Construction Grammar. Following that, Chapter 4 will specify the types of semantic representations that will be used to describe the semantic contributions from constructions and from verb and argument lemmas.

# 3 CONSTRUCTION GRAMMAR

## 3.1 INTRODUCTION

Since verb alternations are a phenomenon at the syntax-semantics interface, it is useful to be able to model them jointly on both levels. Whether a verb can participate in an alternation seems to be impacted by the extent to which its semantic content is compatible with the different syntactic realizations associated with the alternation. In Construction Grammar (Goldberg 1995, 2013), sentence patterns themselves are regarded as having their own semantic content that is instantiated and extended by the specific verbs that appear in them. Constructions are "conventional, learned form-function pairings at varying levels of complexity and abstraction" (Goldberg 2013: 17). The sentence patterns that characterize the verb alternations under investigation in this thesis are called argument structure constructions. However, constructions also exist on other levels, including words, morphemes, or idioms. Goldberg argues for argument structure constructions as a special subclass of constructions that provides the basic means of clausal expression in a language.

This chapter will summarize some key ideas of Construction Grammar according to Goldberg, Goldberg, with a focus on the interaction between verbs and argument structure constructions. In the context of modeling alternation behavior, as shown in Part III of this thesis, the mechanism that allows verbs to be integrated into constructions is of particular interest. Wherever the constructions involved in an alternation have distinct meanings, it is essential to be able to predict how the meaning of alternating verbs will combine with the meaning of the alternating variants.

The point of departure for Goldberg is the observation that pairs of sentences instantiating similar, but slightly different syntactic patterns around the same verb are typically associated with differences in meaning. She discusses this with respect to several example sentence pairs that are realizations of English verb alternations, where the meaning differences take the form of different selectional preferences, different inferences that are being triggered, or other semantic requirements that apply to one of the sentences but not the other. While a lexicosemantic approach (e.g., Levin 1985: 6, Levin & Rapoport 1988: 281, Gropen et al. 1989: 241) would attribute these differences to different senses of the verb that are involved in the different patterns, the constructionist perspective instead assumes one central verb sense that then interacts with the semantics of the syntactic patterns individually to form a combined meaning. Thus, the semantic differences can be explained by the alternating sentence patterns having slightly different semantics that are combined with the semantics of the verb.

One set of example sentences from Goldberg (1995: 2), quoted from Anderson (1971: 394), is given in (13). The sentences exemplify the spray-load alternation (or locative alternation), in which the roles of THEME and GOAL can each be expressed either in the direct object position or as a prepositional object.

(13) a. I loaded the hay onto the truck.
     b. I loaded the truck with the hay.

In the typical interpretation of such sentences, (13b) implies that the truck will be covered in hay as a result of the loading event, whereas (13a) has no such implication and only denotes that some non-zero amount of hay is located on the truck as a result of the loading. In later work, Goldberg (2002: 340) presents arguments in favor of generating forms like (13b) directly instead of deriving them from forms like (13a): other instances of the locative construction have no alternative form in this pattern, as shown in (14) (taken from Goldberg 2002: 338-339).

(14) a. They covered the wall with posters.
b. * They covered posters onto the wall.
c. Pat adorned the tree with lights.
d. * Pat adorned lights onto the tree.

It would be possible to associate the locative alternation with a particular set of verbs, and assume that the different entailments of sentences like (13a) and (13b) are due to different senses of these locative-alternating verbs, one of which is compatible with each construction of the alternation. However, beyond such pairs of "rough paraphrases" (Goldberg 2002: 327), verbs can also appear in a wider variety of argument structure constructions, in which meaning differences would be difficult to ascribe exclusively to the verbs themselves. Some examples from Goldberg (1995: 11) are shown in (15).

(15) a. Pat kicked the wall.
b. Pat kicked Bob black and blue.
c. Pat kicked the football into the stadium.
d. Pat kicked at the football.
e. Pat kicked his foot against the chair.
f. Pat kicked Bob the football.
g. The horse kicks.
h. Pat kicked his way out of the operating room.

In the constructionist perspective, each of the occurrences of *load* in (13) and each occurrence of *kick* in (15) is based on the same central sense of the respective verb, and the different sentence patterns each provide their own meaning contributions that are instantiated and extended by the verb and its argument(s). This ultimately results in the different interpretations of *kick* in the example sentences in (15). Note, for example, that sentence (15f) expresses a transfer meaning in which the *ball* is being moved from *Pat* to *Bob* by way of a *kicking* event; Goldberg assumes that this transfer meaning is not contributed by the lexical meaning of *kick* (since it is missing from the other examples listed in (15)), but by the argument structure construction itself.

Goldberg defines constructions as form-meaning pairs (or, in her later work, as conventional, learned form-function pairs, Goldberg 2013: 17) that are the basic unit in language, and whose semantics cannot be further decomposed. For each distinct construction, one or more of its properties cannot be predicted based on its component parts or knowledge from other constructions. In that sense, constructions are idiosyncratic, which warrants viewing them as part of the lexicon, instead of separating syntax from the lexicon completely. The formal definition for constructions given by Goldberg (1995) is as follows:

> C is a CONSTRUCTION iff$_{def}$ C is a form-meaning pair $\langle F_i, S_i \rangle$ such that some aspect of $F_i$ or some aspect of $S_i$ is not strictly predictable from C's component parts or from other previously established constructions. (Goldberg 1995: 4)

While the main focus of Goldberg is on single-clause sentence patterns, she also points out that a definition requiring constructions to be form-meaning pairs also allows morphemes to be viewed as constructions, since their properties are also idiosyncratic and they can also interact with verbs to yield a combined interpretation. Constructions exist at various levels of complexity and abstraction, including words, partially-filled words (morphemes), idioms, and argument structure constructions (Goldberg 2013: 17). Some alternations that are characterized in English by different sentence patterns are realized in other languages with morphemes indicating, for instance, causativity or passives (Goldberg 1995: 22).

Like verbs and other lexical items, constructions can also exhibit polysemy. Goldberg (1995: 4) discusses a number of English argument structure constructions in detail and shows how they can form a "family of distinct but related senses", in which certain generalizations and semantic constraints apply. For instance, argument structure constructions may be compatible only with specific verb classes, or require some specific relationship between the event type of the verb and the event type associated with the semantics of the construction itself. Additionally, she assumes that constructions are interrelated and hierarchically organized in a network, a constructicon (Jurafsky 1992: 28). Constructions are viewed as part of speakers' competence or knowledge of language. The meanings of simple clause constructions are direct reflections of scenes that are basic and relevant to human experience.

Concerning the integration of verbs into sentence patterns, Construction Grammar differs from unification-based grammars like Lexical-Functional Grammar (Kaplan & Bresnan 1982), Generalized Phrase Structure Grammar (Gazdar et al. 1985), and Head-Driven Phrase Structure Grammar (Pollard & Sag 1994). These grammar frameworks rely on the assumption that the verb is the semantic head of the sentence, and its semantic features percolate upwards to determine the semantic features of the sentence (Goldberg 1995: 14). However, as Goldberg (1995: 16) argues, there can be conflicts between the semantic requirements of a verb and those of an argument structure construction, which cannot be modeled well in a framework that does not make explicit the semantic contributions or requirements from the construction. In Construction Grammar, the (lexically specified) semantic features of the argument structure construction itself are also taken into account in order to determine the overall interpretation of an instantiated construction with a particular verb. This is how the different sentence meanings come about in the example sentences in (15).

In the remainder of this thesis, when semantic contributions from verbs and argument structure constructions are discussed, the argument structure constructions will occasionally be refered to simply as constructions, even though following Goldberg's definition, verbs also have the status of learned form-function pairings and thus actually also qualify as constructions. Contrasting lexical and constructional contributions in this way is in line with related work such as Lichte & Petitjean (2015: 187) speaking about "the well-established distinction between lexical and constructional contributions to the overall meaning" or Kallmeyer & Osswald (2013: 267) speaking about "examples of the interaction between lexical and constructional meaning". In the implementation presented in Part III of this thesis, the lexical semantics of verbs will be stored in a lemma lexicon, while the semantics of unfilled constructions will be stored alongside their syntactic structure in a separate resource.

## 3.2 THE SEMANTICS OF CONSTRUCTIONS

In order to usefully analyze the semantic interactions between a construction and a verb that instantiates it, a semantic framework providing a suitable level of detail is necessary. Goldberg (1995: 29) notes that Construction Grammar grew out of Frame Semantics and an experientially based approach to language (Fillmore 1975, 1977, 1982, 1985), and that frameFrame Semantics in the style of Fillmore (1982) is an appropriate approach to semantics in this context because it also focuses on speaker-centered "construals" of situations in the sense of Langacker (1987: 128, 1991: 93). Rich frame-semantic knowledge is necessary because it provides the details needed to draw inferences, impose selectional preferences, or determine whether the event type denoted by a verb is compatible with the event type denoted by a construction.

Goldberg (1995: 28) contrasts frame-semantic representations of verb meaning with semantic decompositional structures like X CAUSES Y TO RECEIVE Z. These representations lack the level of detail that is needed to characterize verb meaning in the context of Construction Grammar, and are argued by others (e.g., Lakoff 1966: VIII-6, Levin 1985: 56, Pinker 2013: 205) to instead only reflect the "syntactically relevant aspects of verb meaning". Goldberg points out that the "syntactically relevant aspects of verb meaning" are more appropriately regarded as properties of argument structure constructions, which are independently motivated. Thus, semantic decompositional structures of that type are used by Goldberg to describe constructional meaning, whereas the meaning of verbs is expressed with frame-semantic representations that encode world knowledge and the relevant concepts at a higher level of detail.

Fillmore (1975, 1977, 1982) views frames as conceptual representations of scenes, which he defines as "a coherent individuatable perception, memory, experience, action, or object" (Fillmore 1977: 84). Meanings are defined relative to a particular background frame, whose participants or substructures can be brought to the foreground or background depending on whether attention is being directed at them. Substructures are made more prominent by profiling (Langacker 1987: 183, Langacker 1991: 331), which will be discussed in the following section.

## 3.3 INTEGRATING VERBS INTO CONSTRUCTIONS

Goldberg (1995: 43) proceeds to describe how the (more general) meanings associated with constructions are made more concrete by the (more specific) meanings of the verbs that appear in them. The relevant syntactic constituents in a construction express argument roles, such as `agent`, `patient` or `goal`, while the arguments of a verb express frame-specific participant roles, such as `hitter` or `buyer` (Goldberg 1995: 43, Goldberg 2002: 342).[1] When an argument structure construction is instantiated with a specific

1 In Goldberg's notation, argument roles and participant roles are represented in normal font, with profiled roles appearing in bold-face normal font. This thesis will represent these roles with `monospace font`, or **`bold-face monospace font`** for profiled argument roles.

verb, that verb's argument roles function as instances of the more general argument roles. Selectional restrictions apply on both levels.

In order for a participant role to instantiate an argument role, the two must be compatible with each other. Each verb can profile certain participant roles that belong to the frame that the verb evokes. Profiled roles are those that are normally obligatorily expressed in finite clauses. Profiling means that the role functions as a focal point within the frame and that it has a certain degree of prominence or salience. These profiled roles are a crucial factor in determining whether the verb is compatible with a construction: Goldberg (1995: 48) explains that profiled participant roles must always be associated with argument roles that are realized as direct grammatical functions in a construction.

Goldberg illustrates this with example sentences involving the two verbs *steal* and *rob*, both of which are based on a shared background frame of unlawful taking of `goods` from a `target` by a `thief`, but differ in the roles they profile. The examples from Goldberg (1995: 45) are given in (16) and (17).

(16) a. Jesse robbed the rich (of all their money).
  b. * Jesse robbed a million dollars (from the rich).

(17) a. Jesse stole money (from the rich).
  b. * Jesse stole the rich (of money).

While *rob* can appear with the `target` as the direct object, it is infelicitous with the `goods` in that position. The reverse is true for *steal*. Goldberg's explanation for the difference is that *rob* is lexically determined to profile the roles **`thief`** and **`target`**, while the verb *steal* is lexically determined to profile the roles **`thief`** and **`goods`**. The respective non-profiled participant roles can be expressed in adjunct phrases, but are not felicitous in positions with direct grammatical relations to the verb.

Concerning constructional profiling, Goldberg states the following:

> Every argument role linked to a direct grammatical relation (SUBJ, OBJ, or $OBJ_2$) is constructionally profiled. (Goldberg 1995: 48)

For the ditransitive construction, the three argument roles associated with the subject, direct object, and indirect object positions are profiled. The construction is associated with the meaning X CAUSES Y TO RECEIVE Z. Verbs are compatible with this construction if they are lexically specified to have three profiled participant roles that are compatible with the (more general) argument roles that are profiled in the construction. Compatibility essentially boils down to the question whether each of the participant roles provided by the verb can instantiate one of the argument roles provided by the construction. The treatment of cases without a clear one-to-one association between profiled argument roles and profiled participant roles will be discussed in the following.

The profiled roles of constructions and verbs are represented by Goldberg (1995: 51) as exemplified in (18). Each profiled role is typeset in bold font.

(18) a. Ditransitive construction:
  CAUSE-RECEIVE < **`agent`** **`recipient`** **`patient`** >
  b. The verb *hand*:
  HAND < **`hander`** **`handee`** **`handed`** >

| | | | | | | |
|---|---|---|---|---|---|---|
| Sem | CAUSE-RECEIVE | < | **agt** | **rec** | **pat** | > |
| | R | | | | | |
| R: instance, means | HAND | < | **hander** | **handee** | **handed** | > |
| Syn | V | | SUBJ | OBJ | OBJ2 | |

Figure 1: Combined representation of the ditransitive construction and the verb *hand* (from Goldberg 1995: 51).

### 3.3.1 *Fusing participant roles and argument roles*

Goldberg (Goldberg 1995: 50,Goldberg 2002: 342) uses the term *fusion*, coined by Jackendoff (1990: 50), to describe how the argument roles from the construction and the participant roles from the verb are unified. For the construction and the verb shown in (18), the meaning of the construction is "entirely redundant with the verb's meaning", and the verb further specifies the event denoted by the construction, so that there is a one-to-one-correspondence between the roles. The combined structure as represented by Goldberg (1995: 51) is reproduced here as Figure 1.

The representation format in Figure 1 and all following construction figures in this section is not meant to be a full formalization by Goldberg (1995). As she says in later work, she has "avoided using all but the most minimal formalization in [her] own work" (Goldberg 2013: 29) but finds feature-value matrices to be the most complete and well-worked out formalism to represent constructions. The function of the figures from Goldberg (1995) reproduced in this section of the present work is to visualize the fusion of participant roles and argument roles.

In Figure 1, the semantic contribution of the construction is expressed in the top row. The bottom row denotes the syntactic constituents that are involved in the construction. Each argument role is connected with an arrow to a grammatical relation; for instance, the **agent** role is connected to the syntactic SUBJ position. In the middle row, the contribution from the verb is expressed. The relation between the construction and the predicate denoted by the verb that is being integrated into it is referred to as R; in the example, R is an "instance, means" relation. Thus, the CAUSE-RECEIVE meaning contributed by the construction is instantiated by the HAND meaning contributed by the verb *hand*. Furthermore, the participant roles provided by the verb are inserted into the construction such that each participant role is directly connected to an argument role. In this case, all three profiled argument roles are instantiated by profiled participant roles. Whether a particular participant role is compatible with a particular argument role is determined by their meaning; for instance, the **recipient** role is expected to be animate, but not agentive, which renders the **handee** participant role the most likely one to instantiate it, out of all available participant roles from the verb (since the **hander** role is expected to be agentive, and the **handed** role is not agentive and not necessarily animate).

The fusion mechanism also covers scenarios without such a one-to-one relationship between profiled argument roles and profiled participant roles. The main principles guiding the fusion of argument roles and participant roles are the semantic coherence

| | | | | | | |
|---|---|---|---|---|---|---|
| Sem | CAUSE-MOVE | < | **cause** | goal | **theme** | > |
| | PUT | < | **putter** | **put.place** | **puttee** | > |
| Syn | V | | SUBJ | OBL | OBJ | |

Figure 2: Combined representation of the caused-motion construction and the verb *put* (from Goldberg 1995: 52).

| | | | | | | |
|---|---|---|---|---|---|---|
| Sem | CAUSE-RECEIVE | < | **agt** | **rec** | **pat** | > |
| | R | | | | | |
| R: instance | MAIL | < | **mailer** | mailee | **mailed** | > |
| Syn | V | | SUBJ | OBJ | OBJ2 | |

Figure 3: Combined representation of the ditransitive construction and the verb *mail* (from Goldberg 1995: 53).

principle and the correspondence principle (Goldberg 1995: 50, Goldberg 2002: 342). The semantic coherence principle requires two roles to be semantically compatible in order to be able to be fused. This is the case if either one of the roles can be construed as an instance of the other. This is determined by general categorization principles. The correspondence principle requires each profiled participant role that is expressed to be fused with a profiled argument role of the construction, but admits the exception of cases where the verb has three profiled participant roles, and the construction has a nonprofiled argument role with which one of the participant roles can be fused.

The latter scenario is illustrated by Goldberg (1995: 52) with the verb *put* instantiating the caused-motion construction. The caused-motion construction is represented as CAUSE-MOVE <**`cause`** `goal` **`theme`**>, where the roles **`cause`** and **`theme`** are profiled because they are linked to the grammatical functions SUBJ and OBJ, but the role `goal` is not profiled because it is linked to an oblique function. Nevertheless, the verb *put*, which is represented as PUT <**`putter`** **`put.place`** **`puttee`**>, can appear in this construction. The **`putter`** and **`puttee`** participant roles can easily be fused with the **`cause`** and **`theme`** argument roles, respectively. In addition, the **`put.place`** role can fuse with the non-profiled `goal` argument role because the **`put.place`** role is a type of goal, which means that they are semantically compatible. Goldberg's combined representation of the construction with the verb is shown in Figure 2.

Verbs can also appear in constructions if the construction has one more profiled argument role than the verb has profiled participant roles. This is the case in another example provided by Goldberg (1995: 53). In that example, the verb *mail*, which is represented as MAIL <**`mailer`** **`mailed`** `mailee`>, is integrated into the caused-motion construction, represented as CAUSE-RECEIVE <**`agent`** **`recipient`** **`patient`**>. This is unproblematic because the profiled status of an argument role can generally be passed on to a nonprofiled participant role that is fused with it. The combination of the construction and the verb is shown in Figure 3.

| | | | | | | |
|---|---|---|---|---|---|---|
| Sem | CAUSE-RECEIVE | < | **agt** | **rec** | **pat** | > |
| | R | | | | | |
| R: means | KICK | < | **kicker** | | **kicked** | > |
| | ↓ | | ↓ | ↓ | ↓ | |
| Syn | V | | SUBJ | OBJ | OBJ2 | |

Figure 4: Combined representation of the ditransitive construction and the verb *kick* (from Goldberg 1995: 54).

Profiled argument roles thus do not obligatorily need to fuse with profiled participant roles. This is because constructions can contribute roles that are not associated with a participant role of verbs that appear in them. Goldberg (1995: 54) illustrates this with a sentence like (15f) *Pat kicked Bob the football*, where the verb *kick*, represented as KICK <**kicker kicked**>, is integrated into the ditransitive construction, represented as CAUSE-RECEIVE <**agent recipient patient**>. The semantic coherence principle determines which roles of CAUSE-RECEIVE can be instantiated by which roles of KICK. None of the participant roles of *kick* can instantiate a **recipient**. This role is thus exclusively contributed by the construction in this scenario, and no sense of *kick* with an obligatory `recipient` needs to be assumed to be part of the lexicon. The combined representation of the construction and the verb is shown in Figure 4. The integration of this verb into this construction is what licenses sentences like (15f) and makes them interpretable.

Fusing argument roles with participant roles allows constructions to impose selectional requirements on the arguments of the verb. This is illustrated by Goldberg (1995: 55) with an example of the verb *send*, represented as SEND <**sender** `send.goal` **sent**>, integrating into the ditransitive construction. Since the `send.goal` participant role fuses with the **recipient** argument role, an animacy requirement is imposed on the entity filling that role. In sentences like *John sent Chicago a letter*, only a metonymical reading is possible in which *Chicago* refers to certain people in Chicago, in order to fulfill the animacy requirement. Because the animacy requirement is not contributed by the verb itself, alternative sentences like *John sent a letter to Chicago* can be interpreted without the metonymical extension, so that *Chicago* can refer to a place instead.

This analysis reflects the claims about the dative alternation that appeared in Chapter 1 of this thesis. The caused possession meaning ascribed to the ditransitive variant may be linked to the presence of the `recipient` role provided by that construction; the prepositional object variant does not (necessarily) involve an animate recipient who is being caused to possess something, so that this variant is associated with a caused motion meaning instead.

Polysemy can cause verbs to have different sets of profiled participant roles, depending on the sense of the verb. As a result, the same argument roles in a construction can be fused with different participant roles. An adapted example from Goldberg (1995: 56) is given in (19). The argument role associated with the SUBJ position is fused either with the **tenant** participant role of the verb *lease*, as in (19a), or the **landlord** participant role, as in (19b).

(19) a. Cecile leased$_1$ the apartment from Ernest. (**tenant** is profiled)
b. Ernest leased$_2$ the apartment to Cecile. (**landlord** is profiled)

The different senses of *lease* are represented by Goldberg (1995: 56) as shown in (20).

(20) a. lease$_1$ <**tenant property** landlord>
b. lease$_2$ <tenant **property landlord**>

In this case, Goldberg acknowledges that the different constructions that are available to the two senses of *lease* (and similarly to other verbs, like *rent*) may constitute an alternation that is due to real lexical polysemy of the verb. She ascribes this to the different sets of profiled participant roles for the different verb senses, which is not applicable to verbs like *kick* appearing in constructions like the ditransitive construction. Levin (1993) categorizes verbs like *lease* and *rent* as participating in the alternations illustrated by the example sentences in (21), which are generated to resemble the examples from Goldberg above, with annotations added to distinguish between the two relevant verb senses.

(21) a. Dative alternation (Levin 1993: 45):
i. Ernest leased$_2$ the apartment to Cecile.
ii. Ernest leased$_2$ Cecile the apartment.
b. Benefactive alternation (Levin 1993: 48):
i. Dora leased$_1$ the apartment for Cecile.
ii. Dora leased$_1$ Cecile the apartment.
c. Sum of money subject alternation (Levin 1993: 83):
i. Cecile leased$_1$ the apartment for $2000.
ii. $2000 leased$_1$ Cecile the apartment.

Levin lists *lease* as a GIVE verb in the context of the dative alternation, a GET verb in the context of the benefactive alternation, and a VERB OF OBTAINING in the context of the sum of money subject alternation.[2] These different categories reflect the different verb senses mentioned by Goldberg. Note that Levin does not posit a *to/from* alternation resembling the contrast shown in (19) above.

It is also possible for constructions to prevent profiled participant roles of a verb from being expressed. The three mechanisms that can cause this are called shading, cutting, and merging (Goldberg 1995: 56). Shading refers to a process that "deprofiles" roles, which may then be expressed by an adjunct; an example for this is the passive construction, which deprofiles a participant role that typically fuses with an **agent** argument role. Cutting is a similar process, but in contrast to shading, the deprofiled role cannot be expressed. An example is the impersonal passive construction in German, which licenses sentences with a deprofiled agent like *Es wird im Treppenhaus geputzt* (roughly: *There is cleaning being done in the staircase*). Finally, merging allows two participant roles to fuse with a single argument role and be linked to the same grammatical function. An example for this are reflexive constructions in Romance languages.

2 For the purposes of this example, the VERB OF OBTAINING sense of *lease* is grouped with the GET sense of the verb.

Verbs can also be lexically specified to allow indefinite or definite null complements in certain participant roles. Indefinite null complements are roles that may be left unexpressed because they do not receive prominence, and whose referents' identities are irrelevant in the given context. This is the case, for instance, for optionally intransitive verbs like *eat*. Definite null complements must be given in discourse in order to be left unexpressed.

Goldberg (2002: 345–346) notes that roles that are profiled neither by the construction nor by the verb behave like traditional adjuncts, while the behavior of roles that are profiled exclusively by the verb falls "somewhere in between that of traditional arguments and traditional adjuncts". She contrasts sentences like *She loaded the wagon with hay* with sentences like *She broke the window with a hammer*, and argues that the *with* phrase in these sentences has different functions: the phrase would normally be regarded as an adjunct, but in the *load* sentence, it expresses the profiled role **`loaded-theme`** of the verb. Because no such profiled role is associated with *break*, the sentences behave differently when additional adjunct phrases are inserted before the *with* phrase. Sentence (21b-i) above falls in the same category as the *load* sentence, since the profiled participant role **`tenant`** of *lease*$_1$ is expressed in the form of a prepositional object.

### 3.3.2 *Semantic conditions on fusion of verbs and constructions*

Goldberg (1995: 60) examines the conditions that must be met in order for a verb to be felicitously integrated into a construction. Beyond the number of profiled participants and the satisfaction of the principle of semantic coherence and the correspondence principle, which relate to the fusion of argument roles with participant roles, the semantics of the verb and the construction itself must also be compatible.

In the trivial scenario, the basic sense of the verb simply instantiates the semantics of the construction, as is the case for verbs like *hand* in combination with the ditransitive construction, or verbs like *put* in the caused-motion construction.

Other relationships between the meanings of constructions and the meanings of verbs are discussed under the term *conflation patterns* (Talmy 1985: 60), which denote different types of semantic mismatches between the two senses. They allow verbs to instantiate constructions even when no direct instantiation relation exists between them.

Verbs can be compatible with a particular construction if the verb expresses the means by which the event or situation denoted by the construction comes about. Goldberg (1995: 60) illustrates this with the verb *kick* appearing in the ditransitive construction. The basic meaning of the construction denotes a successful transfer of an object to a recipient, and the basic sense of the verb denotes the means by which that transfer is being conducted.

Alternatively, verbs can express the result of an event denoted by the construction. This is the case for verbs integrating into causative constructions. Such constructions are then represented by Goldberg as CAUSE <`agent patient PRED` < `...`>>, where the specifics of `PRED` < `...`> are contributed by the semantics of the verb. This type of construction will be discussed in more detail in Chapter 4 of this thesis. The present work does not assume that causation events always involve an agent and a patient, and will instead focus on the components CAUSE and EFFECT, both of which are events that may or may not involve agents and patients.

A more general perspective on these relations is formulated by Goldberg as follows:

> *Causal Relation Hypothesis*: The meaning designated by the verb and the meaning designated by the construction must be integrated via a (temporally contiguous) causal relationship. (Goldberg 1995: 62)

This principle can also be violated by certain constructions. For instance, the *way*-construction denotes a motion along a path and can be instantiated by verbs that do not denote the means of motion, but merely the manner in which the motion is performed. For instance, the verb *knit* can be integrated into the construction without the *knitting* event contributing to or causing the motion denoted by the construction, as in (22) (reproduced from Goldberg 1995: 62).

(22) I knitted my way across the Atlantic.

An extension of the causal relation hypothesis is necessary to account for scenarios where the verb denotes a necessary precondition for the event denoted by the construction. This is the case when a verb of creation, like *bake*, instantiates the ditransitive construction. The creation event denoted by *bake* does not cause and is not caused by the transfer event denoted by the construction, but an item must first be created before it can be transferred. A possible representation of the event denoted by this combination of construction and verb will be shown in Chapter 4.

Goldberg summarizes the relationship between the semantics of constructions and the semantics of verbs instantiating them as follows:

> The semantics associated with the construction defines a semantic frame, and the verb must inherently designate a particular salient aspect of that frame. (Goldberg 1995: 65)

This thesis will operate on this assumption about constructions and verbs and the frames they evoke. Alternations will be regarded as sets of argument structure constructions with distinct but related semantics, into which the basic sense of participating verbs is integrated. The interpretation of a sentence that instantiates a construction is derived from the composition of the semantics of the construction and the semantics of the observed verb. This process is no different for alternation-specific constructions than for any other construction that can be instantiated by a verb.

### 3.3.3 *Semantic conditions involving arguments of the verb*

Beyond these considerations revolving around the semantics of constructions and the semantics of verbs, Montemagni (1994: 352) points out that (alternation-specific) constructions are also sensitive to the semantics of the observed arguments of the verb. There are cases in which a verb is generally assumed to participate in an alternation, but one of the alternants requires a specific type of verb argument. For arguments that are incompatible with the required type, the verb cannot appear in all alternation-specific constructions.

An example given by Montemagni is the verb *gather* and its ability to participate in the causative-inchoative alternation. The examples in (23) and (24) are adapted from Montemagni (1994: 352).

(23) a. John gathered his friends.
b. John gathered the animals.
c. John gathered his papers.
d. John gathered his maps.
e. John gathered berries.

(24) a. His friends gathered.
b. The animals gathered.
c. * His papers gathered.
d. * His maps gathered.
e. * Berries gathered.

Montemagni (1994: 352–353) argues that the difference between the acceptable and unacceptable sentences in (24) is due to the semantics of the different verbal arguments. She suggests distinguishing "volitional" themes, as in (24a) and (24b), from "non-volitional" ones, as in the remaining examples. A similar distinction holds between possible themes for verbs like *begin*; in this case, only arguments of the type *event* license the alternation. However, Montemagni (1994: 351) argues against making these distinctions on the level of thematic roles, as this would lead to a larger and larger inventory of thematic roles as more and more semantic conditions for particular verbs and their arguments are being added.

Instead, she proposes a solution involving the specification of selectional restrictions for all syntactically expressed arguments of the verbs under discussion. In all examples shown above, the thematic role assigned to the relevant verb argument is the THEME role. Her specification for *gather* is reproduced in (25); (25a) describes the alternating argument structure for the verb, while (25b) describes the non-alternating argument structure.

(25) a. gather(SUBJ, OBJ<+animate, +volitional>)
agent theme
b. gather(SUBJ, OBJ<-animate, -volitional>)
agent theme

The specification makes it possible to distinguish between uses of *gather* with animate, volitional arguments, like *his friends*, and uses of the same verb with inanimate, non-volitional arguments, like *berries*. The distinction then makes it possible to specify in the lexicon or grammar which verb uses license the alternation and which ones do not. Montemagni (1994: 353–354) does not claim that this solution is satisfactory, she is mainly concerned with illustrating the problem of taking verb argument properties into account when determining alternation availability.

Levin (1993) does not list *gather* as a verb participating in the causative-inchoative alternation. Beyond the restrictions on the verb's arguments hypothesized by Montemagni, another reason for the exclusion of this verb from the alternation may be that the inchoative construction typically expresses *change of state* events, while *gathering* events seem to be a subtype of *change of location* instead of *change of state*.

Part III of this thesis presents a metagrammar implemented in the XMG framework covering a number of verb alternations. This framework makes it possible to define constraints to predict the behavior of verbs like *gather*, as observed by Montemagni. In order to model these verbs, one could include the different event types *change of location* and *change of state* in the frame type hierarchy and enable the inchoative construction to be instantiated by verbs whose lexicon entry is compatible with *change of state* (like *dry*), while another construction with a name like *volitional translocation construction* can be instantiated by verbs whose lexicon entry is compatible with *change of location* (like *gather*). A causative construction for sentences like *Ahab dried the harpoon* or *John gathered his friends* can then be instantiated by both *dry*-type verbs and *gather*-type verbs. Aspects of meaning and structure that are shared between the inchoative construction and the volitional translocation construction can be defined once and be inherited by both constructions.

## 3.4 CONCLUSION

This thesis relies on constructions in the sense of Goldberg (1995, 2013) as the relevant level of analysis for modeling and predicting the alternation behavior of verbs. Each alternation is characterized by the set of argument structure constructions that are associated with it. Verbs can participate in an alternation if they can successfully be "fused" with each construction belonging to the alternation. This is the case if the verb and the construction are semantically compatible in a way that allows the verb's participant roles and the construction's argument roles to be mapped to each other.

The semantic representations used for verbs and constructions in this thesis will not follow the format used by Goldberg (1995). Instead, frames will be represented as recursive, typed feature structures in the style used by Osswald & Van Valin (2014), Kallmeyer & Osswald (2013), Lichte & Petitjean (2015), Kallmeyer et al. (2016), Seyffarth (2019b) and others. The meaning of a construction that is instantiated by a particular verb results from the combination of the partial frame provided by the construction with the partial frame provided by the verb. In the case of type incompatibilities, the unification fails. This view on frame semantics will be elaborated upon in Chapter 4.

Part III of this thesis will be concerned with modeling the interactions between constructions and verbs in the context of verb alternations. Particular attention will be dedicated to cases in which constructions share a syntactic form, but are semantically characterized by different frames. Part IV discusses ways to identify whether a verb participates in an alternation or not, given that the verb can be observed in all relevant syntactic configurations that are associated with the alternation. Here, the challenge is to determine whether the observed instances of the verb in a syntactic pattern do in fact instantiate the expected, alternation-specific construction.

# 4 FRAME SEMANTICS

## 4.1 INTRODUCTION

In the constructionist perspective, the interpretation of a linguistic utterance results from interactions between the semantic information provided by the instantiated constructions and the semantic information provided by the lexical material that appears in the constructions. Goldberg (1995, 2002, 2013) argues that construction knowledge is part of the lexicon due to the idiosyncratic nature of constructions, and that constructions are part of speakers' linguistic competence. She uses semantic decompositional structures to represent the meaning of constructions, and semantic frames in the style of Fillmore (1982) to represent the more specific meaning of verbs and the roles of their arguments. This is justified by Goldberg (1995: 40) with the notion that both constructions and verbs are associated with meanings that constitute basic scenes that are relevant to human experience. Semantic frames denoting verb meanings are particularly well-suited to represent speaker-centered construals of situations, as well as to model inferences, selectional preferences, and event type compatibility.

The proposal by Goldberg (1995: 65) assumes that "the semantics associated with the construction defines a semantic frame, and the verb must inherently designate a particular salient aspect of that frame". Under this perspective, it is justified to represent both constructional meaning and verbal meaning in the form of semantic frames. The frames can represent different levels of detail. Goldberg (1995: 43) notes that the argument roles contributed by constructions are typically more general than the participant roles contributed by the verb; in a frame representation, this can be reflected by using types that are located at different levels of specificity in a type inheritance hierarchy.

This chapter of the thesis is concerned with motivating and specifying the style of semantic frames that will be employed for the TAG implementation modeling certain verb alternations and related constructions in Part III. The frames need to be structured and detailed enough to reflect the phenomena under investigation.

Goldberg's point of departure, frames in the style of Fillmore, are not deemed sufficient for these purposes by Osswald & Van Valin (2014), who highlight the benefits of defining and representing semantic frames at a greater level of detail and with more internal structure.

A formalization of frames as typed, recursive feature structures is proposed by Petersen (2015). These frames are a step towards a decompositional representation and analysis of predicates and event structure, but do not include all aspects proposed by Osswald & Van Valin. For instance, the formalization by Petersen does not allow frame descriptions to contain non-functional relations or feature structures with multiple base nodes.

A formalization of semantic frames by Kallmeyer & Osswald (2013) introduces additional aspects like base labels and non-functional relations. This is the formalization that is implemented in the frame compiler for the XMG framework (Lichte & Petitjean 2015: 187), which will be used for the metagrammars developed in Part III of this thesis.

The following sections will briefly summarize the main properties of the semantic frames assumed by Osswald & Van Valin, Petersen, and Kallmeyer & Osswald, respectively.

## 4.2 OSSWALD AND VAN VALIN'S ARGUMENT FOR FRAMES AS TYPED, RECURSIVE FEATURE STRUCTURES

Osswald & Van Valin (2014) examine the extent to which the representation of frames in FrameNet (Fillmore, Johnson & Petruck 2003, Fillmore & Baker 2009) is sufficient to model verb meanings, as the stated goal of FrameNet is to create a hierarchical network of interrelated frames in order to facilitate the development of an empirically grounded theory of the syntax-semantics interface. In particular, Osswald & Van Valin are concerned with the amount of internal structure of frames that is represented and with the relations between frames in FrameNet. They argue (Osswald & Van Valin 2014: 139–140) that it is beneficial to represent frames in a way that makes their internal structure accessible and uses that internal structure to encode relationships between frames, which is not the case in FrameNet.

FrameNet is an implementation of Frame Semantics in the style proposed by Fillmore (1982), where frames constitute schematic cognitive structures that represent speakers' knowledge about the world (see e.g. Fillmore 1982: 112), and word senses are characterized by such frames. Each frame that appears in FrameNet is documented in terms of a frame name, an informal description of situations that are described by the frame, a set of lexical units that can instantiate the frame, and the mandatory core roles and optional non-core roles that describe participants in events described by the frame. FrameNet organizes frames in a hierarchy based on eight different relations (Fillmore & Baker 2009: 330), but Osswald & Van Valin (2014: 130) criticize that there is no "top-down" strategy guiding the design of the frame hierarchy and the decision whether a particular relation exists between two frames or not.

According to Osswald & Van Valin (2014: 137), the structure of the FrameNet frame hierarchy would benefit from a more decompositional approach to frame semantics, which would be able to explicitly reflect the internal structure of an event or a state of affairs. This is in contrast to the current approach, which simply labels constituents of sentences with semantic roles and does not systematically establish relations between frames.

FrameNet uses frame-specific semantic roles that describe the semantics of the different elements in each valency pattern; as Osswald & Van Valin (2014: 128) comment, FrameNet lacks a universal role inventory. Recall that Goldberg (1995: 43) draws a distinction between (more general) argument roles that characterize construction meanings, and (more specific) participant roles that characterize verb meanings. In order to model interactions between constructions and verbs, as implemented in Part III of this thesis, a more structured role inventory would be advantageous in which both universal roles and frame-specific roles are accounted for and hierarchically ordered.

Osswald & Van Valin specifically investigate the frame type *event*[1] in FrameNet, arguing that this type is not used consistently to structure the frame hierarchy. They criticize(Osswald & Van Valin 2014: 134) that *event* is characterized as a subframe of a frame called *change of state scenario*, which seems to imply that every event is necessarily a part of a change of state.

Osswald & Van Valin (2014: 134) argue that frames that inherit from the type *event* should exhibit some important common properties. They suggest that the type *event* could either be regarded as being representative of all kinds of dynamic situations, including processes and activities like motion in place and directed perception; or the *event* type could be defined in a more narrow sense, and only cover situations that involve a "conceptually salient" change that is typically manifested on one of the participants. Instead of following one of these strategies, the frames that inherit from *event* in FrameNet are very heterogeneous, with some seemingly exemplifying the view of *events* as any sort of dynamic situation, and others contradicting this view. Some frames that should inherit from *event* under the first view are not related to it at all according to the FrameNet hierarchy. The event types *process* and *activity* could be, but are not, used to systematically structure the hierarchy. Osswald & Van Valin (2014: 136) argue that there should also be clear criteria for deciding whether a new event frame that is added to FrameNet should inherit from *process*, *change of state scenario*, or neither of these. In their view, a distinction between telic and atelic events, or a more fine-grained Aktionsart distinction, should be encoded in the frame hierarchy, because these distinctions are relevant both for natural language reasoning and for discovering generalizations about the syntax-semantics interface.

For the purposes of this thesis, the representation of events and the nature of the relationship between more specific frames and this more general frame type are of particular interest. Constructions are regarded as providing very general semantic information (for instance, some frame of the type *event*), while verbs are expected to provide semantic information that is either as specific as the semantics of the construction or more specific. Thus, a reliable frame type hierarchy, including relations between event types, is necessary in order to determine whether a verb and a construction are compatible with each other.

Osswald & Van Valin (2014: 136) devote particular attention to the representation of causativity in semantic frames. The frame relation *is causative of* is used in FrameNet to connect pairs of frames denoting state changes and their causative counterparts. Verbs that evoke both frames in such a pair participate in the causative-inchoative alternation. However, the causativity relation is not used consistently in FrameNet, so that some such frame pairs are not specified to be in this relation according to the frame type hierarchy, even if their associated verbs are generally considered to participate in the causative-inchoative alternation. There is also no inheritance relation between more general causative frame types, like *cause change*, and more specific ones that are intuitively categorized as subtypes of them, like *cause change of phase*.

The improvements suggested by Osswald & Van Valin (2014: 139) for the representation of causativity in semantic frames revolve around representing frames at a greater level of detail with more internal structure. A decompositional analysis of event frames,

1 FrameNet frame names follow a convention of being spelled with an initial capital letter and underscores between words. In the interest of readability, FrameNet frame names will be spelled with lowercase letters in italics and with spaces between words in this thesis.

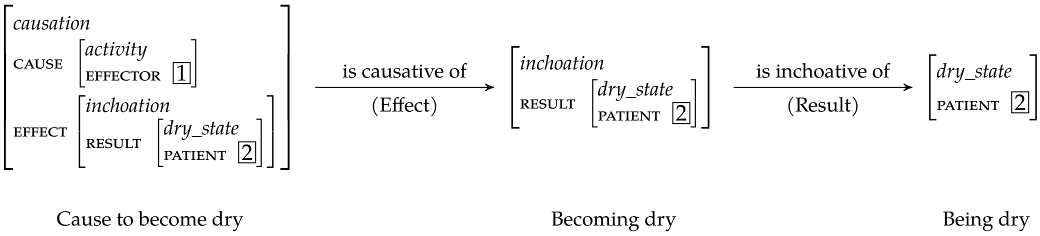


Figure 5: Frames representing *cause to become dry*, *becoming dry* and *being dry* (from Osswald & Van Valin 2014: 140).

according to them, would close some of the gaps in the frame type hierarchy and also benefit the goal of formulating linking generalizations. More specifically, with such an analysis, semantic factors that play a role in argument realization can be made more explicit, since causal and aspectual factors and other issues like volitionality can be represented explicitly. This is why Osswald & Van Valin propose a move away from the plain role frames used in FrameNet and toward a more complex frame structure, in which frames are regarded as typed, recursive feature structures that can also be embedded into each other. As a result, the proposal by Osswald & Van Valin is more promising than the plain role frames in FrameNet for the purposes of this thesis, whose main focus is on the conditions on different argument realizations in the context of verb alternations.

Embedding frames into each other can directly benefit the representation of relations that exist between frames, for instance in the context of causativity. Osswald & Van Valin (2014: 140) provide an example set of frames evoked by the verb *dry*. The verb participates in the causative-inchoative alternation, which is encoded in FrameNet by the lexical unit being able to evoke both the *cause to be dry* and *becoming dry* frames. Additionally, the verb has a zero-related adjective evoking the frame *being dry*.

The representations proposed by Osswald & Van Valin for these frames are shown in Figure 5. They argue that the frames associated with the different senses of *dry* should encode the fact that these senses are systematically related to each other, which they represent by embedding a frame for the sense *being dry* (to which they assign the frame type *dry state*) as the value of a RESULT attribute in the frame for the sense *becoming dry* (to which they assign the frame type *inchoation*), and in turn embedding the frame for *becoming dry* as the value of an EFFECT attribute into a frame for the *cause to become dry* sense (to which they assign the frame type *causation*).

While the existing FrameNet frame *cause to be dry* is in fact in an "is causative of" relation with the FrameNet frame *being dry*, the fact that the frames actually share parts of their internal structure cannot be represented with the plain role frames stored in FrameNet. Another advantage of the decompositional approach suggested by Osswald & Van Valin (2014: 140) is that the roles associated with embedded frames would transparently be accessible from the embedding frames. For instance, the PATIENT of the frame representing the *dry state* is found by following an attribute path in the *cause to be dry* frame along the attributes EFFECT – RESULT – PATIENT. The frame-specific role DRYEE can still be used as a shortcut for the value at this path.

Osswald & Van Valin (2014: 141) go on to show that decompositional frames with an internal structure can also be used to represent more details about the processes involved in events represented by frames. In FrameNet, events are viewed as having a

$$\begin{bmatrix} \textit{process} & \\ \textsc{prog} & \begin{bmatrix} \textit{incr\_change} & \\ \textsc{init} & \begin{bmatrix} \textit{stage} & \\ \textsc{entity} & \boxed{1} \\ \textsc{moisture} & \boxed{2} \end{bmatrix} \\ \textsc{result} & \begin{bmatrix} \textit{stage} & \\ \textsc{entity} & \boxed{1} \\ \textsc{moisture} & \boxed{3} \end{bmatrix} \\ \textsc{lesser} & \left( \boxed{3}, \boxed{2} \right) \end{bmatrix} \end{bmatrix}$$

Becoming drier

$$\begin{bmatrix} \textit{inchoation} & \\ \textsc{init} & \begin{bmatrix} \textit{stage} & \\ \textsc{entity} & \boxed{1} \\ \textsc{moisture} & \boxed{2} \end{bmatrix} \\ \textsc{result} & \begin{bmatrix} \textit{stage} & \\ \textsc{entity} & \boxed{1} \\ \textsc{moisture} & \boxed{3}\ \textit{zero} \end{bmatrix} \\ \textsc{lesser} & \left( \boxed{3}, \boxed{2} \right) \end{bmatrix}$$

Becoming dry

$$\begin{bmatrix} \textit{state} & \\ \textsc{cont} & \begin{bmatrix} \textit{stage} & \\ \textsc{entity} & \boxed{1} \\ \textsc{moisture} & \textit{zero} \end{bmatrix} \end{bmatrix}$$

Being dry

Figure 6: Frames representing atelic, telic, and static senses of *dry* (from Osswald & Van Valin 2014: 141).

pre-state, a central change, and a post-state, out of which the central, changing part of the event is typically in the foreground (Fillmore & Baker 2009: 331). Osswald & Van Valin argue that decompositional frames would also be advantageous for encoding the pre-state and post-state. The representation they propose for the different stages of drying events is reproduced here as Figure 6.

As shown in Figure 6, Osswald & Van Valin suggest a representation which explicitly encodes the fact that the entity undergoing a drying process, whether that process is externally caused or not, has a certain moisture value in the init(ial) stage of the process and another moisture value in the result stage, and that the final value of moisture is in a lesser relation to the initial value of the same attribute. This also allows a distinction between telic drying as a finished process, which implies a final moisture value of zero, and the atelic interpretation, in which the entity's moisture value is reduced but not necessarily brought down to a final value of zero. The different frames, titled *becoming drier* and *becoming dry*, would apply depending on contextual indicators of (a)telicity like the temporal phrases *in an hour* (denoting telicity) and *for an hour* (denoting atelicity). Analogous to the two frames denoting a change of moisture values, the frame for the state of being dry is described with a moisture attribute that continuously has the value zero.

According to Osswald & Van Valin (2014: 153), explicitly encoding these details of event frames facilitates the Aktionsart distinctions that are relevant to understanding the syntax-semantics interface and drawing correct inferences. Additionally, with this decompositional analysis of frames, it is possible to expicitly and transparently encode the relationship between an activity and the resulting state of the affected object within a single frame, or more generally, to represent subcomponents of events and attributes of participants, which is particularly relevant for modeling the syntax-semantics interface.

In a verb alternation, the different argument structure constructions that are available for a verb typically have related meanings or are near-paraphrases. Subtle meaning differences between the variants of the dative alternation and between the variants of the locative alternation have been discussed previously in this thesis. In the causative-inchoative alternation, the meaning of one construction constitutes one of the subevents of the other construction, analogous to the analysis for *dry* proposed by Osswald & Van Valin (2014: 140). Such meaning differences can be modeled with

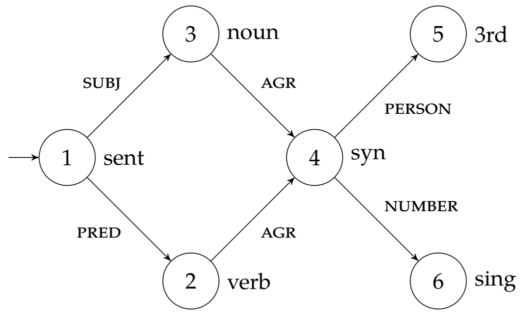


$$\begin{bmatrix} \textit{sent} & \\ \textsc{subj}: & \begin{bmatrix}\textit{noun} & \\ \textsc{agr}: & \begin{bmatrix} \boxed{4}\ \textit{syn} & \\ \textsc{person}: & [\textit{3rd}] \\ \textsc{number}: & [\textit{singular}]\end{bmatrix}\end{bmatrix} \\ \textsc{pred}: & \begin{bmatrix}\textit{verb} & \\ \textsc{agr}: & [\boxed{4}]\end{bmatrix}\end{bmatrix}$$

Figure 7: Example feature structure represented as a visual graph with nodes and arcs (top) and as an attribute-value matrix (bottom) (from Carpenter 1992: 37–38).

decompositional frames, such that the aspects of meaning in which the two variants overlap are represented identically, and only the particular meaning contributions inherent to one construction or the other are added to their respective frames.

### 4.3 Petersen's formalization of frames as typed feature structures

The notion of frames as typed, recursive feature structures as pursued by Osswald & Van Valin (2014) is largely consistent with the work of Gamerschlag et al. (2014), Petersen (2015), Löbner (2015) and others, who view frames as the universal format of representation for concepts in cognition; they assume that the human mind conceptualizes and categorizes objects and situations via functional attributes whose values themselves can be complex frames. This perspective on frames builds mainly on work by Barsalou (1992), who suggests frames as representations of concepts which are recursively composed of attributes for the object they describe and the values for these attributes. These frames can be visualized as acyclic directed labeled graphs, with circles representing nodes and arrows representing attribute arcs; an alternative representation is an attribute-value matrix notation, as used by Osswald & Van Valin (2014) and others for these types of semantic frames. In the graph, or feature structure, each node is labeled with the most specific conceptual class to which the object that is represented by the node belongs. Examples for the different notations of feature structures are shown in Figure 7 (taken from Carpenter 1992: 37–38). As Carpenter points out, the graphical notation can be difficult to typeset and understand, particularly for more complicated feature structures. Instead of the graphical notation, feature structures are usually notated as attribute-value matrices, as in the lower part of the figure.

Petersen (2015: 47) proposes a formalization of frames that is largely consistent with the formalization of typed feature structures and operations on them provided by Carpenter (1992). Petersen (2015: 46) notes that her goal is to present a simple and rigid formalization in which no elements are added for merely technical or computational reasons. Her frames specifically represent concepts.

In the formalization given by Petersen (2015: 48), a frame consists of a finite set of nodes, one central node, a partial transition function determining the available attribute paths between nodes in the frame, a total node typing function that assigns exactly one type to each node, and a symmetric and anti-reflexive inequation relation. Frames can have at most one root node, from which all other nodes in the frame can be reached via an attribute path, as well as source nodes, which have no incoming arcs (attribute paths), and sink nodes, which have no outgoing arcs. Following Carpenter (1992: 37), substructures can be accessed via a composition of the transition function in the graph, which corresponds to following a path along specific attribute arcs starting at a specific node and ending in a target node that is the root of the substructure.

Consistent with Carpenter (1992: 11), Petersen (2015: 48) also assumes a type hierarchy. Types organize feature structures into natural classes in the domain that is being represented, and the inheritance hierarchy orders the types based on their generality. The type hierarchy induces a subsumption order on frames, as each node has exactly one type and the type of a frame's central node is also the type of the entire frame. Subsumption is defined based on a morphism function $h$ which ensures that a subsumed concept is always more specific than the subsuming concept. The alphabetic variance relation is an equivalence relation over the collection of frames that holds for all pairs of frames that mutually subsume each other.

Both Carpenter and Petersen define attributes as relationally interpreted functional concepts that map nodes to other nodes in a graph representing a typed feature structure via labeled directed arcs. While Carpenter (1992: 86) makes an explicit distinction between attributes and types, Petersen (2015: 57) views the set of attributes as merely a subset of the type set. An appropriateness specification determines which attributes are available in frames of a particular type and which values each attribute can have. A frame is well-typed if all defined attributes for each node and their values are consistent with the appropriateness specification.

Carpenter (1992: 45) defines how typed feature structures can be unified. The unification mechanism involves determining whether two pieces of partial information are consistent, and if they are, combining them into a single result. Incompatibility of feature structures is caused by nodes that are located at the same path in the feature structures and have types that cannot be unified. Unifying two well-typed feature structures must also result in a well-typed feature structure. The type of a feature structure resulting from unification is either as specific as or more specific than the types of the original feature structures. Feature structures that result from unification represent neither more nor less information than is contained in the original feature structures; unification is used to represent conjunctive information.

Unification can be conceptualized procedurally as follows, according to Carpenter (1992: 45–46). First, the root nodes of the original feature structures are identified, and the result is labeled with the unification of their types. Then, each time nodes that have been identified have values for identical attributes, those attribute nodes are identified, each time replacing the type of the identified nodes with the join of their original types. This is repeated until closure is reached. As soon as two nodes that are meant to be identified have inconsistent types, the unification fails.

In the terminology used by Carpenter (1992: 37), structure sharing is made possible by tags, like $\boxed{1}$, that label nodes. Labels appearing in multiple nodes signify that the nodes share their value; information that is shared is not duplicated in the notation. An inequation relation explicitly denotes nodes that are not identical, which can be useful for underspecified partial descriptions when unifying them with other partial descriptions.

## 4.4 KALLMEYER AND OSSWALD'S FORMALIZATION OF FRAMES

Kallmeyer & Osswald (2013) propose an implementation of frame semantics that supports the goal of Osswald & Van Valin (2014) to enable a decompositional analysis of predicates and event structure, and which shares certain aspects of the formalization by Petersen (2015) while differing from it in other ways. Kallmeyer & Osswald integrate this theory of frame semantics with Lexicalized Tree Adjoining Grammar (LTAG, Joshi & Schabes 1997) in a way that makes it possible to associate specific (partial) semantic descriptions with specific (partial) syntactic structures. This allows them to describe elementary constructions in terms of joint syntactic and semantic structures, which in turn makes it possible to derive larger syntactic and semantic structures based on these elementary descriptions. In the following, the frame semantics proposed by Kallmeyer & Osswald will be described. For an introduction to tree adjoining grammar, see Section 5.2 of this thesis. Section 5.4 goes into detail about the way Kallmeyer & Osswald combine LTAG and frame semantics to model certain linguistic phenomena.

Kallmeyer & Osswald (2013: 279) view frames as concept-centered representations that are potentially nested, typed, graph-like feature structures. They can be extended with additional subcomponents or constraints. These frames differ in certain ways from the feature structures defined by Carpenter (1992) and the frames defined by Petersen (2015) which were discussed in the previous section. First, Kallmeyer & Osswald's frames allow unification of substructures instead of only unifying complete feature structures by identifying their roots. Second, they argue for assuming that a frame does not necessarily have a unique root node, but can instead have several *base nodes*, and every node in the frame is reachable from some base node. Third, they allow frame descriptions to include non-functional relations between nodes. Furthermore, nodes in Kallmeyer & Osswald's frames can have multiple types as long as the types are not specified by the type hierarchy to be incompatible with each other.

The formalization given by Kallmeyer & Osswald (2013: 280–290) defines frames based on a signature $\langle A, T, R \rangle$ in which $A$ represents a finite set of attributes, $T$ represents a finite set of types, and $R$ represents a finite set of relation symbols that each have an arity of 2 or more. Each typed feature structure with relations over the signature $\langle A, T, R \rangle$ is defined as a quadruple $\langle V, \delta, \tau, \rho \rangle$ in which $V$ is a finite set of nodes, $\delta$ is a partial function from $V \times A$ to $V$ that represents the node transition function, $\tau$ is a function from $V$ to $\wp(T)$ that represents the typing function, and $\rho$ is a function from $n$-tuples of nodes from $V$ to subsets of the available $n$-ary relations in $R$.

One important aspect distinguishing Kallmeyer & Osswald's definition of frames from the formalization by Petersen (2015) is that each node can be assigned not a single type, but a set of types. If the typing function $\tau$ assigns an empty set of types to a node, this means that the node has the most general type, denoted by the symbol $\top$. With respect to their definition of subsumption, this has the effect that a feature

$$
\begin{array}{lll}
\textit{to} & \textit{into} & \textit{along} \\
\boxed{0}\begin{bmatrix} \textit{event} & \\ \text{PATH} & \begin{bmatrix} \textit{path} & \\ \text{ENDP} & \boxed{V} \end{bmatrix} \end{bmatrix} &
\boxed{0}\begin{bmatrix} \textit{event} & \\ \text{PATH} & \begin{bmatrix} \textit{path} & \\ \text{ENDP} & \boxed{V} \end{bmatrix} \end{bmatrix} &
\boxed{0}\begin{bmatrix} \textit{event} & \\ \text{PATH} & \begin{bmatrix} \textit{path} & \\ \text{REGION} & \boxed{V} \end{bmatrix} \end{bmatrix} \\
\boxed{1}\begin{bmatrix} \text{AT-REGION} & \boxed{W} \end{bmatrix} &
\boxed{1}\begin{bmatrix} \text{IN-REGION} & \boxed{W} \end{bmatrix} &
\boxed{1}\begin{bmatrix} \text{AT-REGION} & \boxed{W} \end{bmatrix} \\
\quad \textit{part-of}(\boxed{V}, \boxed{W}) & \quad \textit{part-of}(\boxed{V}, \boxed{W}) & \quad \textit{part-of}(\boxed{V}, \boxed{W})
\end{array}
$$

Figure 8: Frames for the prepositions *to*, *into* and *along* as proposed by Kallmeyer & Osswald (2013: 302).

structure $F_1$ subsumes another feature structure $F_2$ if the set of types assigned to $F_2$ is a subset of the types assigned to $F_1$, in addition to the subsumption requirements already specified by Carpenter (1992: 40) and Petersen (2015: 48) for feature structures generally. Additionally, the function $\rho$ of $F_1$ and $F_2$ must be taken into account to ensure that the morphism function $h$ allows the relations that are part of the descriptions of $F_1$ and $F_2$ to be unified.

A feature structure as defined by Carpenter (1992) and Petersen (2015: 50) can have at most one root node, from which each substructure of the feature structure must be accessible via an attribute path. Kallmeyer & Osswald (2013: 282) adapt this condition and instead require each component to be accesssible via an attribute path from a set of base nodes. The set of base labels is defined as $B = \{\boxed{0}, \boxed{1}, \boxed{2}, \ldots\}$. Base-labeled feature structures are defined with respect to $B$ such that every node in the feature structure is reachable from some base node as defined by the additional partial function $\beta$ from the set of base labels $B$ to the set of nodes $V$. Examples for frames with multiple base-labeled nodes are given in Figure 8 (taken from Kallmeyer & Osswald 2013: 302). In the figure, each preposition's frame involves two base-labeled feature structures and one relational description.

Base labels impact morphisms between feature structures: unifying base-labeled feature structures with different base labels results in a more specific feature structure with respect to subsumption. A feature structure can be relabeled in such a way that the same nodes of the structure are base-labeled as before. Unifying a feature structure with the base label $\boxed{0}$ and another feature structure with the base label $\boxed{1}$ under identification of $\boxed{0}$ and $\boxed{1}$ results in a unified feature structure whose attribute values are determined by the morphism function, and which carries both base labels $\boxed{0}$ and $\boxed{1}$.

In a semantic description in the form of Kallmeyer & Osswald's frames, every node in a base-labeled feature structure can be reached via an attribute path from a base node. This means that refering to the base-labeled node is sufficient to characterize an entire base-labeled feature structure.

Kallmeyer & Osswald (2013: 288) use so-called Horn constraints to formalize attribute-value constraints. Some example constraints are shown in (26). The constraint in (26a) expresses type inclusion (every instance of *activity* is also an instance of *event*). (26b) expresses type exclusion (instances of *causation* are never instances of *activity*). (26c) means that when the role ACTOR is present (literally: when it is an instance of the most general type $\top$), then the identity of the role AGENT is equal to the identity of the ACTOR.

(26) a. *activity* $\preceq$ *event*
b. *causation* $\preceq \neg$*activity*
c. AGENT : $\top \preceq$ AGENT $\doteq$ ACTOR

These constraints can be applied to frame structure descriptions to enable type inference. Instead of defining a type hierarchy directly, Kallmeyer & Osswald (2013: 289) use type inclusion and exclusion constraints to explicitly specify which conjunctive types (combinations of atomic types) are possible. These constraints can either be precompiled into a type hierarchy, or be used to compute type inference ad hoc. Conjunctive types and type hierarchies will be discussed in some more detail in Section 5.3.3.

The frames appearing in the metagrammars presented in Part III of this thesis follow the definitions provided by Kallmeyer & Osswald, since those definitions are the basis of the frame compiler implemented for the XMG framework by Lichte & Petitjean (2015).

## 4.5 REPRESENTING CONSTRUCTIONS AND VERBS WITH FRAMES

The representation of frames as typed, recursive feature structures in order to capture important properties of phenomena like causativity directly relates to the constructionist perspective on such phenomena. Goldberg (1995: 152) presents multiple examples of causative constructions that are instantiated with lexically non-causative verbs, such as (27), where the non-causative verb *sneeze* is integrated into the caused-motion construction.

(27) Frank sneezed the napkin off the table.

An analysis of such cases can benefit from a decompositional analysis of frames in the way discussed by Osswald & Van Valin (2014). Instead of assigning a frame of the type *causation* to the verb, the construction would provide a *causation* frame into which the meaning of the verb, *sneeze*, can be embedded. The construction interpretation for causation given by Goldberg (1995: 61), citing Alsina (1993: 498), similarly embeds the meaning of the verb into a causation event; it is described as CAUSE <`agent patient PRED` <...>>.

By representing both the semantics of the construction and the semantics of the verb with frames in the shape of recursive feature structures, structure sharing can be explicitly encoded using co-indexation. The interpretation of the combination of a construction and a verb can be determined by combining the partial frame provided by the construction with the partial frame provided by the verb. If the verb cannot appear in the construction, the frames should not be compatible with each other, so that the unification fails.

When analyzing the sentence in (27), the goal should not be to unify the entire *causation* frame contributed by the construction with the entire *sneeze* frame contributed by the verb. Causation is typically analyzed as an interplay of at least two events, one causing event and one resulting event. This is reflected in the predicate decompositions according to Rappaport Hovav & Levin (1998: 104), as exemplified in (28):

(28) [ [x ACT] CAUSE [BECOME [ y DRY] ] ]

$$\boxed{1}\begin{bmatrix} \textit{causation} & \\ \textsc{cause} & \boxed{2}\begin{bmatrix}\textit{activity} & \\ \textsc{actor} & \boxed{3}\end{bmatrix} \\ \textsc{effect} & \boxed{4}\begin{bmatrix}\textit{motion} & \\ \textsc{mover} & \boxed{5} \\ \textsc{path} & \boxed{6}\end{bmatrix}\end{bmatrix} \qquad \boxed{7}\begin{bmatrix}\textit{sneezing} & \\ \textsc{sneezer} & \boxed{8}\end{bmatrix} \qquad \boxed{1}\begin{bmatrix} \textit{causation} & \\ \textsc{cause} & \boxed{2}\,\boxed{7}\begin{bmatrix}\textit{activity, sneezing} & \\ \textsc{actor} & \boxed{3}\,\boxed{8}\ \textit{Frank} \\ \textsc{sneezer} & \boxed{3}\,\boxed{8}\end{bmatrix} \\ \textsc{effect} & \boxed{4}\begin{bmatrix}\textit{motion} & \\ \textsc{mover} & \boxed{5}\ \textit{napkin} \\ \textsc{path} & \boxed{6}\ \textit{off the table}\end{bmatrix}\end{bmatrix}$$

Figure 9: Semantic frame for sentence (27) *Frank sneezed the napkin off the table*. The frame on the left shows the meaning provided by the construction, the frame in the middle shows the meaning provided by the verb, and the frame on the right shows the combined frame. The representation of the PATH is simplified because the details are not of interest here.

Recall that Osswald & Van Valin (2014: 140) propose a representation of causation in which a frame of the type *causation* contains embedded frames representing the CAUSE event and the EFFECT event individually. Participants that are expressed once in the construction can appear in both the CAUSE and EFFECT subframes if this is linguistically motivated. This would be the case, for instance, if the entity whose state change is expressed in the EFFECT subframe is also impacted directly by the event expressed in the CAUSE subframe. It would also be possible for a *causation* event to not have an agentive causer, for instance if the event is due to a natural force; in such cases, the type of the CAUSE subframe shown on the left of Figure 5 on page 38 may be something other than an *activity*. With respect to the example sentence in (27), a mechanism is necessary for determining how the semantic contributions from the construction and the verb and the participants are to be combined with each other.

The construction contributes the fact that the situation described in the sentence is a *causation* scenario. It also specifies that the effect that is being caused is a motion of some entity along some path, which is expressed via a directional prepositional phrase. The activity that causes this motion is not specified by the construction, and neither is the nature of the motion, the identity of the impacted entity, or the specific path that the entity takes. These pieces of information are all contributed by the specific lexical material that appears in the construction. Thus, a frame representing the sentence could look as shown in Figure 9. The frame on the right is color-coded to show the source of attributes and values: types, attributes and labels that are contributed by the construction (left frame) are shown in purple. Types, attributes and labels that are contributed by the verb (middle frame) are shown in orange. Attribute values that are provided by the arguments are shown in black.

The frame that is associated with the lexical item *sneeze* is represented here as belonging to a type named *sneezing*. The *sneezing* type is assumed to be compatible with the *activity* type, which is the reason why the identification of the feature structures labeled as $\boxed{2}$ and $\boxed{7}$ succeeds. The name *sneezing* for this type is transparently derived from the verb itself. While *sneeze.v* is not included in FrameNet's lexical unit list, similar verbs like *cough.v* or *gasp.v* are included and associated with a frame named *make noise*. In Sentence (27), the production of a noise does not seem to be the most relevant property of the *sneezing* event. In a systematically organized frame hierarchy as envisioned by Osswald & Van Valin (2014), multiple inheritance links could connect a frame type like *sneezing* to relevant frame types like *make noise*, *breathing* and possibly *excreting* in parallel.

In the frames presented throughout this thesis, type names associated with verbs will typically take the form *v-ing*. This is inspired by certain FrameNet frame types like the aforementioned *breathing* or *excreting*. However, as already noted by Osswald & Van Valin (2014: 137), "there is obviously no general convention for naming frames [in FrameNet]", and there are pairs of intuitively closely related lexical units whose associated FrameNet frames follow different naming patterns. For instance, *absorb.v* evokes the *soaking up* frame while its approximate antonym *leak.v* evokes the *fluidic motion* frame. This thesis will follow the *v-ing* pattern in order to clearly express that these frames are semantic representations of *v-ing* events. The alternative of using a bare infinitive in that context would possibly lead to confusion over *v-ing* events as opposed to zero-derived nouns, as in *jumping*/*jump*. Whether the semantic frame for *jump* (as a noun) should be identical to or significantly different from the frame for *jumping* (as a verb) will not be discussed in this thesis. The *v-ing* frame type naming pattern is also used, for instance, by Kallmeyer & Osswald (2013: 291).

The structure of the combined frame in Figure 9 on page 45 resembles the frame proposed by Kallmeyer & Osswald (2013: 301) for the lexical meaning of *throw*, to which they assign the top-level type *onset causation*. This is because the caused-motion construction can be instantiated either by non-causative verbs, like *sneeze*, or by verbs that lexically encode caused motion, like *throw*. In both cases, the meaning of the sentence in which the construction is being instantiated is accurately represented by a *causation* frame whose value for the EFFECT attribute is a *motion* event with a MOVER moving along a specific PATH as expressed by the directional phrase in the sentence.

The frame in Figure 9 represents the causation event at the level of detail argued for by Osswald & Van Valin (2014). The semantic contribution from the construction is clearly represented, but no specific values for the attributes specified in the frame can be inserted until the construction is instantiated by a verb and all roles are associated with a participant. Note that the participant *Frank* fills both the more general ACTOR role (contributed by the construction) and the more specific SNEEZER role (contributed by the verb). The identification of the labels attached to two roles in the frame signifies that the attributes ACTOR and SNEEZER have the same value. In a system performing these analyses, a type hierarchy is required which specifies whether the types *activity* and *sneezing* are compatible with each other. If they are, then the result of the unification of the nodes may be labeled with only the more specific one of the types without information loss. Thus, the type *sneezing* would be subsumed by the type *activity*, and only the more specific label would need to be included.

When integrating the verb into the construction, the lexicon entry representing the construction must specify which of its features will unify with (or in other words, embed) the frame provided by the verb. A detailed examination of this issue can be found in Part III of this thesis.

Another example sentence (adapted from Goldberg (1995: 141)) that also warrants a closer look is given in (29).

(29) Chris baked Jan a cake.

$$
\begin{bmatrix}
\textit{intended causation} & \\
\textsc{performed activity} & \begin{bmatrix} \textit{creation, baking} & \\ \textsc{actor} & \boxed{1}\ \textit{Chris} \\ \textsc{cook} & \boxed{1} \\ \textsc{creation} & \boxed{2} \\ \textsc{produced food} & \boxed{2}\ \textit{cake} \end{bmatrix} \\
\textsc{intended effect} & \begin{bmatrix} \textit{transfer} & \\ \textsc{recipient} & \textit{Jan} \\ \textsc{goods} & \boxed{2} \end{bmatrix}
\end{bmatrix}
$$

Figure 10: Semantic frame for sentence (29), *Chris baked Jan a cake*. Contributions from the construction are shown in purple, contributions from the verb are shown in orange. The details of the *intended causation* frame are not in focus here; a different set of attribute labels for this event type may be suitable.

In this sentence, the ditransitive construction, which has a basic transfer meaning, is instantiated by the creation verb *bake*. Goldberg views cases like this as instances of a special case of the construction, which is described as "Agent intends to cause recipient to receive patient". A possible frame for the sentence could look like the one shown in Figure 10. Again, contributions from the construction are shown in purple and contributions from the verb are shown in orange.

The sentence that is represented by the frame in Figure 10 is an example of a construction assigning an additional role to a participant that already fills a role contributed by the verb: the *cake* is the PRODUCED FOOD in the *baking* event[2], and at the same time fills the role of the entity that is being given to someone in the *transfer* subframe of the construction. This example illustrates again the advantages of representing the semantics of both constructions and verbs as typed, recursive feature structures that enable structure sharing between nodes.

Chapter 3 of this thesis reported the way Goldberg (1995, 2002) takes constructional profiling and verbal profiling into account in order to correctly integrate verbs into constructions. She explains how this integration can succeed even when there is a mismatch between the sets of profiled roles from the construction and the verb. With a frame-based analysis, these mismatches do not pose a problem. Roles that are contributed by the verb, whether they are profiled or not, can be integrated into the frame by unifying the entire frame representing the verb with some relevant node within the frame representing the construction. Structure sharing between individual substructures, as illustrated in Figure 10, can be employed to specifically link participant roles from the verb to argument roles from the construction.

Which elements of each frame are identified with each other is a result of the association between syntactic elements and specific parts of each frame. For instance, in a basic transitive active sentence, the COOK role of a *baking* event will be identified with the lexical frame of the entity found in the syntactic subject position. Simultaneously, the ACTOR role of the ditransitive construction in the form shown in the figure will also be identified with the entity in the subject position. Therefore, in the example, both the COOK and the ACTOR role are co-indexed and identified with the entity *Chris*. Chapter 5

2 This role name is taken from the FrameNet frame *cooking creation*, which is evoked by the verb *bake*.

$$\begin{bmatrix} \textit{moving in place} & \\ \text{ACTOR} & \boxed{1} \\ \text{THEME} & \begin{bmatrix} \textit{body part} & \\ \text{TYPE} & \textit{head} \\ \text{POSSESSOR} & \boxed{1} \end{bmatrix} \end{bmatrix}$$

Figure 11: Possible frame structure for the lexicon entry for *nod*, ensuring that the involvement of the *head* possessed by the ACTOR is encoded even in constructions where that entity is not explicitly mentioned. Details on the pattern of movement can be expressed, for instance, in an additional MOVEMENT PATTERN attribute.

of this thesis will go into detail with respect to the mechanisms that are employed to link syntactic structures to semantic structures.

The exclusive contribution of roles from verbs is relevant in the context of verb alternations whose constructions express different sets of arguments, yet seem to express (nearly) identical meanings. For instance, the understood body-part alternation (Rice 1988: 205, Levin 1993: 34) allows the inclusion or exclusion of arguments that refer to "understood" body parts, that is, arguments that are so closely lexically related to the given verb that their absence does not detract from conveying the same meaning, as in (30):

(30) a. Ahab nods his head.
b. Ahab nods.

In sentence (30a), the fact that *nodding* involves a movement of the ACTOR's *head* is explicitly expressed. In contrast, the *head* does not show up in the alternate variant (30b). A frame representation of the two sentences should show that there is no (obvious) meaning difference between the alternants. Practically, such a representation can be produced by referring to the lexicon entry for *nod*, which might explicitly specify the involvement of the *head* in this event. Then, that attribute, which is provided in the lexicon entry, would show up in the frames for both sentences: in (30a), the participant role is "fused" or unified with the argument role whose syntactic position corresponds to the position of the constituent *his head*, and in (30b), the participant role is added to the frame based exclusively on the lexicon entry for the verb. Specifying this participant in the lexicon entry can also reflect selectional restrictions, so that only entities which possess a head or something that can be construed as one have the ability to *nod*. A possible frame structure for the verb *nod* is given in Figure 11.

## 4.6 THEMATIC RELATIONS AND SEMANTIC MACROROLES

This thesis, and in particular, the metagrammars developed in Part III, make use of the thematic relations and semantic macroroles posited by Role and Reference Grammar (RRG, Van Valin & LaPolla 1997, Van Valin 2005). Thematic relations are generalizations across verb-specific roles and include roles like AGENT, INSTRUMENT, or THEME. They are identified via their argument positions in decompositional logical structure representations. Semantic macroroles are generalizations across thematic relations: the semantic macrorole ACTOR is "[...] a generalization across agent, experiencer, instrument and other roles, while UNDERGOER is a generalization subsuming patient, theme, recipient and other roles" (Van Valin 2005: 53).

Van Valin (2005: 56) notes that the thematic relation AGENT, which is assumed to be the intentional, volitional and controlling participant in an event, is a special type of EFFECTOR. This is illustrated with sentences (31) through (33) (taken from Van Valin 2005: 56). The sentences show that the verb *murder* requires a perpetrator which is acting intentionally and volitionally, while this is not required by the verb *kill*.

(31) a. The man killed his neighbour.
b. The man intentionally killed his neighbour.
c. The man accidentally killed his neighbour.

(32) a. The man murdered his neighbour.
b. ? The man intentionally murdered his neighbour.
c. * The man accidentally murdered his neighbour.

(33) a. A branch falling from Pat's tree killed his neighbour.
b. * A branch falling from Pat's tree murdered his neighbour.

Semantic macroroles make it possible to generalize across these distinctions, and refer more generally to the more "agent-like" entity in a clause via the macrorole ACTOR. In the context of the alternations that will be modeled in this thesis, this is particularly important, as different thematic relations can be associated with the ACTOR macrorole in an event or subevent. For instance, the instrument-subject alternation licenses sentences in which the ACTOR role in a causing subevent is filled not by an agentive entity such as a human, but instead by an instrument that is being used by an unexpressed agentive entity.

Thus, verbs like *break* may take an AGENT subject, as in *Queequeg broke the window*, or an instrument subject (which Van Valin: 58–59 analyses as the thematic relation EFFECTOR), as in *The hammer broke the window*. In both cases, the entity in the subject position fills the ACTOR macrorole in the causing subevent. Concerning the nature of instruments that can function as ACTORS, as Van Valin & LaPolla (1997) put it, "… INSTRUMENTS are IMPLEMENTS in a causal chain which are also EFFECTORS" (Van Valin & LaPolla 1997: 122). An IMPLEMENT cannot fill the ACTOR role if it is not also an EFFECTOR. This is why sentences like **The spoon eats the soup* are not acceptable.

The causative-inchoative alternation allows UNDERGOER arguments of participating verbs to appear as the single argument in a clause, licensing sentences like *The window broke*. A distinction between specific thematic roles like PATIENT or THEME is not necessary at this level of description, since the behavior of such verbs and their arguments in this alternation can more generally be described with reference to semantic macroroles.

Van Valin & LaPolla (1997: 146) propose an actor-undergoer hierarchy in which thematic relations are ordered by their ability to function as the ACTOR or UNDERGOER in an event. While the ACTOR macrorole applies to more agent-like thematic relations like AGENT and EFFECTOR, the UNDERGOER macrorole applies to thematic relations that are affected, patient-like participants in states of affairs, like THEME or PATIENT. The ACTOR and UNDERGOER roles are sometimes also referred to as the "logical subject" and "logical object". If a verb has an agent argument, that argument will always fill the ACTOR macrorole, and if a verb has a patient argument, that argument will always be the UNDERGOER. An event can have no macroroles, only an ACTOR, only an UNDERGOER, or both an ACTOR and an UNDERGOER. Although a verb can have at most two macroroles, it is possible for a verb to have more than two direct core arguments (direct NPs that appear in the syntax), or to have two direct core arguments but only one macrorole. For more on this, see Van Valin (2005: 62–64).

The phenomena modeled in Part III of this thesis involve various causative constructions. Causative sentences express complex events in which a causing subevent triggers, enables or otherwise brings about a caused subevent. The corresponding frame representation for such sentences, following Osswald & Van Valin (2014), Kallmeyer & Osswald (2013), Kallmeyer et al. (2016), Lichte & Petitjean (2015), and others, takes the form of a *causation*-type frame with a CAUSE substructure and an EFFECT substructure. The metagrammars developed in this thesis cover sentences in which the lexical meaning of the verb is identified with either the entire *causation* frame or its CAUSE or EFFECT subframe.

Each of the three events involved in such semantic structures can involve different combinations of the ACTOR and UNDERGOER macroroles. In certain cases, the ACTOR of the causing subevent will be identified with the ACTOR of the complex *causation* event, and/or the UNDERGOER of the caused subevent will be identified with the UNDERGOER of the *causation* event. The constraints that are responsible for lifting these semantic roles from a substructure to its superstructure follow Kallmeyer et al. (2016). The subevent whose nature is not described by the verb will not always necessarily receive ACTOR or UNDERGOER attributes in the metagrammars presented here.

## 4.7 CONCLUSION

As Osswald & Van Valin (2014: 139) argue, semantic frames in the form of typed, recursive feature structures are well-suited to represent meaning decompositionally. The semantic frame for a sentence in which a particular verb instantiates a particular construction can be determined based on the interplay of the semantics contributed by the construction and the semantics contributed by the lexical meaning of the verb. Similarities between constructions can be encoded with frame structures with an overlap in their substructures, as in the constructions associated with the causative-inchoative alternation as proposed by Osswald & Van Valin (2014: 140). In that approach, the meaning representation of the inchoative construction is embedded as a subframe into the meaning representation of the causative construction.

The interactions between constructions and verbs described here can be implemented using any grammar formalism that allows the definition of lexicon entries as typed, recursive feature structures. Constructions are characterized by certain syntactic patterns, such as the double object pattern in the ditransitive construction. Additionally, each construction needs to be described semantically, and the grammar must also contain some set of constraints describing how the semantic contributions from verbs and arguments are incorporated into the frame representing a construction.

Kallmeyer & Osswald (2013) show that Lexicalized Tree Adjoining Grammar (LTAG) is a useful formalism for these requirements, because LTAG's elementary trees represent full argument projections. Constructions and verbs are described in terms of elementary trees and their corresponding frames in the form of typed, recursive feature structures, and the tree rewriting operations anchoring, substitution and adjunction take the trees as well as the frames into account to yield a unified whole representation. Kallmeyer & Osswald present two applications of these mechanisms, revolving around English directed motion expressions and the dative alternation. Lichte & Petitjean (2015: 188) refer to the work by Kallmeyer & Osswald and implement an extension to the extensible MetaGrammar (XMG) framework for LTAG that allows grammar authors to add a `<frame>` dimension to metagrammars. The `<frame>` dimension describes frames as typed, recursive feature structures as discussed in this chapter. Part III of this thesis will build on the work by Kallmeyer & Osswald (2013) and Lichte & Petitjean (2015) and model a set of alternation-specific constructions as well as constructions unrelated to alternations in LTAG.

For the purposes of this thesis, typed, recursive feature structures representing semantic frames are an attractive instrument to model verb alternations and related constructions, because they are ideally suited to jointly modeling syntactic and semantic processes based on constraint-based and unification-based mechanisms.

Part III

# USING TREE ADJOINING GRAMMAR AND FRAME SEMANTICS TO MODEL ALTERNATIONS AND RELATED CONSTRUCTIONS

This part of the thesis is concerned with modeling the interactions between constructions and verbs in the context of verb alternations. Particular attention is dedicated to cases in which constructions share a syntactic form, but are semantically characterized by different frames. When integrating a verb into a construction, the lexicon entry representing the construction must specify which of its features will unify with the frame provided by the verb. This part builds on work by Kallmeyer & Osswald (2013) and Lichte & Petitjean (2015) and describes a set of alternation-specific constructions as well as constructions unrelated to alternations in an LTAG metagrammar with semantic frames in order to model alternation behavior.

# 5 A TREE ADJOINING GRAMMAR APPROACH TO MODELING VERB ALTERNATIONS

## 5.1 INTRODUCTION

This part of the thesis explores interactions between verbs, alternations, and related constructions by modeling various verbs, arguments and constructions in a metagrammar for a lexicalized tree adjoining grammar (LTAG, Joshi & Schabes 1997) that includes frame-semantic meaning representations. The focus will be on the distinction between alternation-specific constructions and syntactically identical constructions that are not associated with alternations, and whose semantic contribution is different from that of the constructions belonging to an alternation. When constructions involve the same syntactic form, but different semantics, it may be challenging to disambiguate sentences instantiating one of the constructions. Semantic constraints from the different constructions and lexical properties of the verb in the sentence and its arguments can be consulted to disambiguate, or to derive multiple analyses in case the constraints from multiple constructions are all fulfilled.

The constructions and alternations that will be discussed in this thesis part illustrate this. The metagrammar that will be developed covers a set of sentence types which instantiate constructions that differ with respect to intransitivity or transitivity, the presence or absence of certain prepositional phrases, and their semantic contribution and requirements on verbs and their arguments. Whether a verb can instantiate a particular construction depends, among other things, on a semantic type hierarchy in which the compatibility of different types is specified in the form of constraints. The sentence types are parsed to verify that the derived syntactic structure and semantic frame for each sentence correspond to the predicted outcomes. Certain sentences share their syntactic structure, but are associated with different semantic frames due to the different constructions instantiated in each case.

The metagrammar will model the caused-motion construction, the induced action construction (which features in the induced action alternation, Levin 1993: 31), the causative and inchoative constructions (which feature in the causative-inchoative alternation, Levin 1993: 27), and the instrument subject construction (which features in the instrument subject alternation, Levin 1993: 80), as well as basic activity constructions, constructions expressing directed motion in the form of prepositional phrase adjuncts, and constructions expressing the use of an instrument in the form of a prepositional phrase adjunct headed by *with*. Chapter 6 focuses on the part of the metagrammar that is concerned with the caused-motion construction and the induced action construction. Chapter 7 develops the part of the metagrammar that is concerned with the causative and inchoative constructions and the instrument subject construction. Each chapter presents a working, complete metagrammar. Finally, Chapter 8 merges both metagrammars into one larger metagrammar and sketches how more constructions could be added to create a more exhaustive metagrammar covering a wider range of constructions within and outside of verb alternations.

According to Levin (1993: 4), (non-)participation in a particular alternation is governed by certain meaning facets that are inherent to the semantics of each verb (see Section 2.1 of this thesis). She lists participating and non-participating verbs for each alternation she covers, treating alternation behavior as a result of certain lexical properties of each verb. In the metagrammars developed here, alternation participation is not directly marked in the lexicon entries of individual verbs. Instead, constructions are made available to verbs based on the semantic frame type hierarchy that is part of the metagrammar: each verb's semantic type and the semantic types of all arguments are either compatible or incompatible with a set of types that are required for a sentence to instantiate the various constructions. This approach mirrors Levin's idea of meaning facets governing alternation participation. The metagrammar does not mark alternation participation explicitly in the lexicon or in the type hierarchy: alternation participation is exclusively determined by the availability of each of the alternation-specific constructions for a given verb (and its arguments). This construction-focused view is close to the proposal by Goldberg (2002) that argues for studying constructions independently of each other instead of analysing alternations as sets of paraphrases.

However, in multiple cases, the meanings of the constructions involved in an alternation do explicitly share certain aspects or substructures, as illustrated by Osswald & Van Valin (2014: 140) (see Section 4.2 of this thesis). The modular structure of the metagrammar framework used in this thesis makes it possible to define these shared syntactic or semantic properties only once and make them available to the different constructions via inheritance. This shared inheritance from common superclasses can be seen as an indicator of the relatedness of constructions. This construction inheritance structure reflects the constructionist idea that constructions are organized in an inheritance network, the constructicon (see Chapter 3 of this thesis for Goldberg's description of the constructicon, and Section 5.5.3 for a discussion of the way XMG metagrammars can be used to implement such an inheritance network of constructions). The parser takes all constraints from the grammar into account and derives all syntactically and semantically valid analyses for each input sentence.

The metagrammars will be implemented using the eXtensible MetaGrammar framework (XMG, Crabbé et al. 2013, Petitjean, Duchier & Parmentier 2016) and the XMG-2 compilers available online at `https://github.com/spetitjean/XMG-2`. Each metagrammar developed here will compile to a lexicalized tree adjoining grammar. Frame representations will be encoded in the `<frame>` dimension of the metagrammars and processed by the frame compiler for XMG, which was developed by Lichte & Petitjean (2015) based on work by Kallmeyer & Osswald (2013). The TuLiPA parser (Kallmeyer et al. 2008) will be used to parse example sentences based on the compiled grammars. This parser has been extended by Arps & Petitjean (2018) to allow for the simultaneous derivation of syntactic trees and semantic frames.

This thesis part is structured as follows. A brief introduction to LTAG is given in Section 5.2. The metagrammar description language XMG is described in Section 5.3. The way syntax and semantics can be modeled jointly within the XMG framework is discussed in Section 5.4. The choice of XMG and LTAG for this task will be discussed and justified in Section 5.5. Chapters 6 and 7 present two metagrammars specifically developed for this thesis to model separate alternation-related phenomena. Chapter 8 merges the two metagrammars and concludes the thesis part.

Chapter 6 revolves around modeling verbs that instantiate the caused-motion construction and verbs that participate in the induced action alternation. The syntactic patterns involved in these constructions overlap with each other, but the availability of different semantic interpretations varies between verbs. The caused-motion construction is typically considered to be relatively productive, within certain limitations with respect to the event types associated with verbs (Goldberg 1995: 165, Goldberg & Jackendoff 2004: 540). The induced action alternation imposes more specific constraints on participating verbs and is thus less productive and available to a smaller number of verbs (Levin 1993, Cruse 1972). These differences in the availability of the relevant constructions are reflected in the metagrammar in the form of more or less specific semantic constraints imposed by the constructions on verbs instantiating them.

Chapter 7 discusses the causative-inchoative alternation and the instrument subject alternation, and contrasts them with verbs that participate in neither of these alternations. These two alternations both allow different semantic roles to appear in the syntactic subject position, but they differ in terms of the specific semantic roles that can be expressed as the subject. Again, the syntactic patterns involved in the two phenomena are at times superficially identical, but the available semantic interpretations for sentences differ depending on the verb, its arguments, and the specific constructions that are being instantiated.

A previous version of the work presented in Chapter 6 previously appeared as Seyffarth (2019b). The metagrammar discussed in that paper is available online at `https://github.com/eseyffarth/caused-motion-xmg`. The metagrammar that is presented in Chapter 7 is a new contribution.

All code for the work presented in this chapter is available online at `https://github.com/eseyffarth/ltag-mg-for-alternations`.

## 5.2 TREE ADJOINING GRAMMAR

The term tree adjoining grammar (TAG, Vijay-Shanker 1987, Joshi & Schabes 1997, Abeillé & Rambow 2000) refers to a tree-rewriting formalism that is motivated by linguistic and formal considerations. The information presented in this section is mainly based on Joshi & Schabes (1997).

The elementary objects in TAG are a finite set of trees of finite depth that can be composed to form larger structures via the operations adjunction and substitution, which are explained in Section 5.2.1. Each elementary tree in a tree adjoining grammar is either an initial tree or an auxiliary tree. In initial trees, non-leaf nodes are labeled with non-terminal symbols, and leaf nodes are either labeled with terminals, or labeled with non-terminals and marked for substitution. Auxiliary trees must additionally have exactly one foot node at the leaf level, whose label must be identical to the root node's label (and must therefore be a non-terminal). Auxiliary trees can be composed with other trees via adjunction. Examples for initial and auxiliary trees are shown in Figure 12 on page 58. Foot nodes are typically marked with a star symbol $\star$, and substitution nodes are typically marked with a down arrow symbol $\downarrow$ (see e.g. Abeillé et al. 1990: 4).

In a lexicalized tree adjoining grammar, each elementary tree is anchored by a lexical item, called the *anchor*. The lexicon for the grammar specifies for each lexical item a finite set of trees that can be anchored by it. These lexicalized elementary trees can

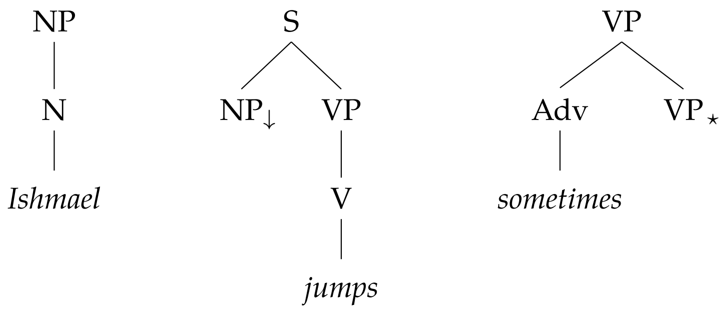


Figure 12: Initial trees anchored by *Ishmael* and *jumps* and auxiliary tree anchored by *sometimes*.

then be composed with other elementary trees via the available operations to form derived trees.

#### 5.2.1 *Operations on trees*

The two tree composition operations in TAG, substitution and adjunction, each require certain nodes of the participating trees to be compatible with each other. Nodes must share the same label in order to be compatible. When nodes are further described with feature structures, those feature structures must be unifiable without conflict in order for the nodes to be compatible.

Substitution is the process by which a tree is extended by another tree that is inserted at a leaf node that is marked for substitution. Substitution nodes are always labeled with non-terminal symbols. The root node of the substituting tree must be compatible in terms of node label and feature structure content with the leaf node at which it is being inserted. In the context of the example trees given in Figure 12, the initial tree anchored by *Ishmael* can be substituted into the tree anchored by *jumps* at the substitution node NP. This operation identifies the NP root node of the former with the leaf NP node of the latter.

Adjunction is the process by which an auxiliary tree is inserted into another tree at a specific non-leaf node labeled with a non-terminal symbol. Only auxiliary trees can be inserted into other trees via adjunction. Substitution nodes and foot nodes are not valid adjunction sites, but auxiliary trees can be adjoined to other auxiliary trees at nodes that are internal and not marked for substitution or as foot nodes. The foot node and root node of the adjoining auxiliary tree must be compatible in terms of node label and feature structure content with the node at which the tree is inserted in the target tree. In the context of the example trees given in Figure 12, the auxiliary tree anchored by *sometimes* can adjoin to the internal VP node of the initial tree anchored by *jumps*. This identifies the root VP node of the former with the internal VP node of the latter, while all nodes in the *jumps* tree that were previously dominated by the original internal VP node are now dominated by the foot node of the adjoining tree.

An example for a tree derivation involving substitution, adjunction, and lexical anchoring is given in Figure 13 on page 59. This figure explicitly shows how elementary trees are anchored by lexical items. Anchor nodes are marked with a diamond symbol ♢. In this figure and all following figures showing TAG derivations, (selected) elementary trees are labeled with a name indicating their function in the grammar. Here, the

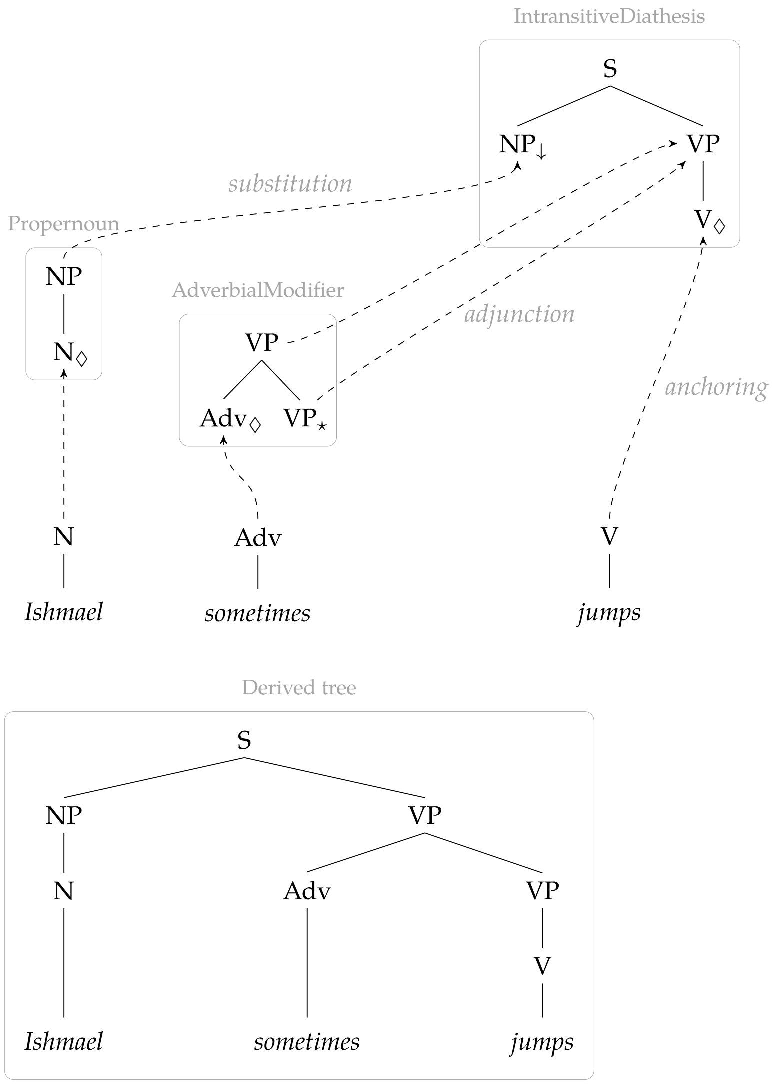


Figure 13: Substitution, adjunction, and lexical anchoring, illustrated with the example sentence *Ishmael sometimes jumps*.

tree anchored by *Ishmael* is called *Propernoun*. The tree anchored by *jumps* is called *IntransitiveDiathesis*. The tree anchored by *sometimes* is called *AdverbialModifier*.

The *Propernoun* tree is composed with the *IntransitiveDiathesis* tree via substitution at the NP node in the subject position. This is possible because the NP node in the verb-anchored tree is marked for substitution and possesses the same label as the root node of the *Propernoun* tree. The *AdverbialModifier* tree is composed with the *IntransitiveDiathesis* tree via adjunction at the internal VP node. That VP node has the same label as the root node and the foot node of the adjoining tree.

Since the operations are constraint-based, they are not performed in a particular order. The derived tree resulting from composing two or more trees is called *completed* if all leaf nodes are labeled with terminal symbols, and no leaf nodes are marked for substitution. This is the case for the derived tree shown in Figure 13.

Adjunction makes it possible for nodes to be far away from each other in a derived tree even when they originally appear in the same elementary tree. This extended domain of locality in tree adjoining grammar has advantages for modeling long-distance dependencies, which are then regarded not as global relationships but as local ones ("complicate locally, simplify globally", Joshi 2004: 638). For instance, in an elementary tree modeling *wh*-movement, certain features will be shared between the nodes in that elementary tree that represent the *wh*-word and the extraction site, and auxiliary trees can be adjoined between them during derivation.

In feature structure based tree adjoining grammars (Vijay-Shanker & Joshi 1988, Abeillé et al. 1990), the nodes appearing in elementary trees are described in terms of untyped feature structures describing the respective node's properties via a set of attribute-value pairs. These properties must be taken into account during substitution and adjunction: the feature structures of the relevant nodes must be unified without conflicts in order for the operation to be successful. The feature structures can thus be used to express additional adjunction constraints.

Each node in each elementary tree can be described in terms of two separate feature structures, named *top* and *bottom*, except substitution nodes, which lack a bottom feature structure (Abeillé et al. 1990: 7). The top feature structure describes the node's relation to its supertree, while the bottom feature structure describes the node's relation to its descendants. During substitution, the top feature structure of the target node and the top feature structure of the substituting tree's root node are unified. The bottom feature structure of the substituting tree's root node becomes the bottom feature structure of the target node. During adjunction, the top feature structure of the target node unifies with the top feature structure of the root node of the adjoining tree. The bottom feature structure of the target node unifies with the bottom feature structure of the adjoining tree's foot node. In derived trees, the top and bottom feature structure of each node are unified, whether or not adjunction has taken place at that node.

The metagrammars that will be developed in this thesis part will not explicitly make use of top and bottom feature structures, but nodes will be characterized with a set of attributes and their values. This effectively corresponds to a grammar in which the top and bottom feature structures of each node in each elementary tree are identical from the start.

Values in feature structures describing nodes in the same elementary tree can be co-indexed to express that different attributes share the same value (see also Chapter 4 of this thesis.) This is useful, for instance, to explicitly require the main verb and the subject NP in a sentence to share certain agreement features. Unification allows attribute values from substituting or adjoining trees to be co-indexed with attribute values from other trees involved in the derivation.

## 5.3 EXTENSIBLE METAGRAMMAR (XMG)

A lexicalized tree adjoining grammar can be described with a so-called metagrammar ("the grammar of a tree grammar", in the words of Crabbé & Duchier 2005: 42). Such a metagrammar reduces redundancy by describing smaller fragments rather than listing all elementary trees separately (Crabbé 2005a: 91, Crabbé & Duchier 2005: 41). Descriptions of the elementary trees of the grammar result from disjunction and conjunction of the fragments described by different metagrammar classes, which form an inheritance hierarchy. Compiling the metagrammar resolves the descriptions and yields all valid elementary trees licensed by the grammar. Metagrammars can be created for various grammar formalisms, provided that a compiler exists that can transform the metagrammar class descriptions into a grammar. This thesis is only concerned with metagrammars for tree adjoining grammars.

As Crabbé (2005a: 84) notes, strong lexicalization in tree adjoining grammars has advantages for parsing efficiency and for representing lexical exceptions or multiword exceptions. At the same time, it is desirable to model generalizations beyond what is possible with strongly lexicalized TAG. Such generalizations can be captured by unanchored trees; the lexical item anchoring these trees is underspecified, and actual elementary trees are generated during compilation.

In a metagrammar, tree fragments are described in terms of dominance, precedence, immediate dominance and immediate precedence relations between nodes. Alternative realizations can be described via disjunctions, and complex descriptions can be assembled based on simpler ones via conjunctions (Crabbé & Duchier 2005: 41). Abstraction allows complex statements to be reused in other fragment descriptions. This factorization makes it possible for the metagrammar to reflect linguistic generalizations, helps in preventing redundancy in the metagrammar, and thus reduces manual effort when updating or changing the grammar. The same mechanisms in metagrammars also provide the possibility of allowing alternative realizations of certain structures. Compiling the metagrammar results in a set of minimal models (Crabbé & Duchier 2005: 43), which are the simplest trees that fulfill all relevant constraints.

With respect to the constructions modeled in this thesis, the factorization allowed by metagrammars is advantageous for modeling the syntactic function and semantic role of a verb argument separately. Verb arguments can anchor specific initial trees that can then be attached via substitution to different nodes of trees anchored by verbs. This can be used to implement different constructions in the sense of Goldberg (1995).

The metagrammars presented in this thesis part will be developed in the XMG description language (Crabbé et al. 2013). Each part of the metagrammars will be compiled with a dedicated compiler, using resources provided by the XMG-2 project (Petitjean, Duchier & Parmentier 2016). Descriptions of partial syntactic trees and partial semantic frames are processed with the `synframe` compiler (Lichte & Petitjean 2015: 207). The compiled grammar files can then be loaded into the TuLiPA parser (Kallmeyer et al. 2008, Arps & Petitjean 2018) to parse input sentences.

### 5.3.1 *XMG: A declarative, notationally expressive, and extensible description language*

eXtensible MetaGrammar (XMG, Crabbé 2005a, Crabbé et al. 2013) is a description language for metagrammars which can be compiled into a grammar with which input sequences can be parsed. Descriptions are organized into named classes, which contain descriptions for (linguistic) structures in various dimensions. Classes can import other classes or inherit from them, which makes the content of the imported or inherited class available to the importing or inheriting class. This inheritance mechanism makes it possible for tree fragments to be organized into a linguistically motivated inheritance hierarchy. Furthermore, uniqueness conditions can be included to signify that a property may only be associated with one node.

Note that the inheritance mechanism in the metagrammar's class hierarchy should not be understood as an exhaustive taxonomic tree (Osswald 2002: 9–10). Instead, the classes form a multiple inheritance hierarchy in the style used in object-oriented programming, where one class can inherit descriptions from multiple other classes. The inheriting class is then defined by the joint features of the inherited classes in combination with all constraints that are specific to the inheriting class. The XMG classes are related by what Meyer (1997: 532) calls *structure inheritance*. In the example class hierarchy from Crabbé et al. (2013) that will be discussed in Section 5.3.2 of this thesis, a class named *IntransitiveFamily* inherits from two classes named *ActiveVerbForm* and *Subject*. This does not mean that intransitive expressions are a more specific sort of active verb form and simultaneously a more specific sort of subject. Instead, the description of the intransitive family class includes all constraints belonging to each of the two inherited classes, as well as possibly some additional constraints that are specific to *IntransitiveFamily* itself. Multiple inheritance hierarchies have also been used in early HPSG work (see Sag 1997: 432 and the references therein).

Multiple dimensions for semantic representations are available in the XMG metagrammar framework; this thesis uses the `<frame>` dimension. Each class describing a construction contains both a specification of the construction's syntactic structure and a specification of the construction's meaning contribution. The semantic frames that are part of the lexicon entries for verbs and their arguments are unified with the semantic frames associated with the nodes they are substituted into. Unification can be blocked if the semantic types for a specific frame attribute in a construction and the corresponding substructure from the relevant substituting or adjoining tree are incompatible with each other according to a frame type hierarchy.

Crabbé et al. (2013: 592) point out three main strengths of XMG. First, the description language is declarative, so that the set of constraints included in a metagrammar encoded by an XMG description all apply simultaneously. This is in contrast to transformational approaches, in which procedural rules are applied in a certain order to derive an analysis for an input sequence (as also highlighted by Lichte & Kallmeyer 2017: see Section 5.5 of this thesis). Second, XMG is notationally expressive to an extent that allows grammar authors to formalize theoretical notions. Crabbé et al. illustrate this expressivity with a series of examples demonstrating how syntactic and semantic analyses for a range of phenomena can be implemented with the help of XMG. Third, they point out that XMG is extensible, so that additional linguistic dimensions can be integrated into its computational architecture if required, and additionally, it can be used to describe grammars in different formalisms.

One such extension to XMG is the frame compiler for XMG developed by Lichte & Petitjean (2015: 207), which makes it possible to process a metagrammar that includes semantic frame descriptions. The TuLiPA parser is able to process descriptions in the `<frame>` dimension together with the other dimensions that are part of XMG (Arps & Petitjean 2018). For more on dimensions in XMG, see Section 5.3.2.

Crabbé et al. (2013: 594) describe a typical architecture of grammars generated with XMG metagrammars, which consists of three distinct layers in which information is stored. The first layer contains unanchored tree schemas that are grouped into tree families according to linguistically-relevant properties they share; the second layer contains morphological information; and the third layer contains a syntactic lexicon associating lemmas with tree families. The available realizations for each lemma and the tree families it belongs to are characterized by the various tree schemas included in those families. The anchor node in a tree schema denotes the node at which the lemma can anchor the tree.

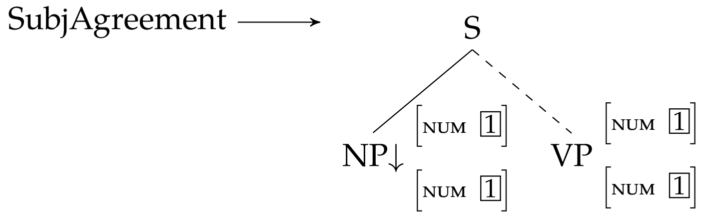


Figure 14: Co-indexation in an unanchored tree ensuring number agreement between a subject and a verb, as proposed by Crabbé et al. (2013: 598).

The organization of tree schemas into tree families makes it possible to specify structural information once and then use it in various different trees. Crabbé et al. (2013: 595) call this structure sharing. The term structure sharing is used in HPSG (Abeillé & Borsley 2021: 10) and LFG (Belyaev 2023: 12) to describe items which are co-indexed and share their identity (see also Section 4.3 of this thesis). An alternative way of describing the concept discussed by Crabbé et al. could be underspecification: tree schemas can be anchored by a number of different words. Which specific word appears during parsing will not impact the structure described in the tree schema, because that structure is shared across all possible anchors. This factorization reduces redundancy and makes grammar maintenance easier, as changes only need to be implemented once. For the purposes of this thesis, organizing partial descriptions in this way is useful to model constructions, which may share parts of their syntactic or semantic descriptions with other constructions. Factoring out these partial descriptions reduces the size of the metagrammar and makes it easier to make changes or add related classes later.

To illustrate the declarative way XMG can be used to describe phenomena like diathesis alternations, Crabbé et al. (2013: 596) contrast their proposed model with a more traditional, rule-based one, citing Flickinger (1987) who proposes handling redundancy in grammatical descriptions with the help of an inheritance hierarchy and a set of lexical rules. In the approach using lexical rules, one realization of an expression can be derived from another, presumably underlying one. In other words, this type of model requires the definition of a "base tree" from which other variants can be generated. Rules are applied sequentially, which means an ordering must be defined in the grammar to prevent derivation loops or the generation of ungrammatical structures.

With XMG, such a transformational approach is not needed. Instead of applying rules sequentially to derive one structure from another, structures are described in terms of constraints that they must fulfill. In this perspective, there are no base structures and no derived structures; instead, the descriptions consist of a series of tree fragments that are inserted into the different structures in different ways. This approach involves no movement or deletion of constituents, and no ordered application of rules.

Moreover, the constraint-based approach is also helpful for defining feature equality, for instance to ensure grammatical agreement between different constituents. Crabbé et al. present an example involving co-indexation of the relevant attribute values of different nodes in a tree fragment; the unification of feature structures then leads to a resulting tree in which these nodes match in the relevant attribute(s). In their example, a tree named *SubjAgreement* uses co-indexation to ensure agreement of certain attributes between a subject and a verb. Their example is shown in Figure 14.

The *SubjAgreement* tree is then incorporated via conjunction into a metagrammar class describing subjects with respect to other properties. Different realizations of subjects are made available in this class via disjunction. In the grammar described by this metagrammar, each possible realization of a subject thus also includes the constraints contributed by *SubjAgreement*. During parsing, the derivation of a tree for the input sequence will only be successful if the subject has the required number agreement with the verb and also fulfills all other constraints that are specified for subjects.

Beyond issues like case, number and gender agreement between nodes in a tree, co-indexation is also useful to connect descriptions across different metagrammar dimensions. Lichte & Petitjean (2015: 190) use co-indexation to link partial semantic frames to nodes in syntactic trees. For more on the integration of frame descriptions in XMG metagrammars, see Section 5.3.3 of this thesis.

### 5.3.2 *Dimensions in XMG metagrammars*

Crabbé et al. (2013: 600) provide a formal definition of the syntax of XMG and specify how the description language can be used to encode constraints in different dimensions. In their example, the three relevant dimensions are `<syn>` for syntactic information, `<sem>` for semantic information, and `<dyn>` for the syntax-semantics interface.

The `<syn>` dimension is where tree fragments are described using tree node variables, feature names, feature values, and feature variables. Relationships between nodes can be expressed with operators denoting node equality, dominance or precedence in the tree fragment. Feature structures can provide additional descriptions of nodes and can also be used to express equality of feature values between different nodes. All models that are licensed by a formula in the `<syn>` dimension fulfill the constraints expressed in that formula. XMG also provides a mechanism to share variables across classes.

Concerning the `<sem>` dimension, the specification given by Crabbé et al. (2013: 601) relates to flat semantics in the style of Gardent & Kallmeyer (2003: 125). As in the `<syn>` dimension, the constraints given in a formula license a set of valid models – in this case, a set of flat semantic formulas – which fulfill these constraints. While this type of semantic representation has seen some applications with TAG in the past, for the purposes of this thesis, the approach involving semantic frames is preferable, as demonstrated by e.g. Kallmeyer & Osswald (2013: 291). The merits of typed, recursive frames to represent semantics are discussed in Chapter 4 of this thesis. Section 5.4 will describe how Kallmeyer & Osswald model certain English constructions in a metagrammar involving semantic frame descriptions.

The `<dyn>` dimension is responsible for providing a mapping between syntactic and semantic elements. Interface features enable a coreference between a variable in the syntactic dimension and the variable referring to its corresponding element in the semantic dimension. The `<dyn>` dimension as described by Crabbé et al. (2013: 601) corresponds to the `<iface>` dimension typically used in implementations with XMG (Simon Petitjean, p.c.).

Crabbé et al. (2013: 602) point out that XMG classes can inherit other classes, are identified with a name, and can export their own variables in order to be imported elsewhere in turn. Only exported variables are visible when the class in which they appear is imported; all other variables are local to the class definition in which they are located. This allows one class to instantiate another one multiple times without the nodes that are included in the instantiated class necessarily being identical. The different dimensions mentioned above can all appear as part of a class description in XMG, although each dimension may be empty for any given class. Finally, the description of a class can include a disjunction or conjunction of other classes. Together with the mechanism for exporting variables, this is how XMG allows a factorization of the grammar.

Using the example of a large feature structure based tree adjoining grammar for French called SemTag, Crabbé et al. (2013: 612) model a suggested approach to using XMG to create a grammar. The strategy described by the authors will largely be followed in the following chapters in this thesis part, although the `<sem>` dimension will be replaced by the `<frame>` dimension describing semantic frames, instead of flat semantic representations.

For the `<syn>` dimension, the authors begin by defining tree fragments that represent either a possible realization of a verb argument or a possible realization of a verb. These tree fragments are then used as building blocks to build the different elementary TAG trees. A hierarchical structure with four levels makes maximal factorization possible. The first level contains the basic tree fragments encoding verb or verb argument realizations. On the second level, grammatical functions are defined as disjunctions of argument realizations. The third level contains definitions of diathesis alternatives, in the form of conjunctions of verb realizations and grammatical functions. The final, fourth level gathers diathesis alternatives into tree families.

Concerning the first level of description in the `<syn>` dimension, the authors describe how SemTag organizes the tree fragments in an inheritance hierarchy to enable maximal sharing of common information. The most basic tree fragments are concerned with representing different realizations of, for instance, subjects or objects, as well as verb realizations (e.g. in different tenses or moods).

On the second level, syntactic function names are used to group alternative ways to express each type of verbal argument. For instance, a *Subject* is defined as a disjunction of two specific subject realizations that are defined on the first level, named *CanonSubj* and *RelatSubj*.

The third level is used to represent diathesis alternatives, by referring to the abstractions provided on the second level. Crabbé et al. (2013: 615) assume that the observed form of a verb constrains how its predicate arguments are realized in syntax. The abstractions encoded on this third level contain descriptions in which the same predicate arguments – as described by the classes in the second level – are combined with different verb realizations.

The fourth level is concerned with tree families, which capture alternative realizations of a given verb type. Families can express subcategorization frames.

An attempt at a visualization of the hierarchical structure of the example metagrammar described by Crabbé et al. (2013: 612–616) is shown in Figure 15 on page 66. Each class is located at one of the four levels of the hierarchy. Classes are connected to each other based on conjunction and disjunction. For the sake of readability, the class *TransitiveDiathesis*, which is defined by Crabbé et al. as a disjunction of three different conjunctions of various classes, is presented in two parts here. *TransitiveDiathesis* consists of a disjunction of the added classes **TransitiveDiathesis1*, **TransitiveDiathesis2* and **TransitiveDiathesis3*, and each of these three classes consists of a conjunction of the relevant classes on other levels of the class hierarchy.

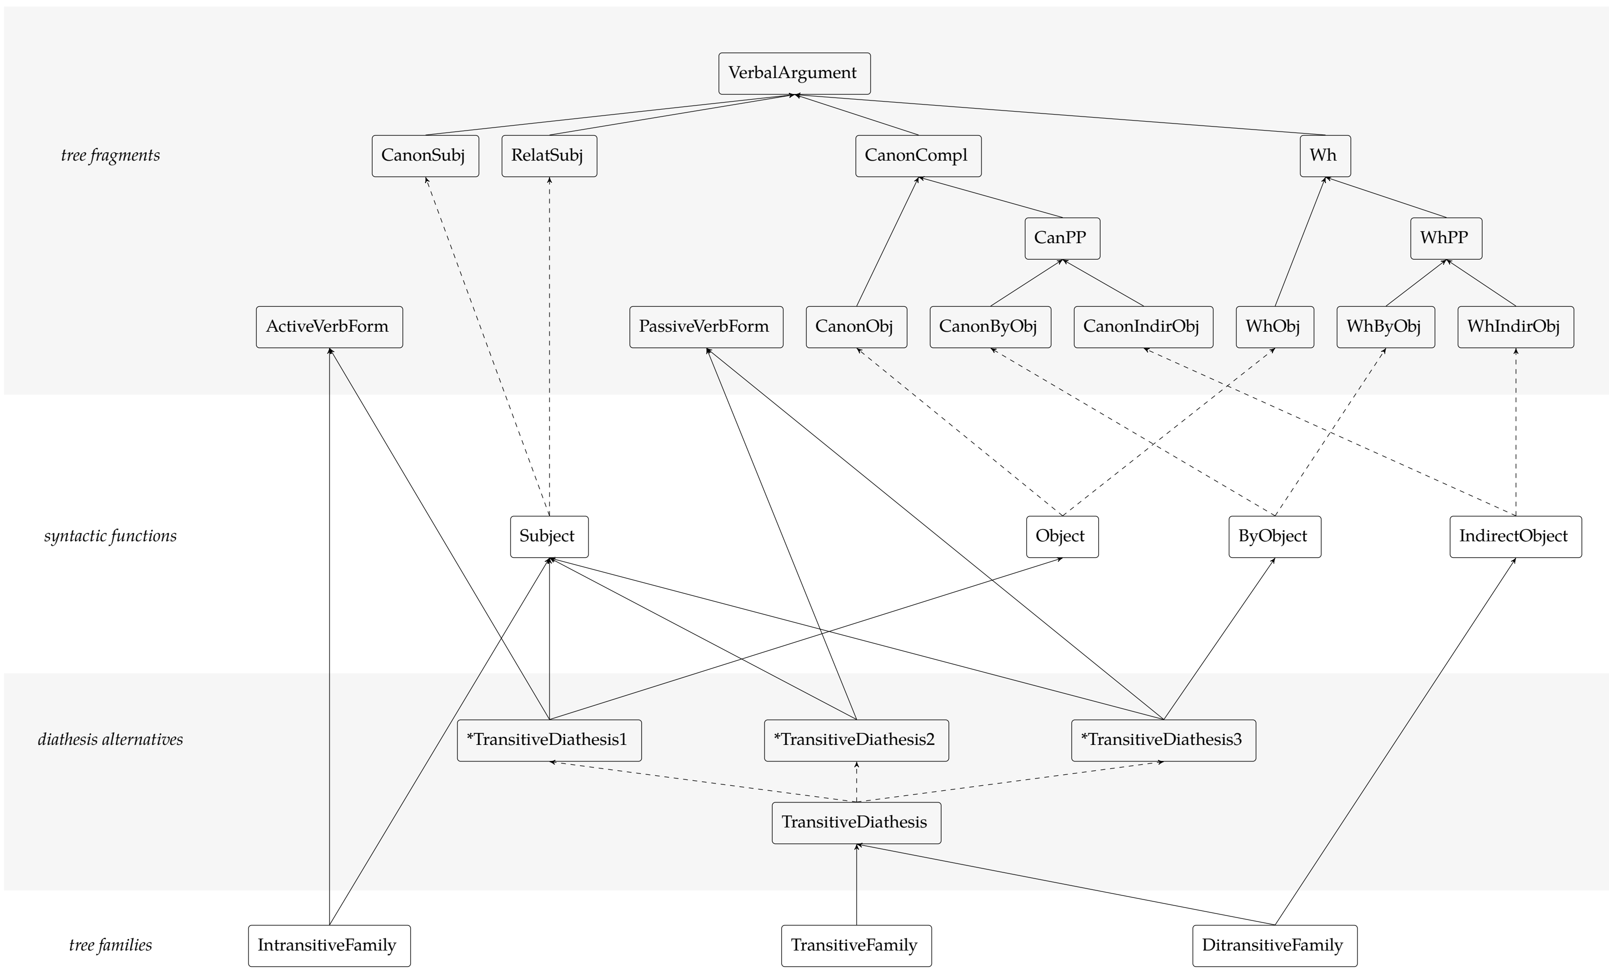


Figure 15: Visualization of the hierarchical structure of the metagrammar described by Crabbé et al. (2013: 612–616). Solid arrows denote structure inheritance. Dashed arrows denote disjunction.

Crabbé et al. (2013: 617) proceed to detail how semantic information can be added to an XMG metagrammar created in the way described above, by associating each elementary tree with a corresponding semantic representation and ensuring with interface features that the different parts of a semantic representation are associated with the correct nodes in the tree. Since the approach pursued in this thesis will rely on semantic frames instead of flat semantic representations, the inclusion of semantics in the metagrammars presented here will not make use of the `<sem>` dimension, but instead rely on the `<frame>` dimension that was added to XMG by Lichte & Petitjean (2015: 207). The following section is concerned with describing the `<frame>` dimension.

### 5.3.3 *The `<frame>` dimension in XMG*

Lichte & Petitjean (2015: 207) present an extension to XMG that allows grammar authors to add a `<frame>` dimension to the metagrammar. Like the `<sem>` dimension used by Crabbé et al. (2013), the `<frame>` dimension also represents semantics and makes it possible to link (partial) semantic descriptions to (partial) tree descriptions in the `<syn>` dimension. Descriptions in the `<frame>` dimension take the form of typed, recursive feature structures, while feature structures describing elements of syntactic trees are untyped. The benefits of a frame-based representation of meaning have been discussed in Chapter 4 of this thesis.

The work by Lichte & Petitjean is based in part on earlier work by Kallmeyer & Osswald (2013), who develop an approach to jointly modeling syntactic and frame-semantic descriptions of constructions in an XMG metagrammar for a lexicalized tree adjoining grammar. The frame formalization proposed by Kallmeyer & Osswald has been described in Section 4.4 of this thesis. Frames following this formalization are not required to have a single root node, but can have multiple base nodes; substructures can be unified; reference to base-labeled nodes is sufficient to completely characterize a feature structure; and frame descriptions can involve non-functional relations. Section 5.4 of this thesis will discuss the XMG implementation of Kallmeyer & Osswald in more detail.

Lichte & Petitjean (2015: 187) and Kallmeyer & Osswald (2013) are concerned with interactions between lexical and constructional meaning, through which the overall meaning of a sentence comes about.[1] The phenomena covered by Kallmeyer & Osswald are directed motion expressions in English, like *John rolled the ball into the goal*, and the English dative alternation. The aim of the `<frame>` dimension implemented by Lichte & Petitjean (2015: 192) is to allow the factorized implementation of approaches like the one described by Kallmeyer & Osswald. Descriptions of lexical items include descriptions of their lexical frames. Unanchored elementary trees are associated with partial semantic frames, whose attribute values are unified with the lexical frames of lexical elements that are involved in the derivation for a given input sequence.

The `<frame>` dimension implemented by Lichte & Petitjean (2015: 202) also relies on the inheritance mechanisms provided by XMG, as described by Crabbé et al. (2013). In the context of adding frame descriptions to the metagrammar, inheritance can be

1 Although Construction Grammar allows verbs and other words to be regarded as constructions themselves, it is common in work in the XMG and LTAG area to distinguish between lexical and constructional meaning. Meanings of words are stored in a lemma lexicon, while meanings of constructions are added to the syntactic dimension of the grammar.

used to factor out partial syntactic or semantic descriptions into separate classes, which are then instantiated by other classes to form combined tree and frame descriptions. Exporting node variables that appear in a class makes the variables and their values accessible to classes inheriting from that class.

The linking between descriptions in the `<syn>` dimension and the `<frame>` dimension (and other dimensions) is based on variables that are local to each metagrammar class. Variables that are shared across dimensions constitute "a direct interface between otherwise separated dimensions" (Lichte & Petitjean 2015: 204). Assigning one variable name at the same time to a node in the `<syn>` dimension and to a frame in the `<frame>` dimension thus links that node to that frame.

TAG's tree composition operations identify certain nodes of the trees that are being composed with each other, as shown in the example in Figure 13 on page 59. When these nodes are associated with descriptions in the <frame> dimension, the identification of the nodes also leads to a unification of their associated frames. For instance, in the context of one tree being added to another tree at a specific node via substitution, the frame linked to the root node of a substituting tree will be unified with the frame of the other tree's substitution node.

Lichte & Petitjean (2015: 204) also discuss whether the `<sem>` dimension could be used to describe frames instead of the flat semantic descriptions it was originally designed for. They argue that the `<sem>` compiler in XMG allows, for instance, the coexistence of multiple instances of one predicate with conflicting values, which is not desired for the frame description; multiple appearances of the same predicate (applied to the same entity) should either be unifiable or lead to the rejection of the resulting structure. Beyond this, the typed nature of semantic frames can also not be represented as accurately by flat descriptions as it can by typed feature structures. Instead of updating the `<sem>` dimension to prevent these issues, the authors opt to create a new dimension called `<frame>`, which is specifically designed to represent frame information and to fulfill all formal requirements that come with the frame approach.

The `<frame>` dimension implemented by Lichte & Petitjean (2015: 209) makes use of two new global fields named `frame_types` and `frame_attributes`, which are responsible for specifying the available attributes, types, and base labels for frames. These global fields are defined outside the metagrammar classes, and can be referenced by the descriptions in the `<frame>` dimension of the classes to specify the signature of each frame. A third global field named `frame_constraints` is responsible for encoding subtype relationships, incompatibilities between frame types, and constraints on the existence or on the values of specific attributes in structures of a particular type.

The specification language developed by Lichte & Petitjean (2015: 208) to describe feature structures is reproduced in Listing 1 on page 69. As the authors note, their implementation of the `<frame>` dimension is oriented on existing XMG syntax. The set of descriptions contained in the `<frame>` dimension of an XMG class is unordered, and an arbitrary number of type expressions is possible to allow for conjunctive types.

```
<frame>{Descriptions;Descriptions;...}
Descriptions ::= var? '[' Description (',' Description)∗ ']'
Description  ::= type | PathEquation | AVPair
PathEquation ::= attr+ '=' var? attr+ (':' Value)?
AVPair ::= attr+ ':' ValueValue  ::= var | type | Descriptions
```

Listing 1: Description language for feature structures from Lichte & Petitjean (2015: 209).

$$\boxed{0}\begin{bmatrix} \textit{causation} & \\ \text{ACTOR} & \boxed{1} \\ \text{THEME} & \boxed{2} \\ \text{CAUSE} & \begin{bmatrix} \textit{activity} & \\ \text{ACTOR} & \boxed{1} \\ \text{THEME} & \boxed{2} \end{bmatrix} \\ \text{EFFECT} & \boxed{4}\begin{bmatrix} \text{MOVER} & \boxed{2} \\ \text{GOAL} & \boxed{3} \end{bmatrix} \end{bmatrix}$$

```
<frame>{
?0[causation,
    actor: ?1,
    theme: ?2,
    cause: [activity,
            actor: ?1,
            theme: ?2],
    effect: ?4[mover:?2,
               goal: ?3]
]}
```

Figure 16: Example: frame description in XMG for the frame component of a transitive tree with a directional prepositional phrase, as proposed by Lichte & Petitjean (2015: 210).

Their Figure 11 (reproduced here as Figure 16) demonstrates how a frame represented as a typed attribute-value matrix can be described in their <frame> syntax. Base labels are written as boxed numbers in the frame and as variables prefixed with ? in the code.

Their Figure 12 (Lichte & Petitjean 2015: 211, reproduced here as Figure 17 on page 70) illustrates how base labels for the <frame> dimension can be represented with XMG variables and thus implement the linking between the <frame> dimension and other dimensions. The class shown in the figure contains descriptions in the <syn> and <frame> dimensions. The <syn> dimension is described in terms of various nodes with variable names (e.g. ?S), attributes and their values (e.g. [cat=s]), optional node marks (e.g. (mark=anchor)), and a set of operators indicating dominance (->∗), immediate dominance (->), and immediate precedence (>>). A fourth operator which does not appear here is available to indicate precedence (>>∗).

The syntax for descriptions in the global field frame_constraints is reproduced in Listing 2 on page 70. Each constraint has the form of a logical implication. Each side of a constraint can contain a set of types, the specification of a path identity in a frame, or an attribute with a specific value.

If the type hierarchy allows it, two frame descriptions belonging to different types can be unified so that the result belongs to a conjunctive type, that is, it satisfies the conditions imposed on each of its constituent types. Lichte & Petitjean point out that only a subset of conjunctive types may be valid given the description of a type hierarchy; they proceed to describe how the set of valid conjunctive types can be determined (Lichte & Petitjean 2015: 215). In case no constraints are given, the set of valid conjunctive types is equal to the power set of elementary types. Otherwise, if constraints apply, the constraints are used to filter down the power set of elementary types, leaving only valid

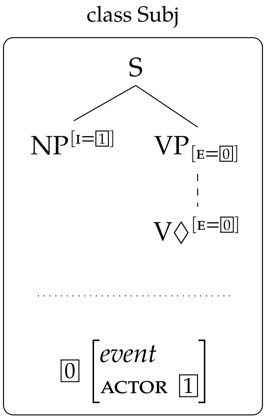


```
class Subj
...
<syn>{
   node ?S [cat=s];
   node ?SUBJ [cat=np,
            top=[i=?1]];
   node ?VP [cat=vp,bot=[e=?0]];
   node ?V (mark=anchor)
         [cat=v,top=[e=?0]];
   ?S->?SUBJ; ?S->?VP; ?VP->*?V;
   ?SUBJ»?VP
}
<frame>{
?0[event,
   actor:?1]
}
...
```

Figure 17: Example: metagrammar class with `<frame>` description, illustrating the representation of base labels in the `<frame>` dimension with XMG variables, as proposed by Lichte & Petitjean (2015: 211).

types in the set. This is achieved either through a top-down filtering approach that generates bit vector patterns for all non-valid conjunctive types, or through a more efficient matrix-based approach. In the context of this thesis, constraints will be necessary to block the unification of certain descriptions, for instance in the form of specifying a verb to be incompatible with a particular construction via type membership. In other words, any two or more types from the hierarchy can conceivably be combined into a conjunctive type, unless a constraint exists preventing that combination.

```
frame_constraints = {Constraint,Constraint,...}
Constraint ::=
  %% type constraint
  type+ '->' type+ |
  %% appropriateness condition
  type+ '->' Descriptions+  |
  %% feature-value constraint
  ('[' (AVPair|PathEquation) (',' AVPair|',' PathEquation)* ']')+
    '->' Descriptions+
Descriptions ::= '[' Description (',' Description)* ']'
Description  ::= type | PathEquation | AVPair
PathEquation ::= attr+ '=' attr+ (':' Value)?
AVPair ::= attr+ ':' Value
Value  ::= type | Descriptions
```

Listing 2: Description language for frame constraints from Lichte & Petitjean (2015: 212).

The added `<frame>` dimension in XMG developed by Lichte & Petitjean (2015) thus provides all the functionality that is required to implement metagrammars involving frame-semantic descriptions, like those proposed by Kallmeyer & Osswald (2013). Modeling syntactic structures and semantic frames jointly in this way is ideal for representing phenomena at the syntax-semantics interface, such as the alternations and constructions this thesis part is concerned with.

### 5.3.4 *Compiling XMG metagrammars*

Before a metagrammar written in XMG can be used to parse input sentences, it must be compiled into a grammar. Different specialized compilers are available for the different dimensions of XMG descriptions. In this thesis part, the following compilers will be used.

Lexical entries are described in the `<lemma>` dimension, which will be compiled with the `lex` compiler. Each lemma entry is linked to a frame, a syntactic category, and a set of trees that can be anchored by it.

Inflected forms of lemmas are described in the `<morpho>` dimension in terms of their morphological features and will be compiled with the `mph` compiler. Each entry in this file is linked to its corresponding entry in the `lemma.mg` file and to a syntactic category.

This separation of morphological and lexical information is yet another factorization that prevents redundancy (Kallmeyer & Osswald 2013: 295): morphological properties of inflected forms, such as case, number and gender features, are stored in the `<morpho>` dimension, while lexical information, such as the semantic frame expressing the meaning of the lemma, is stored in the `<lemma>` dimension, and `<morpho>` entries point to `<lemma>` entries. During derivation, each form that is encountered in the input sequence can be looked up in the `<morpho>` resource, its morphological features can be read, and its semantics can be looked up in the corresponding `<lemma>` entry in the lexical resource.

The frames representing the lexical semantics of each lexicon entry are described in the `<frame>` dimension and will be compiled with the `synframe` compiler. The descriptions in this dimension represent frames in the form described by Lichte & Petitjean (2015: 210). The `<iface>` dimension makes it possible to link frames to nodes in syntactic descriptions. Whenever a lemma with a description in the `<frame>` dimension anchors a tree, the lexical frame for that lemma is stored in the interface, and the anchored class can access the interface to assign the frame to one or more nodes in the tree.

An important property of the interface dimension is that no structural isomorphism between syntactic and semantic components is required (Kallmeyer & Osswald 2013: 270): how the lexical frames will be incorporated into constructional frames is described at the level of construction descriptions. This allows constructions to, for instance, either unify their entire frame with the frame contributed by an anchoring verb, or unify the verb's frame with some substructure of the constructional frame. The advantages of this flexibility will become apparent in Sections 6.6 and 7.5, which describe how the semantics of the various constructions covered in the metagrammar developed in this thesis will be modeled.

The type inventory for the frames in each metagrammar will be described in a separate frame hierarchy file, which will be compiled using the `synframe` compiler. The type hierarchy contains an exhaustive list of all types that appear in the grammar, as well as a set of constraints describing subtype relations, incompatibility relations and attribute constraints for individual types (Lichte & Petitjean 2015: 215). Note that all types that are not explicitly specified to be incompatible with each other can form conjunctive types during derivation.

Finally, descriptions of all elementary constructions and elementary trees without semantic descriptions are described in a subset of the dimensions `<syn>`, `<frame>` and `<iface>`, and will be compiled with the `synframe` compiler.

Descriptions in the `<syn>` dimension specify properties of nodes and node ordering constraints that describe the tree structure of a class. Underspecifying the order of nodes results in the grammar allowing more than one syntactic structure for the class. Nodes can be marked with certain symbols denoting specific functions such as substitution nodes, null adjunction, or foot nodes in auxiliary constructions. Nodes can also be described further with untyped feature structures defining values for certain properties.

The XMG compiler uses the logical programming language Prolog to solve the descriptions in a given metagrammar (Crabbé et al. 2013: 604). The solving of descriptions in XMG is implemented as an extensible definite clause grammar (EDCG, Van Roy 1990) in the form of dedicated accumulators, or pairs of unification variables, per dimension. An accumulator is a device used in Prolog that can hold intermediate results (Blackburn, Bos & Striegnitz 2006: Chapter 5.3) – in this context, it will accumulate an arbitrary number of types of literals, each of which represents formulas of a description language. The function of an interpreter for a specific dimension is to take the literals in that dimension's accumulator and solve the constraints expressed by these literals by computing models that satisfy the given formulas.

Even with dedicated accumulators per dimension, the unification variables can still share information across dimensions. After the EDCG execution, each class in the metagrammar, or axiom, yields a tuple of descriptions for each dimension of the given class, according to the given constraints. In the next step, the valid models for each dimension-specific description list are computed by a solver for that dimension. Crabbé et al. (2013: 606) distinguish, for instance, the `<syn>` solver, which yields a set of trees that are minimal models of the description, and the `<sem>` solver, which yields a set of flat semantic descriptions; in the case of the `<frame>` dimension, this solver yields a set of minimal models for the semantic frames that fulfill the constraints contained in the description of that dimension. By taking unification of variables into account, the solvers generate only models that are able to satisfy all dimension-specific constraints as well as all unification constraints.

Lichte & Petitjean (2015: 207) detail how XMG compilers function. XMG compilers are composed of modular parts that are called bricks. They are implemented in YAP (Yet Another Prolog) and use Gecode bindings for solving constraints. Automatic code generation with Python ensures extensibility.

Lichte & Petitjean (2015: 214) proceed to explain how an XMG compiler generates representations based on the descriptions provided in the `<frame>` dimension. The resulting representations differ from those generated for the `<syn>` dimension in that the latter can be underspecified, which means that a constraint solving procedure is required to compute minimal models that satisfy the given constraints; since the `<frame>` dimension contains no underspecification in that sense, the constraint solving step is not necessary in the compiler for that dimension.

The representations generated by the XMG compiler for descriptions in the `<frame>` dimension take the form of attributed variables in Prolog and consist of two fields: the first one encodes the type, the second one contains an association list with key-value pairs that represents the attribute values and relationships expressed in the underlying description. As Lichte & Petitjean (2015: 214) point out, the association list data structure in Prolog allows values to be attributed values themselves and thus enables recursion.

The unification mechanism of the compiler takes both fields, the type and the feature structure, into account in order to determine how two descriptions in the `<frame>` dimension can be unified. The type attribute undergoes type unification, while the association lists undergo set union. The resulting structures must be well-formed in order for a model to be produced based on the given descriptions. Since the `<frame>` compiler does not search for minimal connected models that satisfy the given constraints, variable equations must be used to specify where unification should take place.

The ability to parse sentences based on LTAG metagrammars involving semantic descriptions in the `<frame>` dimension was added to the TuLiPA parser by Arps & Petitjean (2018). The parser takes grammars compiled from metagrammars with XMG-2 (Petitjean, Duchier & Parmentier 2016) as input, and outputs for each input sentence the derived tree(s) and frame structure(s), as well as the trace of tree fragments involved in the derivation and an overview of the contents of the interface dimension.

## 5.4 JOINTLY MODELING SYNTAX AND SEMANTICS WITH LEXICALIZED TREE ADJOINING GRAMMAR

Different strategies for including semantics in TAG descriptions have been proposed over the years (e.g., Kallmeyer & Joshi 2003, Gardent & Kallmeyer 2003, Kallmeyer & Romero 2008). An approach involving frame-semantic descriptions in metagrammars was first presented by Kallmeyer & Osswald (2013). They propose a metagrammar that models certain constructions jointly on the syntactic and the semantic level. This allows them to jointly derive the syntactic structure and the frame-semantic representation for the types of input sentences they are concerned with.

In Kallmeyer & Osswald's approach, syntactic and semantic (partial) descriptions are linked to each other using the mechanisms provided by XMG. The elementary parts of the semantic frames do not necessarily correspond to the individual parts of the elementary trees that they are associated with; the interface dimension is used in their metagrammar to explicitly state which components of a frame are linked to which nodes in the relevant tree.

### 5.4.1 *Creating LTAG metagrammars with frames*

Kallmeyer & Osswald (2013) define a set of elementary constructions that each consist of an elementary LTAG tree and a decompositional frame. These classes can be used to derive compositional LTAG trees and frames for sentences instantiating these elementary constructions. The meaning of an input sentence is derived based on the interactions between lexical meaning and constructional meaning. This factorized approach is illustrated by Kallmeyer & Osswald with regard to directed motion expressions and the English dative alternation, but is so flexible that it can be applied to other phenomena at the syntax-semantics interface as well.

Kallmeyer & Osswald (2013: 324) motivate the combined use of LTAG and frame semantics with the ability of each of these frameworks to encode constraints that can be unified to derive all valid analyses for a given sequence, based on the lexical and compositional rules encoded in the metagrammar. The extended domain of locality of LTAG is well-suited to expressing relationships between different constituents in a tree, while the process of adjunction makes it possible to represent long-distance dependencies in derived trees that are the result of recursive or iterative effects of local dependencies from the elementary trees involved in the derivation.

The derivation of frames for input sequences in Kallmeyer & Osswald's grammars occurs jointly with the derivation of trees, as elementary constructions are composed via substitution and adjunction. Each elementary tree can be paired with a complex semantic representation whose constituent parts can be mapped to the constituent parts of the tree if desired, but whose overall content does not need to be fully compositionally derived from the nodes in the tree. Then, as trees are composed, the frames associated with the relevant nodes are unified.

By describing in the metagrammar both the syntactic and semantic properties of the fragments appearing in the grammar, as well as rules for their respective composition, Kallmeyer & Osswald (2013: 268) are able to process phenomena at the syntax-semantics interface that involve interactions between lexical and constructional meaning. Since those interactions also play a role in the phenomena to be modeled in this thesis, the general approach of Kallmeyer & Osswald will be followed for the creation of the metagrammars to be developed here.

Kallmeyer & Osswald use frames as semantic representations that capture the meaning of lexical items in the form of typed feature structures. The formalization they provide for the semantic frames they assume is summarized in Section 4.4 of this thesis.

In the implementation by Kallmeyer & Osswald (2013: 278), elementary morphosyntactic trees are paired with their corresponding elementary meaning structures. This linking is made possible by interface features, which are meta-variables that are included in the feature structures associated with the nodes of a given tree. The individual components of the semantic frame for an input sequence are contributed by the frames provided alongside the different trees that are involved in the derivation for that sequence.

The attribute values of the elementary semantic frames are partly underspecified. In an example derivation for the sentence *John eats pizza*, Kallmeyer & Osswald (2013: 291) use an initial tree-frame pair anchored by the sentence's verb, *eats*, where the semantic frame for *eating* contains attributes for the ACTOR and THEME roles, but no values are assigned to these attributes in that initial frame. Substituting other elementary constructions in the subject and object position of the verbal initial tree allows the frames associated with the verbal arguments *John* and *pizza* to unify with their respective attribute slots in the frame provided by the verb. The derivation presented in Figure 10 from Kallmeyer & Osswald (2013: 291) is reproduced here as Figure 18 on page 75.

Interface features are used by Kallmeyer & Osswald (2013: 290) to link partial semantic frames to nodes in elementary trees. Their interface feature E stores the event frame contributed by the verb. This interface feature allows different tree fragments that are involved in a derivation to specify different parts of the semantic frame for the input sentence, which are unified following the unification rules of substitution and adjunction.

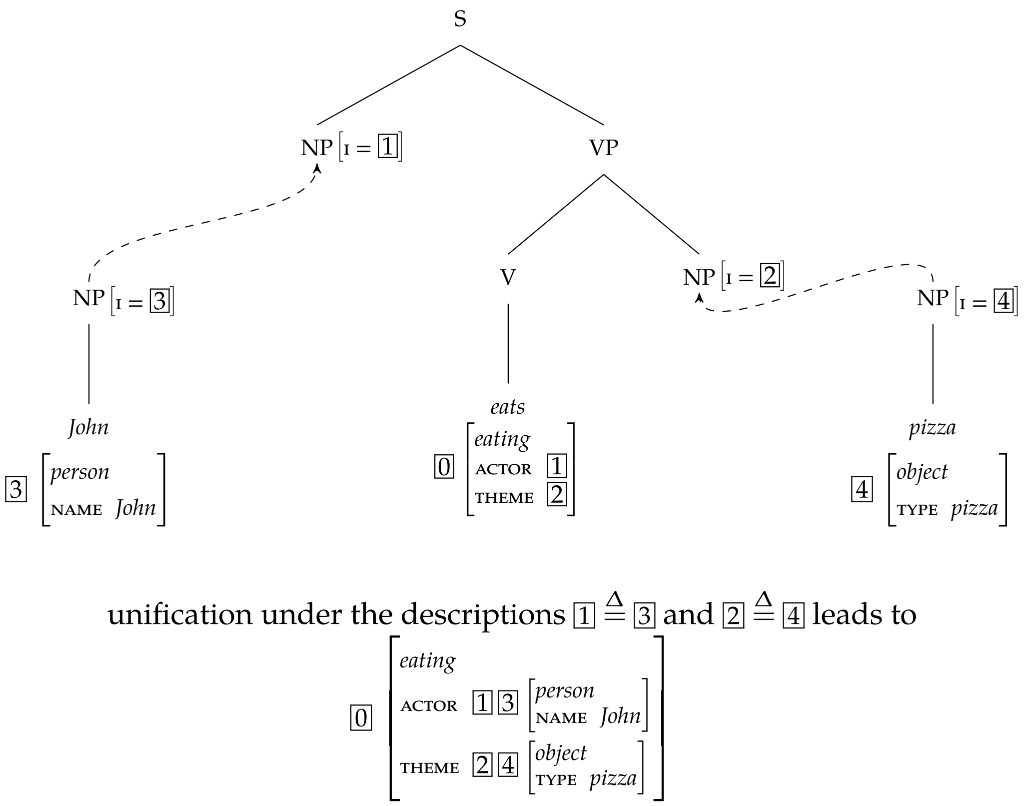


Figure 18: Derivation for the sentence *John eats pizza* as proposed by Kallmeyer & Osswald (2013: 291).

In the metagrammar by Kallmeyer & Osswald (2013: 296), the possible syntactic realizations of a verb's semantic arguments are specified not by the verb's lexicon entry, but by the tree anchored by the verb and the trees that can be combined with it via substitution. This is a crucial advantage with respect to the modeling of alternations and other phenomena in which the same set of arguments can be expressed in different syntactic positions in the different available constructions. Therefore, the metagrammars presented in this thesis will also follow this general architecture.

### 5.4.2 *Kallmeyer & Osswald's application I: Directed motion expressions*

Kallmeyer & Osswald illustrate their proposal with two applications focusing on different linguistic phenomena. The first application (Kallmeyer & Osswald 2013: 297) is concerned with directed motion expressions in English. In these expressions, parts of the directed motion meaning are contributed by verbs of motion, and other parts are contributed by the directional prepositional phrases that they appear with. While some verbs of motion encode the manner of motion, but no path-related information, others encode the direction of the motion, but no manner-related information. The authors focus on manner-encoding verbs. They also consider verbs of transport and caused motion, in which the entity that moves is not identical to the entity that causes the motion. Like manner-encoding motion verbs, these verbs also do not lexically contribute a direction, but can be combined with directed prepositional phrases that do

$$
\textit{walk}\;\boxed{0}\begin{bmatrix} \textit{locomotion} & \\ \textsc{actor} & \boxed{u} \\ \textsc{mover} & \boxed{u} \\ \textsc{manner} & \textit{walking} \\ \textsc{path} & \begin{bmatrix} \textit{path} & \\ \textsc{startp} & \top \\ \textsc{endp} & \top \end{bmatrix} \end{bmatrix}
\qquad
\textit{throw}\;\boxed{0}\begin{bmatrix} \textit{onset-causation} & \\ \textsc{cause} & \begin{bmatrix} \textit{activity} & \\ \textsc{actor} & \top \\ \textsc{theme} & \boxed{u} \\ \textsc{manner} & \textit{throwing} \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{translocation} & \\ \textsc{mover} & \boxed{u} \\ \textsc{path} & \begin{bmatrix} \textit{path} & \\ \textsc{startp} & \top \\ \textsc{endp} & \top \end{bmatrix} \end{bmatrix} \end{bmatrix}
$$

Figure 19: Semantic frames for *walk* and *throw* events, as proposed by Kallmeyer & Osswald (2013: 301).

contribute this meaning. The authors use the term *translocation* to refer to motion events that involve a trajectory, trace, or path. Different aspects of a translocation event, e.g. its direction, goal, or source, can be expressed by chaining different prepositional phrases together.

Whether the directional prepositional phrases that appear with the verbs in question take the function of a complement or an adjunct matters with regard to their treatment in the LTAG grammar, as Kallmeyer & Osswald note. These different functions need to be modeled differently: obligatory elements should show up in the lexical entry for a verb, while optional elements will contribute their meaning on the constructional level instead. Kallmeyer & Osswald (2013: 299) take the perspective of Gehrke (2008: 213), in which bounded directional prepositional phrases are complements and unbounded ones are adjuncts.

The directed motion expressions discussed in this first example illustrate how useful the frame-semantic approach to representing meaning is in this constructional context. For expressions involving motion-causing verbs, the lexical meaning contributed by the verb corresponds to a *causation* type event with two subevents, a CAUSE (the action performed by the ACTOR) and an EFFECT (the *translocation* event undergone by the MOVER in response to the CAUSE). Kallmeyer & Osswald (2013: 301) present semantic frames for the example verbs *walk* and *throw*, which are reproduced here as Figure 19. In these frames, the attributes ACTOR and MOVER are not filled with specific values yet, since they need to be filled by concrete verb arguments. Boxed letters are used instead of base labels here to signify path identity. For instance, in the *walk* frame, the ACTOR and MOVER attributes do not refer to base-labeled nodes, but they must have the exact same value. The attribute PATH exists in the frames for both *walk* and *throw*, but its nature is not specified lexically in the case of these manner-encoding verbs; for path-encoding verbs like *enter* or *leave*, a certain value for this attribute would be included in their lexical frames.

The frames for *walk* and *throw* differ from each other in the co-indexation of different participants of the respective events described by the verbs. *Throw* expresses a caused motion, while *walk* expresses a locomotion that is not externally caused; the ACTOR and MOVER of the internally-caused locomotion refer to the same entity and are thus co-indexed in the frame. In the caused-motion frame, the ACTOR is not the entity that moves. Instead, the causing event has a THEME argument which is identical to the MOVER of the resulting *translocation* event.

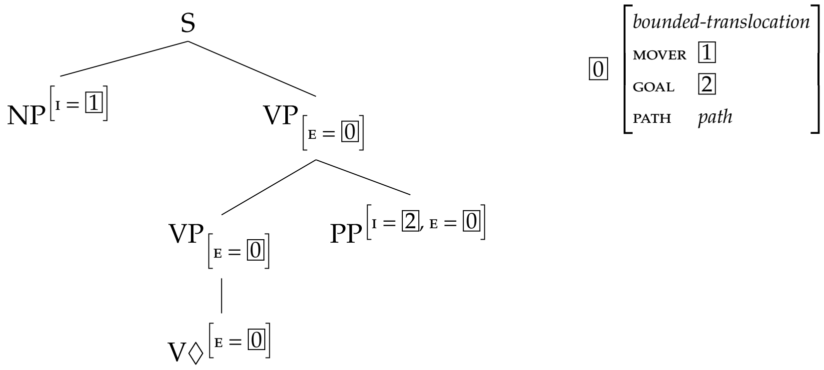


Figure 20: Unanchored tree and frame for intransitive manner-encoding verbs of motion, as proposed by Kallmeyer & Osswald (2013: 304).

Kallmeyer & Osswald (2013: 302) propose semantic frames for a set of directional prepositions. These prepositions differ with regard to whether they contribute information about the trajectory, goal, end point, or some other aspect of a path. Their frames underspecify the *translocation* event that they are associated with, but contain specific values describing the path. As a result, a unification of the verb frames and the prepositional frames should lead to complete frames in which the translocation event as well as its path are fully specified to the extent possible based on the structures involved in the derivation.

The next step in Kallmeyer & Osswald's development of an LTAG metagrammar for the constructions in question is the creation of a frame type hierarchy (Kallmeyer & Osswald 2013: 303). That hierarchy contains all event types that appear in the metagrammar, and describes them in terms of subtype relationships, attribute constraints, and mutual incompatibility constraints. The type hierarchy and the constraints expressed in it need to be included in the metagrammar because the information contained in it determines whether an attempted unification of two partial frame descriptions will be successful or not.

As mentioned above, the authors treat prepositional phrases that cannot be iterated as complements, and other prepositional phrases as adjuncts. This difference plays a role in the design of the unanchored tree for the verb and the elementary trees that are available to anchor it: complements are obligatorily integrated into the unanchored tree and are not part of the lexical elementary tree for a verb. Adjunct prepositional phrases are instead integrated into the final derived tree and frame based on constructional elementary trees that are co-anchored by the preposition, and which are unified with the semantics of the verb via adjunction.

Kallmeyer & Osswald (2013: 304) give an example of an unanchored elementary tree for an intransitive manner-encoding motion verb, which is reproduced here as Figure 20. The interface features `E` and `I` are responsible for integrating the semantic frames for the verb and its arguments into the derived structure in the appropriate way. The tree in the example is linked to a frame of the type *bounded-translocation*, and the `PP` node is associated with the goal attribute of the frame via co-indexation. The two `VP` nodes are both associated with the event frame for this unanchored tree, which makes it possible to add modifying trees via adjunction whose meaning contributions are then unified with that event frame.

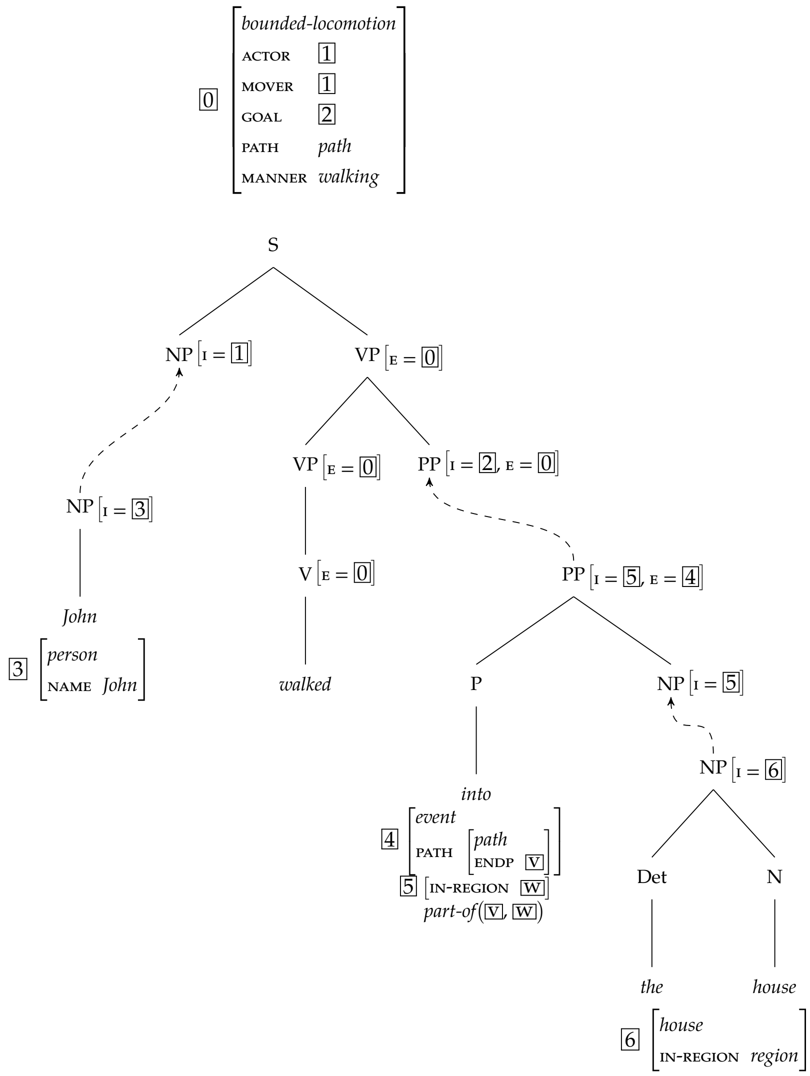


Figure 21: Derivation for the sentence *John walked into the house*, as proposed by Kallmeyer & Osswald (2013: 305).

Since the unanchored tree shown in Figure 20 contains a PP node that is not an anchor node, and final derived trees can only have leaf nodes labeled with terminal symbols, an elementary tree must be substituted at this position to obtain a complete derivation. The same holds for the NP node. The information from the frame of the elementary tree that is substituted at the NP node is unified with the interface feature index $\boxed{1}$, while the contribution of the PP is unified with the interface feature index $\boxed{2}$. In this unanchored construction, the element with the index $\boxed{1}$ is assigned to the MOVER attribute of the event frame; in other words, this construction describes scenarios in which the syntactic subject is the semantic MOVER. The element with index $\boxed{2}$, which will be filled with a concrete value according to the PP tree that is substituted at the PP node, expresses the GOAL of the event frame.

Kallmeyer & Osswald (2013: 305) use this unanchored tree and frame, together with lexical elementary trees, to derive a tree and frame for the sentence *John walked into the house*. The derivation is reproduced here as Figure 21.

$$
\boxed{0}\begin{bmatrix} \textit{bounded-locomotion} & \\ \text{ACTOR} & \boxed{1}\begin{bmatrix}\textit{person} & \\ \text{NAME} & \textit{John}\end{bmatrix} \\ \text{MOVER} & \boxed{1} \\ \text{GOAL} & \boxed{2}\begin{bmatrix}\textit{house} & \\ \text{IN-REGION} & \boxed{W}\end{bmatrix} \\ \text{PATH} & \begin{bmatrix}\textit{path} & \\ \text{ENDP} & \boxed{V}\end{bmatrix} \\ \text{MANNER} & \textit{walking} \end{bmatrix}
$$

$$\textit{part-of}(\boxed{V}, \boxed{W})$$

Figure 22: Derived semantic frame for the sentence *John walked into the house*, as proposed by Kallmeyer & Osswald (2013: 306).

The lexical elementary tree for the preposition *into* has one lexical leaf node, *into*, as well as a non-terminal leaf node with an NP label, to which the NP *the house* is substituted. The *into* tree is associated with a frame representing the semantics of that preposition, including attributes that are at this point not yet filled with specific values. Due to the substitution of the tree for *the house* at the NP node, the frame associated with that tree provides a value for the ENDPOINT attribute of the prepositional frame, as the values indexed as $\boxed{5}$ and $\boxed{6}$ are unified in the process of substitution. Additionally, the prepositional frame is unified with the event frame for the verb as the PP construction is substituted to the S tree at the PP node, leading to a unification of the values indexed as $\boxed{5}$ and $\boxed{2}$. In the same way, the attribute values indexed as $\boxed{3}$ and $\boxed{1}$ are unified as the tree for *John* is substituted at the NP node of the S tree. Since there are no type conflicts and no other conflicts regarding the values of attributes that need to be unified, the unification of the different trees and frames that are involved in this derivation is successful. The derived frame is reproduced here as Figure 22. In it, the features appearing in the frame contributed by the verb *walked* are unified with features contributed by the arguments during the derivation.

The frame associated with the preposition *into* illustrates some of the properties of the frame semantics assumed by Kallmeyer & Osswald. This frame has two base nodes, labeled as $\boxed{4}$ and $\boxed{5}$, instead of one single root node. The frame also contains a non-functional relation named `part-of`, which specifies that the end point of the path described by this frame is part of the IN-REGION of the entity labeled as $\boxed{w}$. In other words, the path ends inside the entity $\boxed{w}$.

For sentences in which the prepositional phrase is not a complement, but an adjunct, Kallmeyer & Osswald (2013: 306) provide the following analysis. First, since the prepositional phrase is not obligatory, it is not a part of the unanchored tree for the verb. Instead, the unanchored tree for the verb only provides a slot for the syntactic subject. The contribution of PATH information from a prepositional phrase is still possible, but is now implemented via the adjunction of another elementary construction to the VP node. As argued by the authors, adjunct PPs can be iterated and combined, which is possible with TAG's adjunction mechanism: with the binary left-branching structure assumed for the VP by Kallmeyer & Osswald, multiple trees can be adjoined to the VP node to integrate multiple adjuncts into the derived structure. As long as the

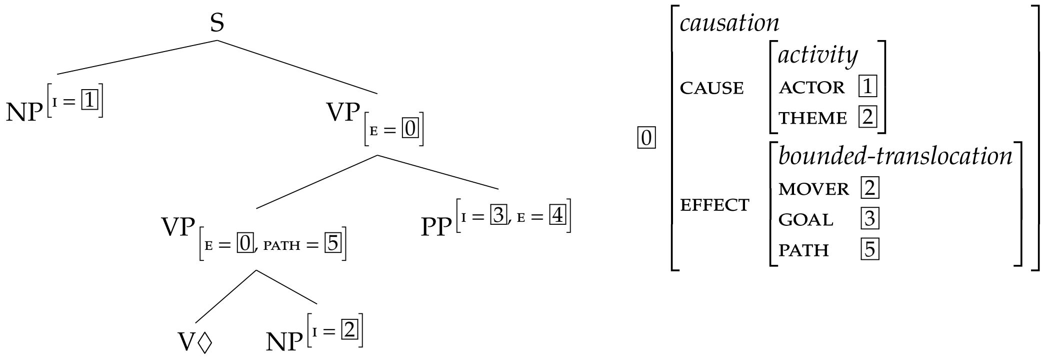


Figure 23: Unanchored tree and semantic frame structure for verbs describing caused motion, as proposed by Kallmeyer & Osswald (2013: 310).

semantic contributions of the different adjuncts can be unified without contradictions, all available attribute values will be integrated into the final frame.

Both types of derivations – with PPs as complements or as adjuncts – are also possible with motion verbs that do not lexically contribute a path. In such cases, the same unanchored tree (with or without a slot for the PP) can be used as for translocation verbs, but no information about the path is contributed by the lexical frame for the verb. Instead, the only available information about the path in the final derived frame is contributed by the PP.

Kallmeyer & Osswald (2013: 310) proceed to illustrate how their grammar handles verbs of transport and caused motion. Their Figure 26 is reproduced here as Figure 23. As mentioned above, these verbs assign the actor and mover roles to different arguments. The authors implement this by assigning different indices to those roles in the unanchored tree that will be used for these verbs. Then, the integration of the different elementary trees during derivation is guided by the interface features and results in an overall frame in which the actor and mover roles are filled by distinct values. The event frame associated with the unanchored tree is of the type *causation*, with an effect subframe that is of the type *bounded-translocation*.

The interface features play a special role in this case because different modifiers contribute meaning to different parts of the event frame. Prepositional arguments and modifiers can contribute information to the *translocation* frame, while other modifiers should integrate into other parts of the event frame. This is solved by associating the PP node's interface feature E with the *translocation* subframe, and the interface feature E of the VP node with the full event frame.

Another challenge is that path modifiers can also attach between the verb and the PP, and modify the embedded event, as in *Mary kicked the ball along the line into the goal*. In the approach chosen by Kallmeyer & Osswald (2013: 310), such modifiers only need access to the path, and not to the event frame above it, since they do not contribute a participant to the motion event. Instead, the authors introduce a new interface feature `path` that is accessible at nodes where path modifiers can adjoin, so that the information contributed by these modifiers can be integrated into the overall frame in this way. A side effect of the introduction of the `path` interface feature is that the directional PP modifier trees now make use of this feature as well instead of including their own (underspecified) *event* frame. As a result, the PPs can modify the path they refer to directly.

In addition to some basic classes that represent intransitive and transitive sentence patterns, Kallmeyer & Osswald (2013: 311) define one class, *DirPrepObj*, that describes trees and frames for directional prepositional objects, which contribute the goal of directed motion events. By not requiring the event associated with the frame of this class to be identified with the frame contributed by the verb by way of an interface feature, the authors avoid having to specify whether the trees in this class are anchored by verbs of caused motion or transport, or by verbs that express a motion of the ACTOR. Instead, export variables are used to determine how the event frame provided by the verb and the frames for the different semantic roles are combined into one final derived frame. This way, the motion expressed by the frame associated with the *DirPrepObj* class can be identified with the top-level event frame for directed motion verbs, or with the embedded frame in the case of caused motion verbs.

### 5.4.3 *Kallmeyer and Osswald II: The English dative alternation*

The second application presented by Kallmeyer & Osswald (2013: 313) is concerned with the English dative alternation, in which two event participants can be expressed in different syntactic positions (see also Chapter 1 of this thesis). Goldberg (1995: 142) understands the double object construction as expressing caused possession, and the prepositional object construction as expressing caused motion; however, as Kallmeyer & Osswald (2013: 315) note, the lexical meaning contributed by the specific verb can override or block these interpretations in individual expressions. Their perspective is supported by Rappaport Hovav & Levin (2008: 138), who view verbs like *give* as having a meaning that expresses caused possession in both constructions, as well as Beavers (2011: 10–11), according to whom four main types of results can be encoded by verbs that license the constructions associated with the dative alternation.

Kallmeyer & Osswald (2013: 318) choose to treat the double object construction as expressing prospective possession, while actual possession can only be contributed by the lexical meaning of a verb that occurs in this construction. The interactions between meaning contributed by the construction and meaning contributed by the verb are then modeled within an LTAG metagrammar.

As in the first application, Kallmeyer & Osswald (2013: 317) first present the lexical frames they intend to implement in their metagrammar. Three example verbs, *send*, *throw* and *give*, are represented with frames (reproduced here as Figure 24 on page 82) that are all of the type *causation*. The verb *throw* is of the type *onset-causation*, since it expresses a *causation* event whose CAUSE subevent only overlaps with its EFFECT subevent at the onset of the latter. They also introduce a new event type *undergoing*, which is responsible for ensuring that the MOVER of a non-active motion event is co-identified with the THEME role of that event. The EFFECT subframe in the frame for *throw* is an example of a frame with a conjunctive type, since it is described as a conjunction of the types *translocation* and *undergoing*. The result is that the constraints for each of these types must be fulfilled.

The three frames differ in their precise event structure and the attributes they provide for the different subevents, due to the difference in the lexical meaning contributed by the different verbs. For instance, only the frame for the verb *throw* encodes a *send*

$$\boxed{0}\begin{bmatrix} \textit{causation} & \\ \textsc{cause} & \begin{bmatrix} \textit{activity} & \\ \textsc{actor} & \top \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{bounded-translocation} \wedge \textit{undergoing} & \\ \textsc{theme} & \boxed{u} \\ \textsc{mover} & \boxed{u} \\ \textsc{goal} & \top \end{bmatrix} \end{bmatrix}$$

*throw*

$$\boxed{0}\begin{bmatrix} \textit{onset-causation} & \\ \textsc{cause} & \begin{bmatrix} \textit{activity} & \\ \textsc{actor} & \top \\ \textsc{theme} & \boxed{u} \\ \textsc{manner} & \textit{throwing} \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{translocation} \wedge \textit{undergoing} & \\ \textsc{theme} & \boxed{u} \\ \textsc{mover} & \boxed{u} \end{bmatrix} \end{bmatrix}$$

*give*

$$\boxed{0}\begin{bmatrix} \textit{causation} & \\ \textsc{cause} & \begin{bmatrix} \textit{activity} & \\ \textsc{actor} & \top \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{change-of-possession} & \\ \textsc{theme} & \boxed{u} \\ \textsc{recipient} & \boxed{v} \\ \textsc{result} & \begin{bmatrix} \textit{possession} & \\ \textsc{possessor} & \boxed{v} \\ \textsc{possessed} & \boxed{u} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Figure 24: Frames for various motion verbs, as proposed by Kallmeyer & Osswald (2013: 317).

MANNER in the CAUSE subevent, while the frames for the verbs *send* and *give* contain less information about the CAUSE subevent and at the same time encode more details for the EFFECT subevent. The EFFECT subevent has *translocation* types for the verbs *send* and *throw*, while it encodes a change of possession instead of a translocation for the verb *give*. The path of the translocation is not specified in either of the two verbs that express a change of location, but will instead be contributed by the prepositional phrase that appears in a construction that is anchored by one of these verbs.

The unanchored trees and semantic frames proposed by Kallmeyer & Osswald (2013: 319) for the constructions associated with the dative alternation are reproduced here as Figure 25 on page 83. They analyze the prepositional object construction like the caused-motion construction in the first application. The double object construction is represented with a similar unanchored elementary tree, which is associated with a *causation* frame with an EFFECT subevent of the type *change-of-possession*, and which provides two NP slots under the VP node into which concrete arguments can be substituted during derivation.

Since the two constructions assign different types to the EFFECT subframe, but the grammar is designed to allow verbs that participate in the alternation to appear in both constructions, the question arises how the type provided by the verb and the type provided by the construction can be unified. According to Kallmeyer & Osswald, the resulting frame is

> [. . . ] a causation with effects along two dimensions: there is a directed motion of the theme and at the same time the theme undergoes a change of possession. (Kallmeyer & Osswald 2013: 318)

In other words, the EFFECT subevent of an anchored tree in this case would belong to both of these types. The authors note that the types do not necessarily exclude each

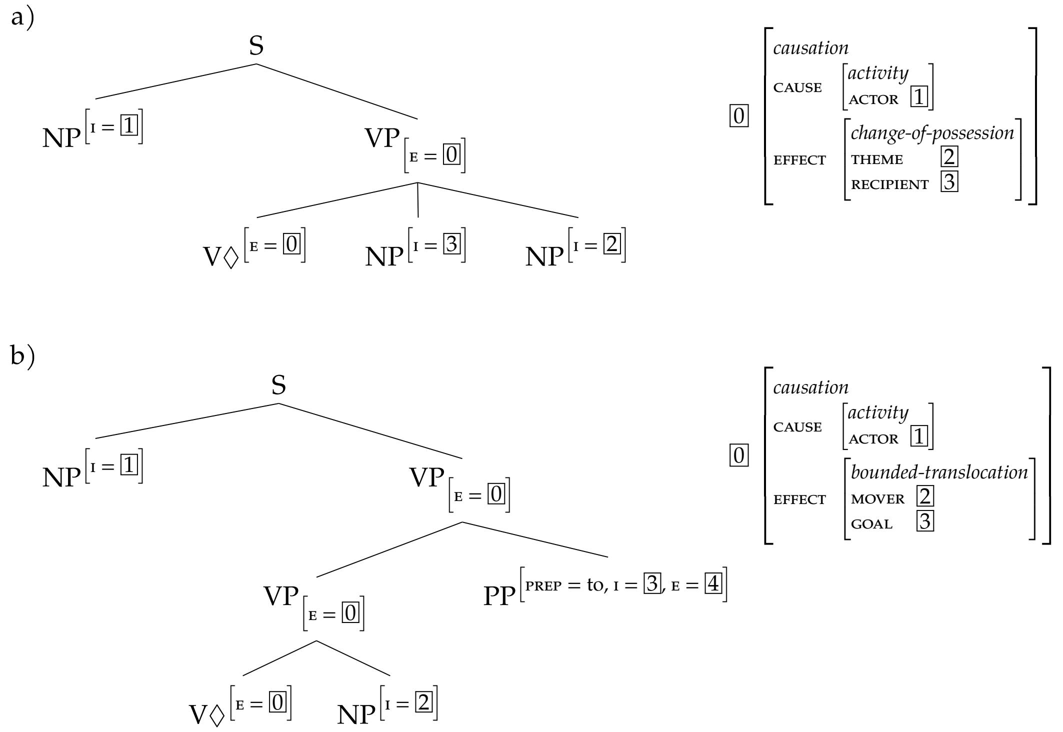


Figure 25: Unanchored trees and frames for the dative alternation: a) the double object construction, b) the prepositional object construction, as proposed by Kallmeyer & Osswald (2013: 319).

other and can be combined with a conjunction. The unification of the features contributed by different constituents follows constraints that are given in the metagrammar and the type hierarchy, to ensure that e.g. the entity associated with the goal of the *bounded-translocation* frame is at the same time the recipient of the *change-of-possession* frame.

Kallmeyer & Osswald implement the metagrammar in a way that ensures optimal factorization, that is, in a way that enables them to

> [...] generalize from the two phenomena [...] and to use the class for directional PP arguments [...] in both the prepositional object construction of the dative alternation and constructions with verbs of directed motion. (Kallmeyer & Osswald 2013: 320)

In this factorization, the metagrammar classes for the indirect object and the prepositional object are defined separately from the classes for the constructions themselves, and export variables are used to ensure appropriate argument linking when the metagrammar is compiled.

The alternation between the two constructions is realized with a separate class in the metagrammars called *DOPOConstr*. This class makes use of the class for transitive verbs defined earlier, as well as one of the two classes *IndirObj* and *DirPrepObj-to*, one of which implements the double object construction, while the other implements the prepositional construction with *to*. The use of these two construction-specific classes is implemented as a disjunction. Upon compilation of the metagrammar, each of these possibilities is generated as a tree class. The frame associated with this metagrammar class is underspecified enough to be able to integrate either of the two constructions and unify the features of the available arguments with those of the frame of the verb that anchors the tree. Information about the *causation* event is unified with information provided by the class for transitive verbs, while information about the embedded *undergoing* event in the main event's EFFECT subframe is unified with information from the class for transitive verbs and from the observed construction. The type of the *undergoing* subevent is combined via conjunction with the type provided by the construction class, either *change-of-possession* or *translocation*.

Kallmeyer & Osswald point out that their model for representing the dative alternation focuses mostly on the selection and combination of types and the interaction between frames provided by the verbs and frames provided by the different constructions in this alternation, and does not take into account other factors that have been discussed as influencing the distribution of the construction, such as "discourse structure effects, heaviness constraints, and the definiteness, pronominality, and animacy of recipient and theme" (Kallmeyer & Osswald 2013: 322). These and other issues are also mentioned in Chapter 1 of this thesis. Kallmeyer & Osswald suggest that a full grammar model would have "an information structure component, ordering constraints which are sensitive to constituent length, and so on" (Kallmeyer & Osswald 2013: 322), and that the use of probabilistic constraints would additionally be required to model the availability of the different constructions in certain environments.

The work by Kallmeyer & Osswald illustrates how well-suited metagrammars for LTAG with semantic descriptions in the form of typed, recursive frames are to model phenomena that involve interactions between lexical meaning and constructional meaning. It also shows how describing the required classes on the metagrammar level allows grammar authors to encode the meaning contributions of single argument realizations and of their combinations in a principled way and without redundancy.

Kallmeyer & Osswald use interface features and export variables to ensure that semantic contributions from the different syntactic arguments in a sentence are incorporated into the semantic frame for a sentence in a way that is appropriate for the observed construction. This is also a useful strategy for modeling the alternations and other constructions under investigation in this thesis part, since the semantic roles for a verb that participates in an alternation can be expressed in different syntactic slots. Linking the syntactic and semantic dimensions to each other is then simply a matter of selecting a specific realization of a tree family for the observed alternation-specific construction.

## 5.5 WHY USE TREE ADJOINING GRAMMAR TO MODEL ALTERNATIONS AND CONSTRUCTIONS?

The extended domain of locality and constraint-based, modular nature of TAG makes it a useful framework for modeling different environments in which verbs can appear. On the syntactic side, this can be used to model diathesis alternations without expressing meaning differences between the alternants. Including a semantic dimension like the `<frame>` dimension introduced by Lichte & Petitjean (2015: 207) makes it possible to also model the semantic properties of different constructions: elementary trees become elementary constructions ("form-function pairs"), and a parser can jointly derive syntactic trees and their corresponding meaning. Lexical frames contributed by the words in an observed input sequence are taken into account as well as constructional meaning, which is made available in the form of partial frame descriptions in the definitions of unanchored elementary construction classes.

Crabbé (2005a) demonstrates how diathesis alternations can be captured on the syntactic level with metagrammar classes. The example he gives is concerned with the realizations of French verbs in their active or passive form. In his words, "[i]n a diathesis alternation the actual form of the verb constrains the way arguments of the predicate are realized in syntax" (Crabbé 2005a: 93). The constraints for the available realizations in this example are implemented as a disjunction of the active case, in which a subject, an active verb form, and an object are referenced with a conjunction operation, and the passive case, in which a subject, a passive verb form, and a by-object are referenced. In addition, a realization with a subject and a passive verb form, but without an object or a by-object, is available. The referenced classes are fragments described elsewhere whose structure can be shared to any class that requires access to them.

For the English verb alternations discussed in this thesis, the form of the verb cannot be used as an indicator to identify which syntactic function each semantic argument takes. However, the general approach of creating classes that group alternating variants by way of disjunction is also applicable for the alternations investigated here. As discussed in Chapter 2 of this thesis, alternations are characterized by the argument structure constructions involved in them, and the meaning of sentences instantiating a construction results from interactions between the meaning of the alternating verb and its arguments and the meaning contributed by the construction. Kallmeyer & Osswald (2013) demonstrate how the constraints guiding these interactions can be spelled out in a metagrammar. Elementary trees that represent the syntactic structure of the alternating constructions are paired with elementary frames, and a derivation for an input sequence consists of a valid tree representing the structure of the whole sequence, paired with a valid frame representing the semantics of the whole sequence.

### 5.5.1 *Inheritance in the metagrammar is not inheritance in the grammar*

Due to the modular, inheritance-based approach of XMG, one may wonder whether the actual grammars yielded by compiling an XMG metagrammar also rely on inheritance. This is not the case. Structure inheritance is at play exclusively at the level of metagrammar class description. The grammatical sequences of the language described by the compiled grammar are not derived via inheritance.

Several arguments have been brought forward against inheritance-based approaches to morphology. This section will focus on two of these arguments. First, Krieger & Nerbonne (1993: 115) point out that what they call "naive" inheritance cannot account for how the iteration of derivational processes can lead to different results, for instance in the case of German *Vor+version* and *Vor+vor+version* (*pre-version* and *pre-pre-version*, respectively).

Second, Krieger & Nerbonne (1993: 115) note that an inheritance-based approach to derivation cannot account for competing analyses of derived words like *undoable* as either *undo-able* or *un-doable*.

These concerns do not apply to the approach proposed in this thesis because the derivation of sequences with an LTAG grammar does not rely on inheritance. Inheritance in the metagrammar is a tool to make it more convenient to describe grammar

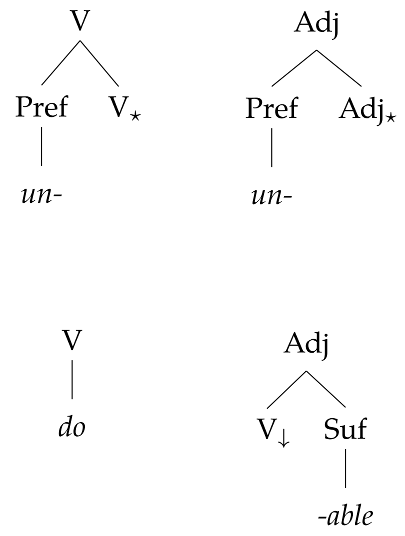


Figure 26: Required trees in a compiled grammar to derive *undoable*.

symbols by expressing generalizations that apply to multiple grammar symbols, without needing to describe the relevant properties again and again for each item independently.

As described in detail by Crabbé (2005a: 84), the main advantage of working with a metagrammar is that it allows us to explicitly encode generalizations in grammatical representations. As he puts it (Crabbé 2005a: 84), "Any substantial modification in a large sized grammar such as the representation of subject-verb agreement requires to modify an important number of independently described units." Employing inheritance in the metagrammar lets us encode such generalizations exactly once and apply them simultaneously to the set of all symbols in the compiled grammar to which they are relevant.

Each symbol of the compiled grammar is an elementary tree of any depth, characterized by its specific structure (nodes in sibling relationships or parent/child relationships), feature structures, node categories, and any other properties distinguishing it from other trees in the grammar. Deriving a tree analysis for a sequence is possible if and only if there is a possible combination of the grammar's trees that accounts for every terminal symbol in the sequence.

In other words, sequences in the grammar are not derived via inheritance from other sequences or words. Instead, combining the grammar's symbols via unification yields one tree for each grammatical sequence – or a set of trees in case the sequence has multiple possible derivations.

Both examples given by Krieger & Nerbonne (1993: 115) can be correctly represented with an XMG approach. Consider the word *undoable*, which can be interpreted either as *un-doable* (something is impossible to do) or *undo-able* (it is possible to reverse something). A minimal LTAG grammar to derive the different senses of this word must contain the trees shown in Figure 26.

All possible combinations of these trees are available as analyses of the sequence *undoable*. The combination of trees does not happen in a particular order. Nor are the sequences *undo* or *doable* in some way "inherited from" the elementary trees anchored

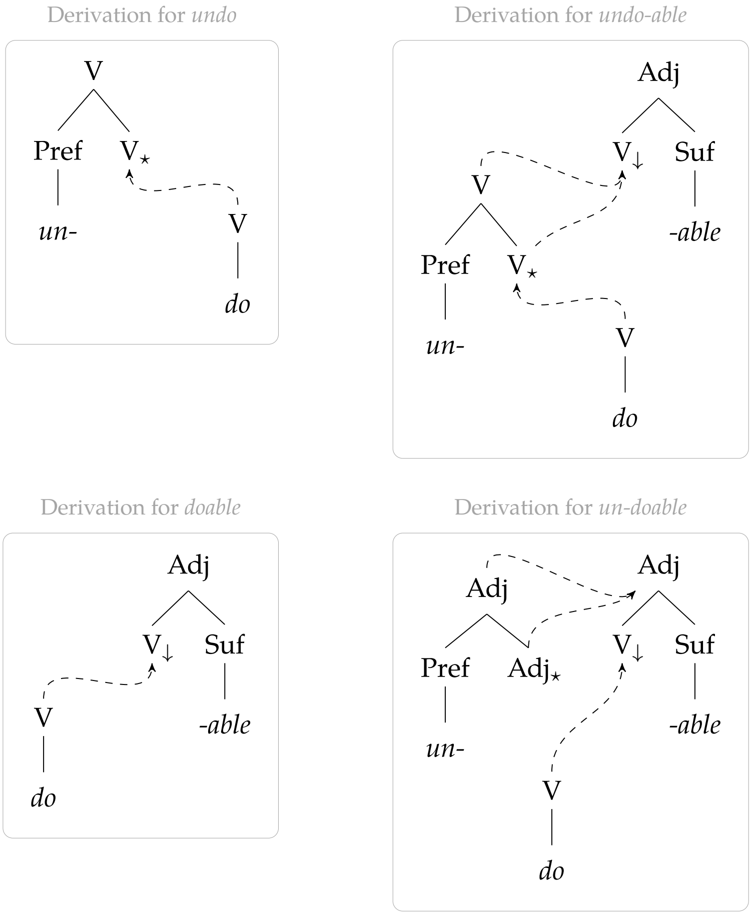


Figure 27: Derivations for *undo*, *undo-able*, *doable*, and *un-doable*

by *un-*, *do*, or *-able*. Within other frameworks, such as HPSG, the prefixation of *un-* may be conceptualized as applying a prefixation rule to a stem and thereby generating the derived word. In this perspective, the stem is the input to the rule and the derived word is the output. Analogous to that process, adjoining the *un-* tree to a tree representing a stem yields the derived tree representing the derived word.

Figure 27 shows the available derivations for *undoable*. Depending on whether the Adj or V variant of the *un-* tree is used, it is possible to either prefix it to a verb, like *do*, or to an adjective, like *doable*. Each variant of the *un-* auxiliary tree contains all information that is necessary to determine whether it can be adjoined to a specific node in another tree. Which one is chosen is not a matter of the order in which different operations take place, but instead of matching the properties of the relevant trees and frames to each other.

The grammars developed in this thesis can be parsed with TuLiPA (Kallmeyer et al. 2008, Arps & Petitjean 2018). When multiple analyses are available for a word, as in the case of *undoable*, the parser presents all competing analyses.

Structure inheritance allows us to factorize the metagrammar to avoid redundancy in the description of the required grammar symbols, or elementary trees. The two variants of the *un-* prefix resemble each other in many respects, and differ in the category of their root and foot nodes. We can define a metagrammar class *AbstractUnPrefix* from which *AdjectivalUnPrefix* and *VerbalUnPrefix* inherit. All properties that *AdjectivalUn-*

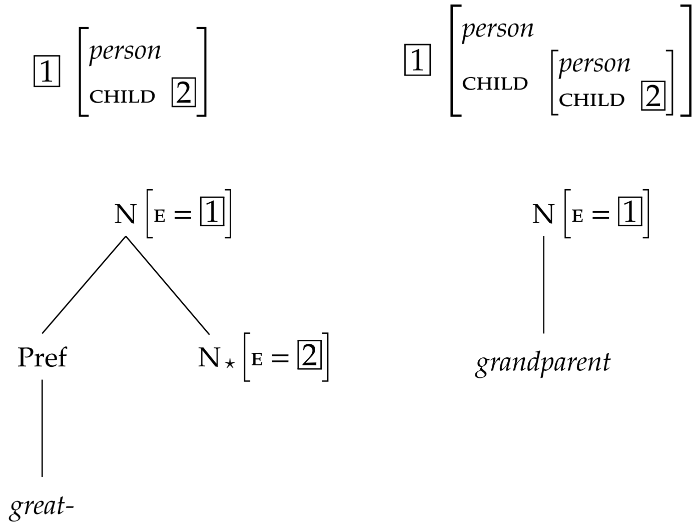


Figure 28: Required trees in a compiled grammar to derive *great-great-grandparent*.

*Prefix* and *VerbalUnPrefix* have in common are only described in the description of the *AbstractUnPrefix* class. All properties in which *AdjectivalUnPrefix* and *VerbalUnPrefix* differ are expressed specifically in the description for each of these two classes.

This example highlights that in the approach taken in this thesis, inheritance matters purely for convenience at the metagrammar level, and the analysis of specific sequences is not based on inheritance, but instead exclusively on the processes of tree substitution and adjunction.

Because the derivation of an analysis for a specific sequence does not rely on inheritance, a single symbol in the compiled grammar can in fact be involved multiple times in one derivation. This allows, for instance, for repeated affixation to account for recursive phenomena. Consider the term *great-great-grandparent*, which is, for the purposes of this argument, analogous to the *Vor+vor+version* example given by Krieger & Nerbonne (1993: 115). This word is not "inherited" from its component parts *great-* and *grandparent*, and the repeated prefixation of *great-* is unproblematic. Figure 28 shows a sketch of a possible TAG grammar fragment for such words, and Figure 29 on page 89 shows the derivation for the term *great-great-grandparent*. The derived tree and frame are shown in Figure 30 on page 89.

As shown in the figures, sequences with multiple instances of a particular elementary tree can be represented without any problems. The expected number of instances of the *great-* prefix appear in the derived tree. The semantic frame of the derived structure results from combining the semantic frames of all elementary trees involved – in this case, two repetitions of the auxiliary tree for *great-* and one instance of the initial tree for *grandparent*. For details on the formal treatment of multiple adjunction in TAG, see Gardent & Narayan (2015). The example shown here follows the standard treatment using so-called dependent derivation (Vijay-Shanker 1987), in which one auxiliary tree adjoins to a node in another auxiliary tree, instead of both adjoining directly to a node in an initial tree. While an argument can be made for using independent derivation here instead, the example illustrates the recursive power of TAG in either analysis.

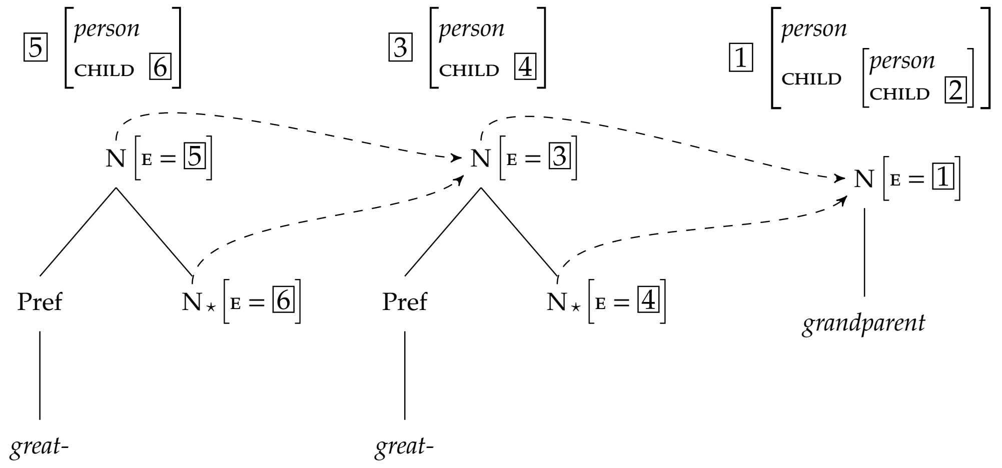

Figure 29: Derivation for *great-great-grandparent*.

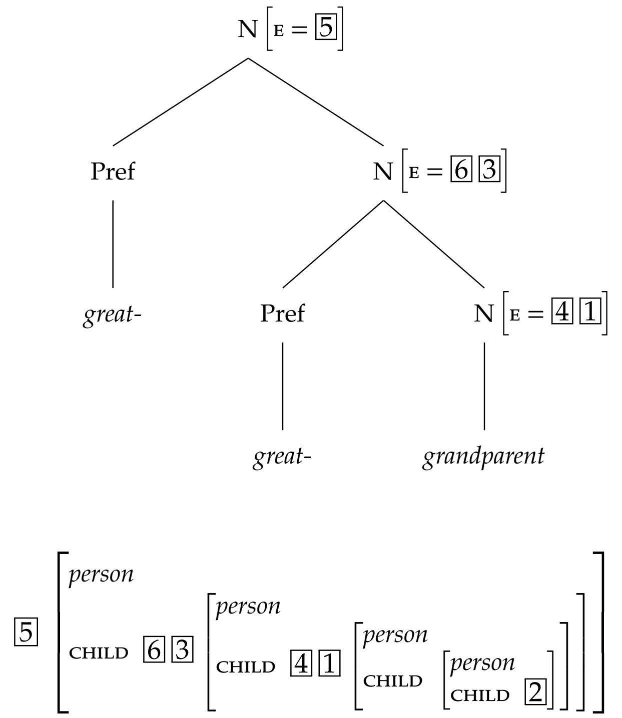

Figure 30: Derived tree and frame for *great-great-grandparent*.

For the sake of illustrating the *great-great-grandparent* grammar fragment, Figure 31 on page 90 shows an example family tree in which Agnes and Albert are the parents of Berta, Berta and Bob are the parents of Clara, and so on down to the youngest family member, Eva. Eva's great-great-grandparents are Agnes and Albert. With respect to the derived tree shown in Figure 30, identifying either Agnes or Albert with $\boxed{5}$ allows us to identify their child, Berta, with $\boxed{6}\ \boxed{3}$. Berta's child, Clara, is identified with $\boxed{4}\ \boxed{1}$. Clara is a grandparent of Eva, so Eva is identified with $\boxed{2}$. Each time a *great-* is affixed to an existing derived form, the adjunction node is extended by the tree contributed by the affix. At the same time, the semantic frame associated with the adjunction node

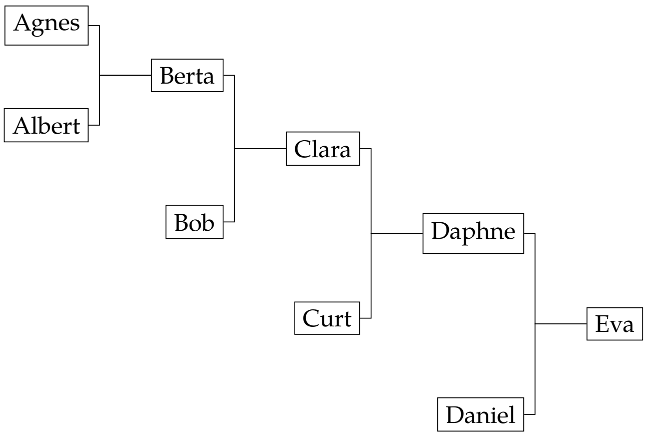


Figure 31: Example family tree.

is embedded into the semantic frame associated with the affix in order to add another ancestor generation.

We can rely on inheritance when designing a metagrammar that generates the grammar fragments for *undoable* or *great-great-grandparent* (or both in one grammar fragment). The relationship between different types of *un-* prefixes has been discussed above. With respect to *great-great-grandparent*, the following inheritance relationships offer themselves in the metagrammar development process.

First, the semantic types involved in the *great-great-grandparent* example can be extended to form an inheritance network. For instance, the frame associated with *grandparent* could be specified to belong to a type *ancestor*, which inherits from a more general type *person*.

Additionally, the descriptions of the trees appearing in this grammar fragment can also be defined based on inheritance between metagrammar classes. For the sake of simplicity, we can restrict ourselves to trees representing morphemes here. A metagrammar class *Morpheme* can provide an underspecified description of structure and semantics of any morpheme that the grammar fragment will cover. A metagrammar class *Affix* inherits from the *Morpheme* class and provides an underspecified description of an auxiliary tree that can extend another tree via adjunction. The *Affix* class does not specify, for instance, at which position the additional nodes will be inserted. Another class *Prefix* does include this information while inheriting all its more general properties from *Affix*. Of course, a parallel class *Suffix* is also conceivable. The possibility of multiple inheritance between metagrammar classes also allows us to add a *Circumfix* class which inherits from both *Prefix* and *Suffix*.

The metagrammar at this point is still missing classes describing the initial trees to which affixes can be added. A possible set of additions could be a *Stem* class inheriting from *Morpheme*, and three distinct classes *VerbalStem*, *AdjectivalStem* and *NominalStem* all inheriting from *Stem*.

Figure 32 on page 91 sketches the metagrammar classes and their inheritance relationships suggested above. Which prefixes and stems in the compiled grammar can combine with each other then depends on the properties of the relevant trees and frames. Are the types compatible? Are the nodes compatible? The rules of unification will prevent derivations for nonsense words like **great-doable* or **un-grandparent*. Implementing the details is outside the scope of this thesis.

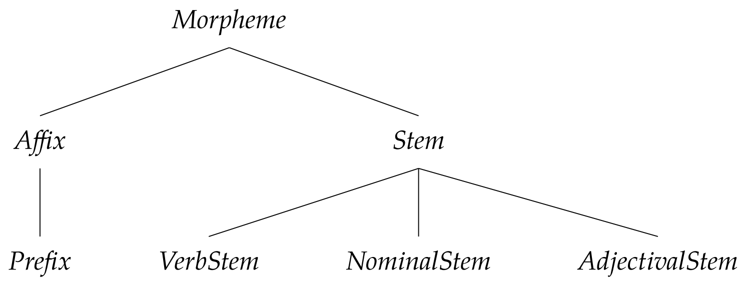


Figure 32: Example metagrammar inheritance relations for the *great-great-grandparent* fragment

Inheritance plays a role exactly as long as we are concerned with designing the metagrammar. Grammar symbols that share a certain set of properties can be described via metagrammar classes inheriting these properties from common superclasses. Compiling the metagrammar yields a set of grammar symbols, each of which is its own independent tree with its own semantic frame. Describing recursive grammar symbols like repeatable affixes (e.g. *great-*) does not require recursion in the metagrammar.

In conclusion, defining metagrammar classes that inherit properties from others does not mean that the symbols of the compiled grammar combine with each other *via inheritance* to derive sequences in the language described by the grammar.

### 5.5.2 *TAG as a framework for modeling natural language*

Whether TAG is a plausible model for natural language is discussed, for instance, by Kroch & Joshi (1985) and Frank (2002). Frank (2002: 35) points out that TAG is a mildly context-sensitive grammar framework, which means that its expressive power is close to the demands of natural language. The generative power of TAG is affected by the use of local constraints on adjunction, such as null adjunction or obligatory adjunction (Kroch & Joshi 1985: 23). Attested phenomena that require mildly context-sensitive grammars to be represented include Dutch verb raising (Joshi 1985), reduplication in Bambara (Culy 1985) and Swiss German cross-serial dependencies (Kroch & Santorini 1991). Linguistic phenomena that can be sufficiently described with context-free approaches can also be modeled in TAG, since every context-free language is also a tree adjoining language (Frank 2002: 32).

TAGs have been applied to a variety of linguistic phenomena, such as raising and equi constructions, passive, *wh-* movement (Kroch & Joshi 1985), extraction and specifically pied-piping in French (Kahane, Candito & de Kercadio 2000), light verb constructions and idioms (Abeillé et al. 1990). Frank (2002: 18) describes how "[TAG's substitution operation] accomplishes effects similar to those of (some of) the generalized transformations from Chomsky (1975) and the Merge operation from Chomsky (1995)". At the same time, applications of substitution and adjunction in a TAG-based theory may not be interleaved with other operations such as transformations; transformational movement cannot be used to create dependencies in TAG that span two elementary trees (Frank 2002: 21).

TAG expresses each syntactic dependency locally within a single elementary tree. As Joshi (2004: 638) puts it, the TAG approach to modeling linguistic structures starts with "complex (more complicated) primitives, which capture directly some crucial linguistic properties and then introduce some general operations for composing these complex structures (primitive or derived)". Joshi (2004: 638) also points out that relying only on a minimum of combining operations has the advantage of remaining language-independent. Frank (2002: 22) takes the operations of substitution and adjunction to be "a universal component of the grammatical architecture" and concludes that "any differences that exist among the grammars of different languages must reside entirely in what elementary trees they take to be well formed".

This approach of starting with complex primitives benefits the representation of phenomena like impersonal passives, as in German *Es wird im Treppenhaus geputzt* (roughly: *There is cleaning being done in the staircase*). Müller (2023: 323) assumes that the impersonal passive comes about in two steps, in which a basic passive form is first derived from an active form, and then the expression of the undergoer in the subject position has to be suppressed to form the impersonal passive. An alternative perspective would be one in which the impersonal passive is regarded as an elementary tree which can be anchored by verbs like *putzen* but does not provide argument slots for the actor and undergoer. In this scenario, it is not necessary to assume a movement of an argument from one slot to another: the elementary trees for active, passive and impersonal passive simply offer exactly the argument slots that are to be filled in each form. There is also no need to assume a derivational relationship between the three variants. Instead, with a metagrammar, the properties that all forms share (for instance, the presence of a verb and a syntactic subject) can be described in more abstract classes, and the properties in which the forms differ (for instance, the content of the syntactic subject) can be described individually in the respective metagrammar classes. Thanks to the inheritance mechanism in XMG, the impersonal passive can inherit a subset of the properties that the basic passive also possesses, instead of necessarily copying *all* properties (which would lead to the issues considered by Müller 2023: 323). For an example of such a description, see Crabbé et al. (2013: 597).

Certain phenomena have been found to be difficult to represent with TAGs in their basic form, such as scrambling. For such cases, multicomponent TAGs have been proposed (Joshi 1987, Weir 1988). They provide the advantage of allowing the addition of sets of trees in a single derivation step, rather than just adding one tree. Kallmeyer (2005) proposes restricted tree-local multicomponent TAGs with shared nodes (RSN-MCTAG) to account for some instances of scrambling in free-word-order languages like German. Balogh (2016) uses tree-local MC-TAG to model verbal fields in Hungarian sentence articulation.

Multiple authors stress the benefits of TAG analyses of certain phenomena with respect to gaining new insights on the linguistic generalizations that are at play (Kroch & Joshi 1985: 6, Joshi 2004: 663). Similarly, one of the goals of the present thesis is to test the presumed linguistic processes and constraints underlying the phenomena under investigation via the implementation of an LTAG metagrammar.

Kroch & Joshi (1985: 6) and Kasper et al. (1995: 93) argue that TAG as a formalism can be regarded to a certain extent as theory-neutral, meaning that the tools it provides can be used to express different linguistic theories and intuitions. The reasoning presented by Kroch & Joshi (1985) refers to the fact that TAG can allow multiple analyses for certain constructions that are incompatible with each other, so that TAG itself cannot be viewed as a theory of universal grammar. They propose multiple analyses, for instance, of the passive, which are roughly analogous to a transformational and lexical analysis in a transformational grammar (Kroch & Joshi 1985: 50). Kasper et al. (1995) propose an algorithm for compiling an HPSG into a TAG. While the two frameworks rely on different notions of localization, it is still possible to translate between them, effectively by performing significant portions of an HPSG derivation at compile-time instead of deriving the structures projected from lexical items at run-time (Kasper et al. 1995: 92-93). Kroch & Joshi (1985: 6) make the general point that "even if the reader rejects particular features of our linguistic analyses he/she should not conclude that the putatively superior analyses are necessarily unstatable within the TAG system". Similarly, Frank (2002: xii) states that "one can pursue TAG syntax using the basic ontological assumptions of any number of frameworks". In other words, it is often possible to achieve the benefits of certain linguistic theories for the representation of certain phenomena within TAG, either in its original form or in a derived form like multicomponent TAG (Joshi 1987), V-TAG (Rambow 2004) or TT-MC-TAG (Lichte 2007).

TAG fragments have been developed for a number of languages, including Arabic (Fraj, Zribi & Ben Ahmed 2008, Ben Khelil et al. 2016), German (Gerdes 2002, Rambow 2004, Kallmeyer & Yoon 2004, Lichte 2007, Kaeshammer & Demberg 2012), English (Kroch & Joshi 1987, Abeillé et al. 1990, Frank 2002, Kaeshammer & Demberg 2012), French (Abeillé 1988, Candito 1996, 1998, 1999, Crabbé 2005b, Bladier et al. 2018), Italian (Candito 1998, 1999), Korean (Han et al. 2000, Kallmeyer & Yoon 2004), Vietnamese (Lê Hồng, Nguyễn & Roussanaly 2008, Le-Hong et al. 2010), Russian (Zinova & Kallmeyer 2012, Zinova 2021), Hungarian (Balogh 2016), Mandarin (Storoshenko 2016), Hindi (Bhatt, Rambow & Xia 2012), and Turkish (Eyigöz 2010).

The properties of TAG discussed above make it a well-suited framework for the modeling of alternation-relevant constructions undertaken in this thesis. While the metagrammars developed here focus on a specific set of constructions and rely on a number of simplifications to avoid distractions from the issues at hand, there are numerous extensions that can be made to achieve a fuller coverage of the English language. As a first step, the metagrammar classes proposed here could be integrated into XTAG (Abeillé et al. 1990), an existing large-coverage TAG metagrammar for English. This integration would require an amount of labor that exceeds what is feasible and appropriate in the context of this thesis, and is therefore left for future work.

### 5.5.3 *TAG as a constructionist grammar framework*

Lichte & Kallmeyer (2017) highlight the benefits of using TAG to model constructions. In particular, they argue that the possibility to jointly describe non-lexicalized syntactic tree fragments and their associated semantics makes TAG a useful framework for capturing insights from Construction Grammar. They point out the parallels between the tree families that can be defined in TAG on the one hand, and the view of language as a "network of constructions" as described by Goldberg (2013) on the other hand.

Like Kallmeyer & Osswald (2013), Lichte & Kallmeyer (2017: 207) pair elementary trees with semantic frames, so that each adjunction or substitution operation can involve a unification not only of the feature structures involved in the trees, but also of the frames associated with them. This is achieved with the use of interface features like `I` and `E` for individuals and events. The type hierarchy is implemented as a set of constraints in the feature logic encoding subtype relationships and type incompatibilities.

Lichte & Kallmeyer relate their view on TAG as a constructionist grammar framework to three ideas by Goldberg (2013). The first, surface-orientation, assumes that there is no "deep" structure, and the focus is on describing and predicting surface structures without structure-destroying operations. According to Lichte & Kallmeyer (2017: 207), this idea is covered by TAG because TAG generates analyses for complex structures directly, that is, not as a sequence of transformation operations, but as a result of combining tree fragments that satisfy the constraints associated with the phenomenon at hand. Additionally, the authors take the view that the process by which TAG models long-distance dependencies – not by movement, but by employing a discontinuous constituent strategy – makes the framework more surface-oriented.

The second idea from Construction Grammar that Lichte & Kallmeyer refer to is the view that constructions exist "at varying levels of complexity and abstraction" (Goldberg 2013: 17). Phrasal constructions and argument constructions are two examples from Goldberg that illustrate the different levels on which these form-function pairs can exist. Lichte & Kallmeyer (2017: 208) distinguish between these terms with regard to whether the constructions are lexicalized or not: "Constructions are pairs of form and meaning that appear not only on the level of words, but also based on bigger syntactic units (phrasal constructions), or unlexicalized abstractions thereof (argument structure constructions)." Lichte & Kallmeyer (2017: 208) point out that the pairing of elementary trees with corresponding semantic frames in TAG "can be seen as constructions on single words or phrasal units". As for unlexicalized argument constructions, which are not anchored by lexical entries, they propose using tree templates, which amount to unanchored elementary trees, and associating these with semantic frames that are then unified with the anchor element's frame during derivation, as shown by Kallmeyer & Osswald (2013: 305). Beyond this, Lichte & Kallmeyer rely on metagrammars to describe the different tree fragments involved in tree templates in a factorized way so that they can be combined as needed, with respect to syntax and semantics. They stress the importance of disjunction as a mechanism to use in tree families to abstract away from specific realizations, which allows one to design tree families to represent argument structure constructions such as the caused-motion construction, as also illustrated by Kallmeyer & Osswald (2013: 274). Crabbé et al. (2013) also point out the benefits of metagrammatical descriptions for the purposes of creating factorized descriptions of (partial) trees, which can be extended with semantic descriptions to yield constructions in the sense of Goldberg.

The third constructionist idea discussed by Lichte & Kallmeyer (2017: 209) is concerned with the ability of constructions to form an inheritance network. In TAG, this inheritance network comes about as a side effect of the factorization that allows tree families to share tree fragments and inherit their properties. The authors note that a metagrammar must be compiled into a grammar before the inheritance network and the tree descriptions contained in it can be used within a tree adjoining grammar. During parsing, the required substitution and adjunction operations are applied to the lexicalized tree templates that represent the set of minimal models described by the metagrammar.

Lichte & Kallmeyer (2017: 210) refer to XMG as an implementation tool that provides the necessary functionalities to implement these ideas. In particular, they draw attention to the availability of different dimensions that can be specified in a metagrammar class in XMG, such as the `<syn>` dimension and the `<frame>` dimension; variables can be shared between dimensions via the interface dimension to implement the desired linking.

For the purposes of this thesis, XMG provides all the functionality that is needed to create a metagrammar describing a set of verb alternations and related constructions that can then be compiled to an LTAG grammar. That grammar can then be used to parse input sequences instantiating the relevant constructions. An initial implementation of a subset of the phenomena discussed in this thesis, presented in Seyffarth (2019b), also shows that XMG and TAG are well-suited to model this type of interaction between constructions and verbs.

Kallmeyer & Osswald (2013) propose metagrammars for directed motion expressions and for the dative alternation in English. In their proposal, elementary trees are paired with partial frames to form elementary constructions, which are then combined via unification as larger structures are formed. This directly parallels the idea expressed by Goldberg (2002: 348) that "an actual expression or *construct* typically involves the combination of at least half a dozen different constructions". According to Goldberg (2002: 348), the following constructions are involved in the sentence *What did Mina buy Mel?*:

a. Ditransitive construction

b. Q-construction

c. Subject-Auxiliary inversion

d. VP construction

e. NP construction

f. Indefinite determiner construction

g. *Mina*, *buy*, *Mel*, *what*, *do* constructions

Each of these constructions can be encoded in the grammar in terms of an elementary tree with an (optional) partial frame description in addition to the relevant syntactic constraints. The constructions listed under g. would be stored as lexical entries, while the other constructions are located in the tree inventory of the metagrammar. Crabbé et al. (2013) organize classes in the metagrammar by placing them on various levels of a class hierarchy, enabling inheritance and the sharing of structural information between constructions. For instance, they would place the constructions listed under d., e. and f. on the hierarchy level reserved for basic tree fragments encoding verb or verb argument realizations. For more on the hierarchy levels used by Crabbé et al., see Section 5.3 of this thesis.

The use of the `<frame>` dimension of XMG to describe elementary constructions has seen several other applications since its introduction by Lichte & Petitjean (2015), for instance in the work by Burkhardt, Lichte & Kallmeyer (2017), who are concerned with modeling depictive secondary predicates in English. These predicates often appear in a sentence-final position and can be ambiguous with respect to the entity or event they target. It is also possible for depictives to be stacked; when multiple depictives occur together, there are certain constraints as to which entity or event can be targeted by each one. Burkhardt, Lichte & Kallmeyer (2017: 21) characterize depictive secondary

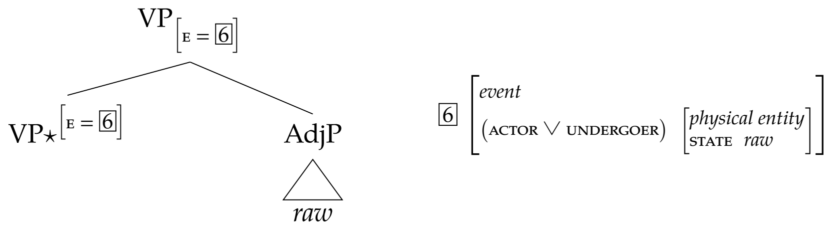


Figure 33: Tree and frame for adverbial constructions that can be anchored by depictive secondary predicates, as proposed by Burkhardt, Lichte & Kallmeyer (2017: 27).

predicates as predicates which "express properties that hold for at least some part of the event time, but do not immediately result from the verb event". Examples for the types of sentences they are concerned with are given in (34).

(34) a. Kim$_i$ ate the steak$_j$ raw$_{*i/j}$.
  b. Kim$_i$ ate the steak$_j$ naked$_{i/*j}$.
  c. Kim$_i$ ate the apple$_j$ unwashed$_{i/j}$.

Burkhardt, Lichte & Kallmeyer present an LTAG metagrammar covering such phenomena that predicts that only certain verbal arguments can be the targets of depictives. The specific verbal arguments that can be modified by depictives are located either at the highest or the lowest level of the ACTOR-UNDERGOER hierarchy proposed by Van Valin (2005: 61). Burkhardt, Lichte & Kallmeyer implement this ambiguity and the relevant constraints for depictives via a disjunction of frame attributes in the adverbial constructions that can be anchored by depictives. Whether the ACTOR or the UNDERGOER is targeted by a depictive in a given input sequence then depends on attribute constraints and other semantic properties of the entity types and the semantics of the specific adverb. For truly ambiguous cases, the parser can derive both analyses. An example for an adverbial construction as proposed by Burkhardt, Lichte & Kallmeyer is reproduced here as Figure 33.

Burkhardt, Lichte & Kallmeyer (2017: 24) argue that the trees anchored by depictives cannot adjoin directly to the elementary tree of the target if it is the subject. In the analysis shown in Figure 33, the tree anchored by the depictive *raw* is an auxiliary tree that can adjoin at the VP node of an elementary tree, for instance one anchored by *ate*. It cannot adjoin to an NP node of a subject or object directly. With this analysis, the semantic frame associated with the depictive tree is unified with the event frame at the adjunction site. The *raw* frame expresses that either the ACTOR or the UNDERGOER of the event will be a *physical entity* that can have a *raw* value in its STATE attribute. For sentence (34a), the subject *Kim* is unlikely to fulfill that constraint, but the object *the steak* will be compatible with the *raw* frame. Analogously, in sentence (34b), the depictive *naked* will express a condition that can hold of the subject *Kim* but not of the object. Finally, the depictive in sentence (34c), *unwashed*, may apply to either the ACTOR or the UNDERGOER of the eating event. This treatment of depictive secondary predicates also allows the target to be unrealized, as in *The game*$_j$ *was played barefoot*$_{i/*j}$.

As Goldberg (2013: 17) points out, morphemes should also be regarded as constructions, because they are "emergent generalizations over existing words in the form of partially filled templates". The work by Zinova (2021) on Russian verbal affixation

$$
e \begin{bmatrix}
\textit{bounded-event} & \\
\textsc{noun-dim} & \boxed{3} \\
\textsc{verb-dim} & \boxed{3} \\
\textsc{m-dim} & \boxed{3} \begin{bmatrix} \textit{scale} & \\ \textsc{min} & \boxed{1} \\ \textsc{threshold} & \boxed{2} \end{bmatrix} \\
\textsc{init} & \begin{bmatrix} \textit{stage} & \\ \textsc{deg} & \boxed{1} \end{bmatrix} \\
\textsc{fin} & \begin{bmatrix} \textit{stage} & \\ \textsc{deg} & \boxed{4} \end{bmatrix} \\
 & \boxed{2} \leqslant \boxed{4}
\end{bmatrix}
$$

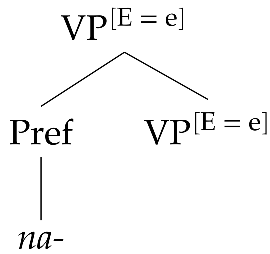


Figure 34: Tree and frame for the Russian prefix *na-*, as proposed by Zinova (2021: 245).

$$
e \begin{bmatrix}
\textit{change-of-state} & \\
\textsc{manner} & [\textit{heat}] \\
\textsc{actor} & \boxed{1} \\
\textsc{theme} & \boxed{2} \\
\textsc{verb-dim} & \boxed{3} \; [\textit{temperature}] \\
\textsc{m-dim} & \boxed{3}
\end{bmatrix}
$$

Figure 35: Frame for the Russian verb *gret′* ‘to heat’, as proposed by Zinova (2021: 246).

shows that morphological processes can be modeled elegantly using LTAG with frame semantics. Her metagrammar captures derivational morphological processes in a way that “allows for a general description of derivational patterns that can be accompanied by a change of the argument structure” (Zinova 2021: 229). The affixes Zinova is concerned with are modeled as tree fragments that can adjoin to a VP node and add affix-specific semantics to the frame that characterizes the VP to which the affix is added. An example metagrammar class for a verbal prefix is reproduced here in Figure 34.

The *na-* prefix in Russian requires a scale that is provided by the verb and is at the same time a parameter of the object. This common scale is identified in the metagrammar class in Figure 34 with the index $\boxed{3}$. The initial stage of the event is indexed with $\boxed{1}$ and corresponds to the minimum point of the measure dimension scale. After the event, the final stage is reached, with the degree of the measure dimension reaching a value indexed with $\boxed{4}$ which is at or above the threshold value of the measure dimension scale. A verb like *gret′* ‘to heat’ can be modified by the *na-* prefix to form *nagret′* ‘to warm up’. The lexical frame Zinova (2021: 246) proposes for *gret′* is reproduced here in Figure 35.

Zinova (2021: 247) proceeds to show how the prefix and the verb combine to form a tree and frame representation for the derived form *nagret′* ‘to warm up’. Her Figure 6.22 is reproduced here as Figure 36 on page 98.

The tree fragments for the verb and the prefix are compatible with each other in terms of their syntactic structure and their semantic description. The prefix is added

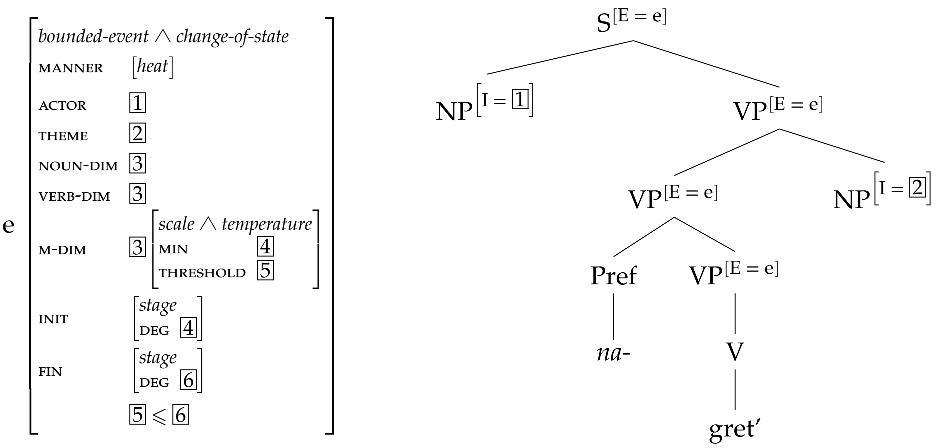


Figure 36: Frame for the Russian verb *nagret'* 'to warm up', as proposed by Zinova (2021: 247).

to the tree anchored by the verb via an adjunction operation at the VP node. Due to the coindexation of the event expressed by the verb and the event the prefix refers to, the respective semantic frames are combined via unification and result in the description shown in Figure 36. The measure dimension is a temperature scale which is provided by the verb. The change of state expressed by the verb is specified to be a bounded event with a final stage at or above some threshold, due to the prefixation with *na-*. The tree includes a slot for a syntactic object, which will be the THEME to which the heating up is applied.

The analysis by Zinova (2021) follows the key ideas from Construction Grammar discussed by Lichte & Kallmeyer (2017): it is surface-oriented, it incorporates constructions at varying levels of complexity and abstraction (since it focuses on modeling verbs and affixes, which can be described as partially-filled words), and it organizes constructions in an inheritance network, using the structure inheritance mechanism provided by XMG.

An XMG implementation of argument linking based on the ACTOR-UNDERGOER hierarchy proposed by Van Valin (2005: 61) is presented by Kallmeyer et al. (2016). Their approach uses the `<iface>` dimension to define language-independent constraints on macroroles which can then be combined with language-specific metagrammar classes that are responsible for the specific syntactic realization of those roles. The work by Kallmeyer et al. constitutes a further factorization over elementary constructions, as macrorole constraints that apply to all constructions in a (language-specific) metagrammar are factored out into separate classes.

### 5.5.4 *Alternative frameworks*

The goal of the implementation presented in this thesis is to explore and predict interactions between verbs, alternations, and related constructions. The metagrammars that will be presented rely on findings reported in linguistic research about the phenomena under discussion. For instance, it is generally agreed upon that semantic requirements are a major factor in determining whether a given verb can or cannot exhibit a specific alternation (see, for instance, the work referenced in Section 6.1.2 of this thesis). Here, participation in an alternation is equivalent to the ability of the verb in question to appear in each of the argument structure constructions characterizing the alternation. Thus, the set of verbs participating in the causative-inchoative alternation corresponds to the set of verbs that can combine with the inchoative construction as well as the causative construction, and the same goes for the other alternations discussed here and the constructions involved in them.

To model the ability of certain verbs to appear in certain argument structure constructions, the requirements identified by linguists must be encoded in the grammar fragment in some way. The semantic type of the event expressed by the verb must be compatible with the type required by the construction (straightforwardly or via coercion). Similarly, the semantic types of the arguments provided by the verb and by the construction must be compatible with each other. The present work puts the focus on encoding these semantic requirements, though of course the syntactic realizations of each argument structure construction under investigation will also be part of the model.

Grammar fragments for similar phenomena as the ones investigated here have been proposed by others in the past. Approaches using LTAG and semantic frames include the analysis of directed motion expressions and the dative alternation in English by Kallmeyer & Osswald (2013), the analysis of the locative alternation in English and Russian by Zinova & Kallmeyer (2012), the analysis of Hungarian verbal fields by Balogh (2016), and the analysis of Russian verbal prefixes by Zinova (2021). In each of these, the grammar fragments benefit from the possibility to connect a syntactic structure with a semantic representation without requiring any structural isomorphism between the two descriptions. Syntactic dependencies are expressed locally within each elementary tree. Some constructions, such as Burkhardt, Lichte & Kallmeyer (2017)'s depictives or Zinova (2021)'s prefixes, are implemented as auxiliary trees that can adjoin to other trees at specific nodes, provided that the frames at that node and at the root of the adjoining tree can be unified. In other cases, argument structure constructions are implemented as initial trees that can be anchored by certain verbs, and whose arguments are added via substitution at the designated nodes. Examples for this design are the analyses for the double object construction and the prepositional object construction proposed by Kallmeyer & Osswald (2013), which are reproduced in Figure 25 on page 83.

There have also been proposals modeling alternations or related constructions in frameworks other than LTAG. Some of those proposals will be summarized in the following. Aranovich & Runner (2001) and Beavers (2005) operate within HPSG, while the proposals presented by Bresnan (1994) and Butt (1993) use LFG. HPSG is an inherently lexicalist framework in which lexical rules are used to realize grammatical function changing operations like passive or dative shift (Aranovich & Runner 2001: 17). LFG assumes a strict separation between the lexicon and syntax (Asudeh, Dalrymple & Toivonen 2008: 69), which is reflected in its basic principle of lexical integrity. Asudeh, Dalrymple & Toivonen (2008, 2013) propose using LFG templates to factor out grammatical information, which makes it possible to access that information either by lexical items or by specific c-structure rules. This way, constructional phenomena can be represented phrasally using c-structure rules that invoke specific templates.

Aranovich & Runner (2001) are concerned with explaining certain differences between dative shift and the spray/load alternation in English, relying on the HPSG framework to model these differences. Their point of departure is the argument by Baker (1997) that dative shift is a transformational rule, and the spray/load alternation is a lexical rule. Aranovich & Runner take the perspective that the qualitative difference between the rules should be analyzed not by assuming a syntax vs. lexicon dichotomy, but instead by recognizing two qualitatively different lexical rule types. They contrast rules that relate lexemes to lexemes (L-to-L rules), which may affect the lexical semantics of a lexical item, with rules that relate words to words (W-to-W rules), which only affect a lexical item's argument structure. In their analysis (Aranovich & Runner 2001: 19), dative shift is covered by a W-to-W rule in which the CONTENT feature of a word is unchanged and only the realization of the arguments is changed. In contrast to this, their analysis for the spray/load alternation makes use of an L-to-L rule which changes the CONTENT attribute's value such that the location argument becomes a patient argument. The reasoning for this distinction is that spray/load alternating verbs have two different lexemes associated with them, while the dative alternation is associated with only one lexeme (Aranovich & Runner 2001: 20).

Beavers (2005) proposes an HPSG analysis of argument alternations that is based on the lexical entailments each verb associates with the participant that alternates between being realized as a direct argument or an oblique. His focus is on locative alternations, but he argues that his approach can be applied to other alternations, based on the generalization that "[d]irect argument variants entail more about the alternating participant than oblique variants" (Beavers 2005: 31). The entailments Beavers focuses on are encoded as constraints on *v-lxm*, involving an attribute ROLES in each verb's CONT value: each NP argument of a verb is assigned some role from the verb's ROLES list. Additional constraints are introduced to determine which obliques can bear which roles.

Bresnan (1994) presents an LFG analysis of locative inversion in English and Chichewa. Locative inversion preposes a locative phrase and postposes the subject NP after the verb, as illustrated in (35) (examples are taken from Bresnan 1994: 75).

(35) a. A lamp was in the corner.
     b. In the corner was a lamp.

For her analysis, Bresnan relies on the lexical mapping theory by Bresnan & Kanerva (1989: 22), which introduces universal constraints on the mapping between semantic roles and syntactic functions. In this mapping theory, argument roles are assumed to be lexically underspecified for the possible surface syntactic functions they can assume (Bresnan 1994: 91). Discourse functions (for instance, presentational focus, which sets a scene and introduces a referent to become the new focus of attention, Bresnan 1994: 90) can trigger alternations of the syntactic functions. In her LFG analysis of locative inversion in English and Chichewa, Bresnan (1994: 92) hypothesizes the same underlying argument structures and the same general principles for mapping a-structure roles into syntactic functions. In both languages, presentational focus allows location to be predicated of a theme argument so that the locative is mapped to the subject and the theme is mapped to an unaccusative object. Additionally, the f-structure analyses for the phenomenon in both languages are virtually identical. At the c-structure level, the languages differ, as locatives are categorized as PPs in English and as NPs in Chichewa.

Butt (1993) presents an LFG analysis of various complex predicates in Urdu. As she points out, the design of LFG allows "a given f-structure to have different c-structure realizations even within the same language" (Butt 1993: 26). This framework, which also allows corresponding expressions in different languages to share their fundamental f-structure while differing in their c-structure, can be used to represent alternations, for instance in languages with relatively free word order. For more on the analysis of argument structure alternations in LFG, see Findlay, Taylor & Kibort (2023).

Beyond various work focusing on representing specific alternations or groups of alternations within LFG or HPSG, as referenced above, there has also been a debate concerning the contrast between lexical and phrasal approaches to argument structure. Some contributions to this debate include Asudeh, Dalrymple & Toivonen (2008, 2013), Müller & Wechsler (2014a), Asudeh & Toivonen (2014), Müller & Wechsler (2014b), Müller (2018). Which of these approaches is chosen to represent argument structure has an impact on how alternations can be analyzed, and whether they are regarded as lexical or phrasal phenomena.

While HPSG, LFG and LTAG offer different tools for modeling linguistic phenomena and come with different assumptions concerning the nature of language itself, it is often possible to translate a grammar fragment from one of these frameworks into another framework. For instance, Kasper et al. (1995) propose an algorithm for compiling a grammar fragment in HPSG into a TAG. Similarly, Findlay (2023a) suggests replacing the context-free grammar backbone of LFG with a tree adjoining grammar in order to achieve "the appropriate level of descriptive freedom to enable LFG to capture substantive idioms" (Findlay 2023a: 202). As Frank (2002: xii) points out, "one can pursue TAG syntax using the basic ontological assumptions of any number of frameworks". Clément & Kinyon (2003) present an approach that allows them to simultaneously generate an LFG and a TAG for French and English from one shared (pre-XMG, Candito 1996 style) metagrammar. Among other things, their metagrammar can generate an LFG without lexical rules, because valency alternations are encoded directly in the metagrammar, and the metagrammar generates all realizations that are valid for each verb.

The present thesis is concerned with modeling alternation behavior as a constructional phenomenon, and aims to develop an analysis that is compatible with the essential ideas from Construction Grammar. Section 5.5.3 has discussed the advantages of using LTAG as a constructionist grammar framework, as argued by Lichte & Kallmeyer (2017). The first basic constructionist idea they list is surface-orientation. Müller (2023: 529) groups LFG, HPSG, TAG and other grammar frameworks together as surface-oriented approaches, because they do not rely on transformations. In this regard, TAG is as suitable for the goals of this thesis as the other frameworks mentioned.

The second key idea from Construction Grammar mentioned by Lichte & Kallmeyer (2017) is the existence of constructions at varying levels of complexity and abstraction. In LTAG, these different constructions can be represented as more or less complex elementary trees associated with (partial) semantic descriptions. LFG also represents form, meaning, and other grammatical information simultaneously, but differs from Construction Grammar in strictly distinguishing between morphology and syntax, based on the principle of Lexical Integrity (Findlay 2023a: 197). However, Findlay (2023a: 202) argues that Construction Grammar's "constructions-all-the-way-down" assumption may be weakened in order to adopt some version of lexical integrity into Construction Grammar, as in frameworks like Sign-based Construction Grammar (Michaelis 2015), which is an HPSG variant (Sag 2010: 486).

The third key idea from Construction Grammar listed by Lichte & Kallmeyer (2017) is the organization of constructions into an inheritance network. In TAG, such an organization is a natural result of the mechanism by which tree families can inherit their structure from various tree fragments. XMG adds another level of possible factorization. Sag (1997: 432) introduces Constructional HPSG, which employs multiple inheritance hierarchies within HPSG to model phenomena such as English relative clause constructions. Syntactic structures in HPSG are typed and organized into an inheritance hierarchy. LFG uses templates, which are "a type of macro which can be used to abbreviate pieces of functional description that are re-used across lexical entries" (Findlay 2023b: 2099). Findlay (2023b) points out that LFG templates are organized in an inclusion hierarchy rather than an inheritance hierarchy, and that templates are a proper part of the grammar while metagrammar classes need to be compiled into a grammar. He also presents two approaches for combining TAG and LFG. He finds that combining TAG and LFG allows one to straightforwardly account for constructional phenomena, which can be represented as either entire (complex) lexical entries or as a part of a lexical entry in the form of a tree template.

The LTAG, LFG and HPSG approaches mentioned above show that there is no one "correct" framework in which to model verb alternations. Each grammar framework provides its own set of mechanisms for predicting alternative surface realizations. All three are surface-oriented. Constructions as form-function pairs can be described in each framework: in LTAG, as elementary trees paired with semantic frames, in LFG, as f-structures that are mapped to c-structures, and in HPSG, as signs that are described in terms of their syntactic properties like argument structure, category and so on as well as their semantic properties like their lexical semantic content. For the purposes of this thesis, the organization of constructions into an inheritance hierarchy matters because various constructions related to verb alternations will be modeled and put into relation with each other. LTAG with frame semantics provides an elegant way to express such a hierarchy, supported by the inheritance mechanism provided by XMG.

The implementations that will be presented in the following chapters will thus most closely resemble the LTAG approach presented by Kallmeyer & Osswald (2013). Analogous to the way Kallmeyer & Osswald formalize insights from linguistic research on the dative alternation in order to model the relevant facets of verbs participating in this alternation and of the argument structure constructions characterizing the alternation, the present work will formalize linguistic insights on the caused-motion construction, the induced action alternation, the causative-inchoative alternation and the instrument subject alternation. A novel aspect of this work in comparison with the LTAG approaches referenced above is the inclusion of ambiguous expressions and the correct treatment of such ambiguous cases based on the semantic type hierarchy and the type requirements of all constructions that could possibly be instantiated by such ambiguous expressions. The coverage of various alternations and additional constructions underlines the benefits of the organization of constructions into an inheritance network, as considerable portions of syntactic and semantic descriptions are shared across constructions. The implementation paves the way for future work on modeling further alternations, although for many alternations discussed by Levin (1993), the semantic requirements on verbs and arguments are not yet sufficiently studied by linguists, and may therefore pose problems in the design of the respective metagrammar classes.

# 6 METAGRAMMAR: THE CAUSED-MOTION CONSTRUCTION AND THE INDUCED ACTION ALTERNATION

This chapter is concerned with developing a metagrammar with frame-semantic meaning representations modeling the behavior of verbs instantiating the caused-motion construction (CMC) and verbs instantiating the constructions involved in the induced action alternation (IAA). This metagrammar will involve a number of constructions relevant to these verbs that share their syntactic form, but are associated with different semantics. The metagrammar needs to allow a parser to recognize not only the syntactic pattern of an input sentence, but also the correct construction in order to output the appropriate semantic frame(s) for that input sentence.

An earlier version of this metagrammar is presented in Seyffarth (2019b). The extended class hierarchy in the metagrammar presented in this thesis, as well as a more elaborate frame type hierarchy and a clearer distinction between semantic contributions from verbs and semantic contributions from constructions, constitute the main improvements of the present work over the results from Seyffarth (2019b).

The implementations described in the following make use of the XMG frame compiler developed by Lichte & Petitjean (2015), which is an application of the formalization proposed by Kallmeyer & Osswald (2013). In the constructions modeled here, no non-functional relations will be used. For the sake of readability and for ease of reference, base labels will be added generously to the substructures in frame descriptions. Frames with multiple base nodes will not appear in the implementation, although this would be possible in the XMG framework.

## 6.1 INTRODUCTION

This section will briefly describe the constructions in focus for the metagrammar to be developed in this chapter. The caused-motion construction and the constructions involved in the induced action alternation share some semantic and syntactic properties, but differ in several ways that need to be modeled by the metagrammar in order to allow the parser to output only valid derivations for sentences instantiating them.

### 6.1.1 *The caused-motion construction*

The English caused-motion construction (CMC, Goldberg 1995: 152, Goldberg & Jackendoff 2004: 540, Kodama 2004, Oyón 2007) is a productive construction denoting a *caused motion*, without requiring instantiating verbs to lexically express a motion-related meaning. The motion sense contributed by the construction is expressed via a directional argument, for instance in the form of a prepositional phrase. The entity that is being caused to move in a certain direction is expressed as the syntactic object, even if the verb otherwise does not typically appear in transitive environments, as shown in (36):

(36) a. Ahab sneezed the map off the table.
b. Queequeg helped Ishmael off the boat.
c. Ahab laughed Starbuck out of his cabin.

The verbs in (36) are not lexically associated with a *caused motion* meaning. For instance, the verb *sneeze* (which belongs to the VerbNet class *hiccup-40.1.1*) has no spatial semantics and does not typically appear with directional prepositional phrases. The caused-motion construction allows a sentence like (36a) to be understood as something like "Ahab's action of sneezing resulted in the map moving off the table." The prepositional phrase *off the table* is the marker of the caused-motion construction. As illustrated by the examples, the construction is available not only for (optionally) transitive verbs, but also for verbs that typically appear in intransitive environments, like *laugh*. The addition of a direct object is valid specifically because of its combination with the directional phrase specifying the direction of the caused motion. Sentence (36c) could be rephrased as something like "Ahab's action of laughing caused Starbuck to leave his cabin".

Goldberg (1995: 152) defines the construction structurally as `[SUBJ [V OBJ OBL]]`, with `V` being a nonstative verb and `OBL` representing a directional phrase, and lists a series of related senses that are all associated with the caused-motion construction. They include "X causes Y to move Z", "X enables Y to move Z", "X prevents Y from moving Comp(Z)", and "X helps Y to move Z". Different verbs appearing in the pattern associated with the construction license different subsets of these available senses. While a lexical meaning component expressing motion is not required, verbs that express (caused) motion can still appear in the construction, like *push* or *shove*, which are associated with the "X causes Y to move Z" sense.

Sentences instantiating this construction express two causally related subevents. The first of these subevents is contributed by the lexical semantics of the verb and denotes the action performed by the ACTOR, which is expressed in the syntactic subject position. For instance, in (36c), the ACTOR *Ahab* is performing the action described by the verb *laugh*. The second subevent is caused by the first one and is underspecified: no verb is present to describe the action performed by the MOVER (the entity in the syntactic object position). Instead, the directional phrase describes the path of a motion that this entity undergoes. In the case of (36c), the MOVER *Starbuck* moves in the direction described by the directional phrase *out of his cabin*. In a semantic frame, the two subevents observed in these types of sentences can be encoded as the CAUSE and EFFECT attributes, respectively, of an event of type *causation* (see Section 4.2 of this thesis and Osswald & Van Valin 2014: 140 for more on such causation frames).

### 6.1.2 *The induced action alternation*

The induced action alternation (IAA, Halliday 1967: 42, Cruse 1972: 521, Levin 1993: 31) is a verb alternation that allows participating verbs to be used in a construction that denotes the primary AGENT (functioning as the ACTOR) inducing, provoking or instigating an action by a secondary AGENT, which is expressed in the syntactic object position. Similar to CMC, the induced action construction also expresses two causally related subevents. In instances of the caused-motion construction, the first subevent is the one described by the verb and executed by the ACTOR, and the second one is a motion event with the MOVER as the main participant. In instances of the induced action construction, the nature of the first, causing subevent is underspecified. The second subevent is described by the verb and denotes the action that is being performed by the entity in the syntactic object position, which is associated with the UNDERGOER role in the CAUSE subevent and with the AGENT role in the EFFECT subevent. While there is a range of possible types of induced actions that can be expressed by sentences instantiating IAA, the examples given by Levin (1993: 31) mainly express *induced motion*, specifically, as illustrated in (37).

(37) a. Sylvia jumped the horse over the fence. (*induced action*)
b. Sylvia jumped the horse. (*induced action, no directional phrase*)
c. The horse jumped over the fence. (*non-induced action*)
d. The scientist ran the rats through the maze. (*induced action*)
e. ? The scientist ran the rats. (*induced action, no directional phrase*)
f. The rats ran through the maze. (*non-induced action*)

Since it is the syntactic object that performs the action described by the verb, the CAUSER can be omitted from sentences like (37a) or (37d) to yield the alternative construction as in (37c) or (37f), respectively. Unlike the caused-motion construction, the induced action construction does not always require the inclusion of a directional phrase, as evidenced by (37b), although this pattern may be less acceptable depending on the verb, as shown in (37e).

Cruse (1972: 521) refers to sentences like (37a), (37b) and (37d) using the term *covert causatives*, and notes that the alternation seems to require a human or human-like causer who imposes their will on an "obedient but independent" agent, who in turn performs the action described by the verb.

Hale & Keyser (1986: 607) refer to verbs that can instantiate the induced action construction as active intransitive verbs of locomotion. They distinguish them from ergatives, which assign accusative case to their syntactic object. Both types of verbs are basically monadic, but active intransitives have an active argument, while this is not the case for ergatives. Hale & Keyser (1986: 609) assume that the transitive variant of such verbs comes about via the addition of an additional external argument. The inherent argument, which features in the intransitive alternant, then cannot assume the subject function and instead assumes the object function.

According to Levin & Rappaport Hovav (1994: 40), the induced action alternation is available specifically to agentive manner of motion verbs. The events expressed by such verbs are considered to not be externally caused. Only the addition of a directional phrase makes it possible to use these verbs transitively. Levin & Rappaport Hovav (1994: 73) view this as a change in the classification of the verb from unergative to unaccusative.

Levin (1993: 31) and Levin & Rappaport Hovav (1994: 40, 70) consider the directional phrase necessary for the transitive use of these verbs, although they do leave open the possibility of the directional phrase being "at the very least understood". Cruse (1972: 521) allows induced action uses of verbs like *gallop* or *march* without any directional phrases, and so do Halliday (1967: 42) and Van Valin (2005: 34).

Although it seems like the variant with a directional phrase is often viewed as the canonical form of the induced action construction, in this thesis, directional phrases

$$\begin{bmatrix} \textit{causation} & \\ \textsc{actor} & \boxed{1}\ \textit{Ahab} \\ \textsc{undergoer} & \boxed{2}\ \textit{the map} \\ \textsc{cause} & \begin{bmatrix} \textit{sneezing} & \\ \textsc{actor} & \boxed{1} \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{translocation} & \\ \textsc{mover} & \boxed{2} \\ \textsc{undergoer} & \boxed{2} \\ \textsc{path} & \textit{off the table} \end{bmatrix} \end{bmatrix} \qquad \begin{bmatrix} \textit{causation} & \\ \textsc{actor} & \boxed{1}\ \textit{Sylvia} \\ \textsc{undergoer} & \boxed{2}\ \textit{the horse} \\ \textsc{cause} & \begin{bmatrix} \textit{activity} & \\ \textsc{actor} & \boxed{1} \\ \textsc{undergoer} & \boxed{2} \end{bmatrix} \\ \textsc{effect} & \begin{bmatrix} \textit{jumping} & \\ \textsc{mover} & \boxed{2} \\ \textsc{actor} & \boxed{2} \\ \textsc{path} & \textit{over the fence} \end{bmatrix} \end{bmatrix}$$

Figure 37: Semantic frames for sentences instantiating the caused-motion construction, like (36a) *Ahab sneezed the map off the table* (left), and the induced action construction, like (37a) *Sylvia jumped the horse over the fence* (right).

will be viewed as optional in this context. This is a simplification. Verbs that participate in the induced action alternation could conceivably be categorized into subsets that impose different requirements regarding the presence of directional phrases. Representing the alternation at this level of detail is beyond the scope of this thesis.

Additionally, while Levin (1993: 31) refers to this alternation as the induced action alternation, all examples from her and from the works cited above involve induced motion, specifically. Because of this, in this thesis, instances of the induced action construction will always be viewed as expressing motion events. This is also motivated by the fact that Levin & Rappaport Hovav (1994: 70) refer to all verbs in this alternation as agentive manner of motion verbs.

### 6.1.3 *Modeling differences between CMC and IAA*

The caused-motion construction and the induced action construction resemble each other in syntactic and semantic terms, but they are not identical. This becomes apparent when both types of sentences are represented in the form of a semantic frame, as illustrated in Figure 37. The frames are simplified to illustrate the essential differences between the constructions; in the actual implementation, the frame representations for such sentences will be more detailed, for instance with respect to the structure of the PATH attribute's value. The frame types for *sneezing* and *jumping* events follow the *v-ing* naming convention introduced on page 46.

As shown in the frames in Figure 37, the subevents involved in the two instantiated constructions are composed in different ways to form different *causation* frames. For the CMC sentence, the event participant expressed in the syntactic subject position, *Ahab*, is the one performing the *sneezing* action. This causes the event participant expressed in the syntactic object position, *the map*, to move in a certain direction. The manner of this movement is not specified in the sentence, but the path is expressed by the directional prepositional phrase *off the table*. On the other hand, in the sentence instantiating the induced action construction, the nature of the causing subevent – an activity performed by the syntactic subject *Sylvia* – is not described further, and the verb instead describes the manner of the motion performed by the syntactic object *the horse* in the effect subframe. This sentence additionally specifies the path of the motion via the directional phrase *over the fence*.

Both example frames contain sets of macroroles for each subevent: the causation, the causing subevent, and the caused subevent. In the caused-motion construction, the causing subevent contains an actor but no undergoer, since the action performed by the actor – in the example sentence, *sneezing* – may not be able to take an undergoer lexically. The caused subevent in this construction contains an undergoer, which is an entity that moves in some direction but does not control the situation. On the *causation* level, the actor is identical to the actor of the cause subframe and the undergoer is identical to the undergoer in the effect subframe.

The induced action construction comes with slightly different macrorole assignments. While the causing subevent is not described directly by the verb, it is assumed to be some action in which an actor impacts the entity that is the actor of the caused subevent (Cruse 1972: 521). This is why the cause subframe in the figure contains both an actor and an undergoer. Following Cruse (1972: 521), the mover in the effect subframe is regarded as an obedient, but independent agent, which also licenses the assignment of the actor macrorole to this entity.

Finally, while the induced action construction is only available to a limited set of verbs in English (according to Levin & Rappaport Hovav 1994: 40, the set of agentive manner of motion verbs), the caused-motion construction is more productive in terms of which verbs it can be instantiated with. Goldberg (1995: 154) points out the prevailing view in the linguistic literature that the caused-motion construction should not be explained by positing additional verb senses – what she calls "rampant lexical polysemy" –, but that the caused-motion meaning instead comes about constructionally. The model developed in the present work involves two mechanisms by which the issue of "rampant lexical polysemy" is avoided. Firstly, the phenomena are implemented with a metagrammar in which recurring structures are only described once and can then be referenced elsewhere multiple times as required (see Chapter 5). Secondly, meanings like motion-causing *laughing* as illustrated in (36c) on page 104 are not regarded as special cases of *laughing*. Instead, the implementation will contain exactly one lexicon entry for *laugh* and a description for the phrasal construction appearing in such sentences. The meaning of (36c) is jointly described by the (single) meaning of *laugh* and the (single) meaning of the construction.

According to Goldberg (1995: 164), while the caused-motion construction is productive and thus does not need to be licensed lexically by the individual verbs that can possibly instantiate it, there are a number of semantic constraints determining whether or not the construction is available in certain environments. For instance, she notes that the causer can be an agent or a natural force, but not an instrument. She goes on to distinguish a number of different scenarios with respect to whether they constitute a valid instance of the caused-motion construction or not. For the purposes of the metagrammar to be developed in this chapter, the focus is on modeling the differences between the general event structure associated with the caused-motion construction and that associated with the induced action alternation. The full range of subtypes of the caused motion construction will therefore not be covered.

Due to the productivity of the caused-motion construction, it is conceivable for verbs that participate in the induced action alternation to also be able to instantiate the caused-motion construction. Unless there are type conflicts that disallow the caused-motion reading for a given sentence whose verb participates in IAA, such sentences are ambiguous: they license both a CMC analysis and an IAA analysis.

Beyond the transitive context with a directional phrase (which will here be viewed as optional), the verbs that can instantiate the constructions under investigation here also license a number of other constructions. The metagrammar to be developed here will cover basic intransitive and transitive uses of these verbs, both with and without directional prepositional phrases. Where more than one construction is available for a sentence involving a specific verb, semantic properties of the verb at hand and of its arguments will be taken into account in order to determine whether the sentence is compatible with each available construction. Basic selectional constraints for the arguments of the verbs under investigation will be implemented as well.

## 6.2 SENTENCE TYPES COVERED BY THIS METAGRAMMAR

The focus of the metagrammar presented in this chapter is on modeling the caused-motion construction and the induced action construction, as well as a number of related constructions. The caused-motion construction and the induced action construction can be instantiated with the same syntactic form, so the metagrammar needs to include the relevant constraints to allow the parser to determine whether a specific derivation for a given sentence is valid or not. Certain sentences can be ambiguous and license competing analyses.

Goldberg (2002: 348) notes that multiple constructions typically interact to form actual expressions, or constructs: for a sentence like *What did Mina buy Mel?*, she lists the following relevant constructions that contribute to the realization of the sentence:

a. Ditransitive construction

b. Q-construction

c. Subject-Auxiliary inversion

d. VP construction

e. NP construction

f. Indefinite determiner construction

g. *Mina*, *buy*, *Mel*, *what*, *do* constructions

The metagrammars developed in this thesis implement this view of surface expressions as the result of an interplay of multiple constructions, each contributing partial syntactic and semantic descriptions that combine to form the overall syntactic and semantic structure for the expression. In the CMC/IAA metagrammar, constructions for different verb uses will be modeled along with constructions covering the realizations of verb arguments. Directional prepositional phrases will be regarded as instantiations of a set of directed motion constructions that contribute different semantic representations to sentences to which they are adjoined.

This metagrammar will cover a range of sentence types that exemplify the differences between the caused-motion construction and the constructions involved in the induced action alternation. Different verb types and arguments will be included in the lexicon in order to make comparisons of minimal pairs possible, and to illustrate how each derivation can be allowed or blocked depending on the verbs and arguments involved. A set of example sentences is given in (38) through (42). Queequeg, Daggoo and Tashtego are human characters. Differences in the semantic interpretation of the different sentences will be discussed in the following.

(38)
a. Queequeg laughed.
b. * Queequeg laughed to Daggoo.
c. * Queequeg laughed Daggoo.
d. Queequeg laughed Daggoo to Tashtego. (*CMC*)

(39)
a. Queequeg danced.
b. Queequeg danced to Daggoo.
c. Queequeg danced Daggoo. (*IAA*)
d. Queequeg danced Daggoo to Tashtego. (*CMC*/*IAA*)

(40)
a. * Queequeg inserted.
b. * Queequeg inserted into the lock.
c. Queequeg inserted the key.
d. Queequeg inserted the key into the lock. (*CMC*)
e. * Queequeg inserted the key from the lock.

(41)
a. ? Queequeg pushed.
b. * Queequeg pushed to Daggoo.
c. Queequeg pushed Daggoo.
d. Queequeg pushed Daggoo to Tashtego. (*CMC*)

(42)
a. Queequeg jumped.
b. Queequeg jumped over the fence.
c. Queequeg jumped the horse. (*IAA*)
d. Queequeg jumped the horse over the fence. (*CMC*/*IAA*)

Sentences like (38a), (39a) and (42a) instantiate a basic intransitive construction, which is only valid for verbs that can appear without a direct object. Thus, sentences like (40a) or (41a) will not be accepted by the parser, since the verbs appearing in them, *insert* and *push*, typically cannot appear in intransitive contexts and will be regarded as obligatorily transitive verbs here.

Sentences like (39b) or (42b) are intransitive and contain a directional prepositional phrase, which instantiates a directed motion construction expressing a motion of the entity in the syntactic subject position and is only valid for verbs whose lexical meaning is compatible with motion. This excludes sentences like (40b) or (41b) because the verbs appearing in them are obligatorily transitive, as well as sentences like (38b) because verbs like *laugh* do not lexically express a motion of the agent along a path.

Transitive sentences like (40c) or (41c) are valid for verbs that can appear with a direct object. Sentence (38c) is not accepted because the verb *laugh* typically only takes one argument.

The verbs featuring in sentences (39c) and (42c) can appear in transitive environments, but only in the induced action sense. In the verb class terminology of Levin (1993), the induced action construction is available to *run* verbs, like *jump*, and also to *waltz* verbs, like *dance*. Both of these classes contain agentive manner of motion verbs. The induced action construction is generally preferred with a directional phrase, but occasionally considered valid without one (Halliday 1967: 42, Cruse 1972: 521, Hale & Keyser 1986: 609–610, Levin 1993: 32, Levin & Rappaport Hovav 1994: 40, Van Valin 2005: 34). A distinction between verbs that require a directional phrase in this context and verbs that do not require one is possible with the mechanisms provided by XMG, but will not be pursued here. The factorized approach and the inheritance hierarchy of semantic types make it easy to add such distinctions in a later version.

Finally, sentences like (38d), (39d), (40d), (41d) or (42d) contain both a direct object and a directional phrase and are thus licensed by the caused-motion construction. In all sentences of this type, the entity expressed in the syntactic object position undergoes some motion that is caused by something that is being done by the entity expressed in the syntactic subject position. Sentences (39d) and (42d) are at the same time instances of the induced action construction, which also involves a direct object and a directional phrase. Based on the metagrammar developed in this chapter, the parser should recognize these cases and derive two competing analyses for each sentence of this type.

Sentence (40e) exemplifies the semantic requirements verbs can impose on their arguments and adjuncts, which also interact with construction availability. In this case, the verb *insert* lexically denotes a caused motion of a theme into some location or position. This is why sentence (40d), whose directional phrase is headed by *into*, is acceptable but sentence (40e) is not. In the metagrammar to be developed in this chapter, the caused-motion construction should be available to verbs like *insert*, and the lexicon entry for *insert* will specify the verb's requirements on the directional phrase.

The sentence types discussed above are categorized into acceptable, unacceptable and questionable sentences, based on the referenced literature on the constructions involved in the two alternations as well as personal intuitions based on experience using the English language. These judgments are not always uncontroversial, for instance in the context of sentences like (39c). A thorough theoretical discussion of these sentence types, taking additional perspectives into account, may yield different judgments. Part IV of this thesis will illustrate that it is not uncommon for verbs to be considered unable to instantiate a particular construction, but nevertheless be observed in corpora in that construction. This thesis part will therefore err on the side of acceptability instead of ruling out less convincing sentence types.

Cases like (41a), which is only indirectly related to the phenomena under discussion in this chapter, will be simplified, in this case by disallowing the intransitive variant of the verb *push*. The following chapter is concerned with verbs that can appear either in intransitive environments or in transitive environments; there, the distinction between intransitive and transitive uses of verbs outside the relevant alternations will be implemented and discussed. Note also that changes in the treatment of such sentences are very easy to implement later if required, due to the modular nature of the metagrammar developed in XMG.

## 6.3 dimensions of the metagrammar

The dimensions discussed in Section 5.3.4 will be relevant for the implementation of the metagrammar. Each dimension will be compiled with a specific XMG compiler. The descriptions for the different dimensions will be arranged in a series of separate files to facilitate processing. The compiled grammar files can then be loaded into the TuLiPA parser to derive trees and semantic frames for the input sentences, as described by Arps & Petitjean (2018). This section gives an overview of the files that will feature in the metagrammar developed here.

Lexical entries for lemmas are stored in a file named `lemma.mg`, whose classes contain descriptions in the `<lemma>` dimension. Inflected forms of lemmas are stored in a file named `morph.mg`, whose classes contain descriptions in the `<morpho>` dimension. The frames representing the meaning associated with each lexicon entry are stored in a file named `frame_dimension.mg`, whose classes contain descriptions in the `<frame>` dimension. The type inventory for the frames in the metagrammar will be described in a separate file named `type_hierarchy.mg`, but will be compiled jointly with the `frame_dimension.mg` file using the `synframe` compiler. Finally, descriptions of all elementary constructions (elementary trees with syntactic and semantic descriptions) and descriptions of elementary trees without semantic descriptions are stored in a file named `syn_dimension.mg`. Each class in this file contains descriptions in a subset of the dimensions `<syn>`, `<frame>` and `<iface>`.

Nodes in the metagrammar classes in the `syn_dimension.mg` file can be marked with certain symbols denoting specific functions such as substitution nodes, null adjunction, or foot nodes in auxiliary constructions. Nodes can also be described further with untyped feature structures defining values for certain properties. Following Kallmeyer & Osswald (2013: 290), nodes representing entities will be linked to (partial) semantic frames via the interface feature `I`, and nodes representing events will be linked to (partial) semantic frames via the interface feature `E`.

The `<frame>` descriptions of classes in the `syn_dimension.mg` file describe the semantics contributed by the construction. The frames may specify frame types or types for values of certain attributes, which effectively constrains which types of events and arguments can instantiate that construction.

The `<iface>` descriptions of classes in the `syn_dimension.mg` file determine how lexical semantic frames will be incorporated into the construction's frame. Interface variables are used to identify lexical frames with `I` or `E` values for individual nodes in the construction's tree. Unanchored classes and classes without their own semantic description do not require an `<iface>` description.

The compiled grammar will be lexicalized, with each class being anchored by a lexical element. In the metagrammar, unanchored "helper classes" exist whose purpose is to enable the sharing of partial structural descriptions between anchored classes. These more abstract classes can add syntactic or semantic descriptions to classes that import them. Variables can be exported in order to be imported elsewhere and thus make a unification of the partial descriptions possible.

Both metagrammars presented in this thesis part will be designed in the modular fashion described above. Each basic vocabulary item is described in terms of its available inflected forms and its lexical frame, and associated with the set of tree families that it can anchor. The files containing this information – `morph.mg`, `lemma.mg` and `frame_dimension.mg` – are mostly independent of the specific phenomena at the syntax-semantics interface. These files will be discussed in Section 6.4. The availability of certain constructions for specific verbs, or of certain roles for specific arguments, is determined by the frame type hierarchy, and additional constraints are contained in the syntactic description of each construction. The files containing these constraints, `type_hierarchy.mg` and `syn_dimension.mg`, will be discussed in Sections 6.5 and 6.6.

## 6.4 LEXICAL DESCRIPTION OF ITEMS IN THE VOCABULARY

The metagrammar focuses on the five verbs that appear in the sentences in (38)–(42) to illustrate the different verb categories with respect to participation in the induced action alternation or compatibility with the caused-motion construction. New verbs can be added to the lexicon to extend the metagrammar. The compatibility of new verbs with each construction must then be described in terms of trees each new verb can anchor and via additional type constraints in the frame type hierarchy.

In the `morph.mg` file, inflected word forms are linked to their lemma forms. An example entry from that file is given in Listing 3. The entry specifies the word form (`“danced”`), its associated lemma (`“dance”`), and the word’s syntactic category (`v`).

```
class MorphDanced
{
  <morpho> {
    morph <- "danced";
    lemma <- "dance";
    cat <- v
  }
}
```

Listing 3: Example verb description in the `morph.mg` file.

The lemmas assigned to the entries in `morph.mg` are described in the `lemma.mg` file. The information contained in this file will allow the parser to anchor the tree families specified for each lemma, as well as look up the lexical frame that is provided for the lemma. An example entry from the `lemma.mg` file is given in Listing 4. This entry specifies the lemma (`“dance”`), the lexical frame it is associated with (`FrameDance`), the syntactic category (`v`), and the name of a tree family that can be anchored by this word (`RegularIntransitiveVerbFamily`).

```
class LemmaDance
{
  <lemma> {
    entry <- "dance";
    sem <- FrameDance;
    cat <- v;
    fam <- RegularIntransitiveVerbFamily
  }
}
```

Listing 4: Example verb description in the `lemma.mg` file.

The entries in the `lemma.mg` file point to frames that represent their semantics, which are in turn described in the `frame_dimension.mg` file. Note that the frames that are stored in this file all represent the lexical meaning of individual lexicon entries – the frames representing constructional meaning are described on the constructional level in the `syn_dimension.mg` file. An example entry from the `frame_dimension.mg` file is given in Listing 5. This entry contains a frame description in which the node variable `?X0` is linked to a frame of the type `dancing`. The `<iface>` dimension specifies that the `?X0` frame is made available in the `E` interface field to any class in the metagrammar that refers to the `FrameDance` class.

```
class FrameDance
declare ?X0
{
  <frame>{
    ?X0[dancing]
  };
  <iface>{
    [e=?X0]
  }
}
```

Listing 5: Example verb description in the `frame_dimension.mg` file.

The `<iface>` dimension is used to store the semantic frame for this lexicon entry in the interface so that it can be made available during parsing.

Only the most essential properties of the lexical frames for verbs are spelled out in the `frame_dimension.mg` file. Other properties, such as requirements on the semantic types of each verb's arguments, will be covered separately in the frame type hierarchy. This is because such selectional requirements can also interact with frames contributed by constructions. In the frame type hierarchy, requirements on frames can be constrained with respect to frames at any stage of the derivation process for input sentences.

In addition to the verbs, the files discussed above also contain entries covering the arguments and prepositions appearing in the example sentences. Example entries for a word that can function as an argument of a verb are given in Listings 6 through 8 below.

```
class MorphDaggoo
{
  <morpho> {
    morph <- "Daggoo";
    lemma <- "daggoo";
    cat   <- n
  }
}
```

Listing 6: Example entity description in the `morph.mg` file.

```
class LemmaDaggoo
{
  <lemma> {
    entry <- "daggoo";
    sem <- FrameDaggoo;
    cat <- n;
    fam <- Propernoun
  }
}
```

Listing 7: Example entity description in the `lemma.mg` file.

```
class FrameDaggoo
declare ?X0
{
  <frame>{
    ?X0[person,
        name: Daggoo]
  };
  <iface>{
    [i=?X0]
  }
}
```

Listing 8: Example entity description in the `frame_dimension.mg` file.

## 6.5 THE LEXICAL FRAME TYPE INHERITANCE HIERARCHY FOR THE CMC/IAA METAGRAMMAR

The `type_hierarchy.mg` file describes a frame type subsumption hierarchy specifying how the frames appearing in the metagrammar are related to each other. The hierarchy is described in terms of constraints on frames, as described in detail by Lichte & Petitjean (2015: 196). These constraints express subsumption relationships, attribute restrictions for specific frames, and incompatibility relationships between specific frames.

Figure 38 on page 115 shows the subtype relationships between the frames that are assigned to the lexicon entries in this metagrammar. More general types are located at the top, more specific types are located at the bottom. The top type $\top$ is the most general type and a supertype of all other types. Types at the leaf level that inherit from the same immediate supertype are mutually incompatible.

In the implementation, incompatibility between types must be explicitly specified. Otherwise, the set of conjunctive types would correspond to the power set of all types in the hierarchy (Lichte & Petitjean 2015: 216). Type incompatibility is crucial to distinguish between verbs or arguments that can or cannot appear in certain constructions. The type requirements imposed by constructions will be added in the following, and an updated type hierarchy that includes the types that are added for this purpose will be presented in Section 6.7.

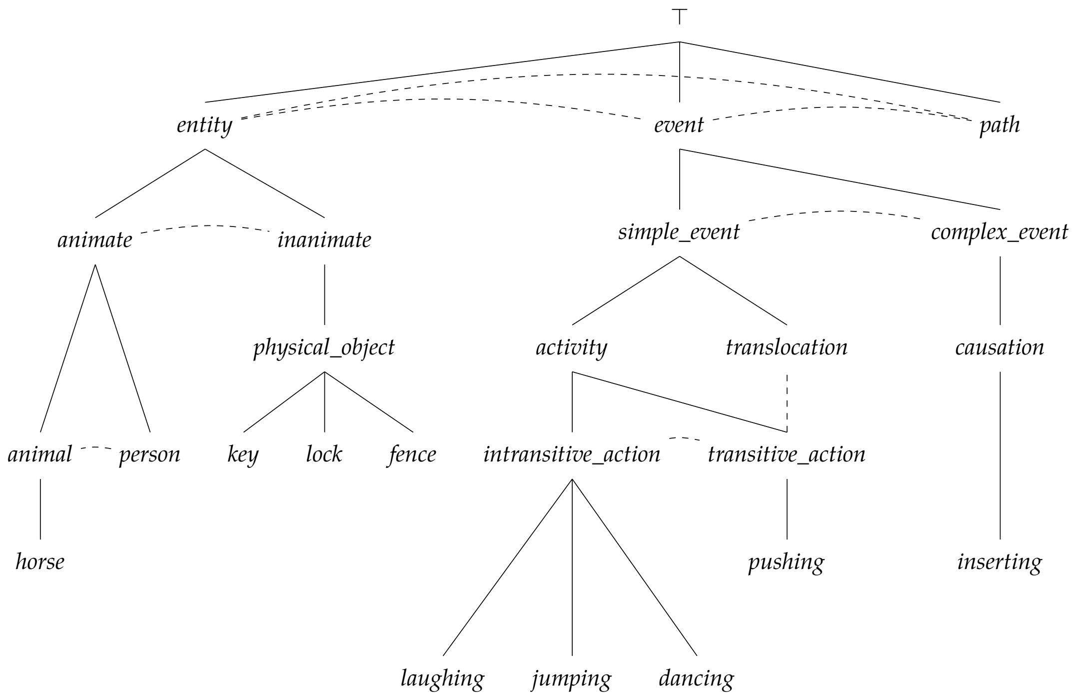


Figure 38: Lexical frame type hierarchy denoting subtype relations and type incompatibilities between the frames assigned to the lexicon entries in the metagrammar for CMC/IAA. Solid lines indicate subtype relations. Dashed lines indicate explicit incompatibility between types. Types at the leaf level that inherit from the same immediate supertype are mutually incompatible.

Beyond the subtype and incompatibility relationships between the semantic types shown in the figure, the type hierarchy also contains constraints that determine the assignment of macroroles in event frames. For instance, if an AGENT is present in a frame, that entity is automatically assigned the ACTOR macrorole. Events of the type *activity* always have an ACTOR. Entities in the AGENT role must be compatible with the type *animate*. With respect to causation events, there are constraints specifying that if the CAUSE subframe contains an ACTOR, that entity is at the same time the ACTOR of the *causation* frame, and if the EFFECT subframe contains an UNDERGOER, that entity is at the same time the UNDERGOER of the *causation* frame (Kallmeyer & Osswald 2013: 279, Kallmeyer et al. 2016: 53).

In the hierarchy shown in Figure 38, the verbs from the example sentences discussed in Section 6.2 fall into three categories: subtypes of *intransitive_action* for verbs like *laugh*, *jump* or *dance*, subtypes of *transitive_action* for transitive verbs like *push*, and subtypes of *causation* for verbs that are lexically causative, like *insert*.

In the type hierarchy, transitivity is one of the facets used to group event verbs into distinct categories. Transitivity is here treated as a semantic notion: a verb is regarded as transitive if the event it describes involves two participants (Hopper & Thompson 1980: 252) – independently of whether these participants are obligatorily expressed at the surface. In fact, the metagrammar fragment presented in Chapter 7 will cover scenarios where a semantically transitive verb, *eat*, can appear in sentences with only one overt participant.

$$\boxed{1}\begin{bmatrix} \textit{inserting} & & \\ \textsc{actor} & \boxed{2} \\ \textsc{undergoer} & \boxed{3} \\ \textsc{cause} & \boxed{4} & \begin{bmatrix} \textit{event} & \\ \textsc{undergoer} & \boxed{3} \end{bmatrix} \\ \textsc{effect} & \boxed{5} & \begin{bmatrix} \textit{translocation} & & \\ \textsc{mover} & \boxed{3} & \\ \textsc{undergoer} & \boxed{3} & \\ \textsc{path} & \boxed{6} & \begin{bmatrix} \textit{path} & \\ \textsc{trajectory} & \textit{into} \\ \textsc{landmark} & \boxed{7} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Figure 39: Lexical frame for the word *insert*

Kallmeyer & Osswald (2013: 298) categorize *push* like *throw*, as a verb of caused motion; in their proposed frame for *throw*, a PATH attribute is contributed by the lexical frame of the verb and is then optionally filled with a specific value during the derivation. However, as Levin (1993: 41–42) notes, *push*/*pull* verbs also participate in the conative alternation, which leaves unspecified whether an attempted action is actually carried out successfully or not. In other words, it is possible to *push* at something without actually causing that thing to move, as illustrated in (43):

(43) Ishmael pushed against the door of Queequeg's room, but it did not move.

Therefore, in the metagrammar to be developed here, only uses of this type of verb which explicitly refer to a PATH with a directional phrase will lead to the inclusion of a PATH attribute in the derived semantic frame. The lexical frame for *push* will not contain a PATH attribute. This makes it possible to add metagrammar classes for the conative alternation later on, if desired.

As discussed on page 110, the verb *insert* lexically specifies a movement of a theme into some location or position. The lexical frame associated with this verb includes a description of that required path, which will effectively block a combination of this verb with a prepositional phrase headed by an incompatible preposition like *over*. Of course, in everyday use, *into* is not the only possible preposition that could appear with *insert*, but this simplification has been made here in order to remain focused on the interactions between verbs, arguments and phrasal constructions. Extensions of the handling of prepositional phrases and of the level of detail dedicated to them are always possible. The lexical frame for *insert* is shown in Figure 39.

Each construction that appears in the metagrammar will be available to a certain set of verbs, possibly under the additional condition that its arguments belong to certain types. These type constraints will be expressed as part of the frame descriptions associated with each construction. The definition of additional event types will become necessary to make the semantic distinctions that are relevant to model the constructions under investigation here. The resulting extensions to the initial type hierarchy shown in Figure 38 on page 115 will be presented in Section 6.7.

The type hierarchy implemented in this metagrammar regards type compatibility as a binary relation that either holds or does not hold between any pair of types. While this has advantages for modeling the constructions, verbs and arguments under investigation here, it does not necessarily fully reflect all facets of type membership as found in the real world. Exceptions, edge cases and subjective interpretations can potentially motivate additions or other changes to the type hierarchy. A large-coverage metagrammar would necessarily also include a large-coverage type hierarchy with more detailed distinctions and additional incompatibility constraints; however, the design of such a large-coverage model is far outside the scope of this thesis. Type coercion is also not implemented in this metagrammar: constructions can conceivably be instantiated with arguments that are lexically specified to be incompatible with a particular semantic role, as long as a human listener is able to reanalyze the argument and assign another type to it that is in fact compatible with the requirements of the construction. This can be the case, for instance, when verbs like *begin* are used with arguments like *book*, which is lexically a physical object or an entity characterized by some informational content, but can be recast into an event type in order to be a valid theme for *begin*. Long, Kallmeyer & Osswald (2022) propose a frame-based analysis for the processes at play in such scenarios. Modeling type coercion or probabilistic type constraints is beyond the scope of this thesis.

## 6.6 construction descriptions

In order to model all example sentences discussed in Section 6.2, a number of constructions must be described in the metagrammar in terms of their syntactic and semantic structure. These descriptions will be located in the `syn_dimension.mg` file. The classes defined in that file are responsible for modeling all initial and auxiliary elementary trees in the grammar, some of them in combination with a semantic description in the `<frame>` dimension. Note that there are cases where sentences have the same syntactic form, but are semantically characterized by different frames; the semantic requirements of the relevant constructions should allow the parser to derive all and exclusively the derivations that are appropriate for each sentence.

As suggested by Crabbé et al. (2013: 612) and discussed in Section 5.3 in the previous chapter, the classes in this metagrammar are structured in an inheritance hierarchy spanning several levels. Classes located at a lower level of the hierarchy can be imported by classes that are on a higher level, which prevents redundancy by making the sharing of structural information possible. The semantic contribution from lexical elements is available to the classes anchored by them, which provides the semantic "building blocks" from which the overall sentence frame can eventually be derived.

Crabbé et al. (2013: 612) implement an example metagrammar across four hierarchical levels: the first one describes tree fragments, the second one describes syntactic functions, the third one describes verbal diathesis alternatives, and the fourth one describes tree families. Their example metagrammar contains purely syntactic classes, i.e., the meaning associated with the resulting structures is not described jointly with the syntactic trees. Both metagrammars developed in this thesis will include an additional level, located between Crabbé et al.'s third and fourth level, which will be responsible for describing constructions. These classes will inherit their syntactic descriptions from classes on lower levels of the hierarchy, and add semantic descriptions. Descriptions in the `<frame>` dimension thus appear exclusively on the construction level. Recall that inheritance relationships among metagrammar classes here are not meant as subconcept relationships, as discussed on page 62. Instead, inheritance links express that all constraints that hold of the inherited class (for instance, *Subject*) also hold of the inheriting class (for instance, *IntransitiveDiathesis*). In other words, *IntransitiveDiathesis* is not a subtype of *Subject*, but it incorporates the constraints provided by *Subject* into its own description and also contains additional constraints that are specific to *IntransitiveDiathesis* itself. All constraints from *IntransitiveDiathesis*, and thus also from *Subject*, are incorporated into the inheriting class *TransitiveDiathesis*, which again adds more constraints specific to itself, and so on.

In the following, the classes located on each level of the hierarchy will be presented and described. An overview of all classes and their inheritance relationships is given in Figure 40 on page 119.

Unanchored classes that contain an `anchor` node will be compiled to lexicalized trees. Unanchored classes without an `anchor` node are for internal use only; they are more abstract, contribute partial descriptions in the `<syn>` and `<frame>` dimensions, and are not evaluated directly but only get imported by other classes. These more abstract classes serve the purpose of factoring out specific aspects of constructions that share parts of their syntactic or semantic structure. For instance, the class *Causation* is unanchored, does not contain any syntactic constraints, and is characterized purely by the *causation* frame that it provides to other classes that inherit from it. The classes *InducedActionConstruction*, *LexicalCausationConstruction* and *General CausedMotion* all inherit from that more abstract class and contain additional descriptions in the syntactic and semantic dimensions. Defining such abstract superclasses thus makes it possible to factor out information that would otherwise have to be repeated in several other classes, and it also makes explicit the syntactic or semantic structures that are shared between classes.

In Figure 40, each unanchored class with an `anchor` node is labeled with the lexical items that can anchor it according to the `lemma.mg` file (for instance, in this grammar fragment, the *Determiner* class can exclusively be anchored by the lexical item *the*). The tree fragments at the first level of the hierarchy will be substituted to certain leaf nodes of other trees during derivation; they are not imported by any other classes. Their semantic contribution is sourced from the lexicon entries for the lemmas that can anchor them. Syntactic functions at the second level of the hierarchy contain information on where arguments are placed in the syntactic structure for a sentence. Tree fragments can be substituted at the substitution nodes in the syntactic function classes to provide information from the lexicon entries of observed arguments. Diathesis alternative classes at the third level of the hierarchy call syntactic function classes to provide a syntactic "backbone" for classes that import them. Constructions at the fourth level of the hierarchy can import other constructions, or directly import diathesis alternative classes or syntactic function classes. The tree families at the last level of the hierarchy exclusively have access to the constructions at the penultimate level.

Here, all constructions that are anchored by prepositions are defined as auxiliary classes that can be adjoined to other classes during derivation. This reflects the fact that prepositional phrases are both logically and syntactically optional for the verbs covered by this metagrammar. They are adjuncts because instead of contributing to the central meaning of the verbs they appear with, they provide additional information (Müller 2023: 32).

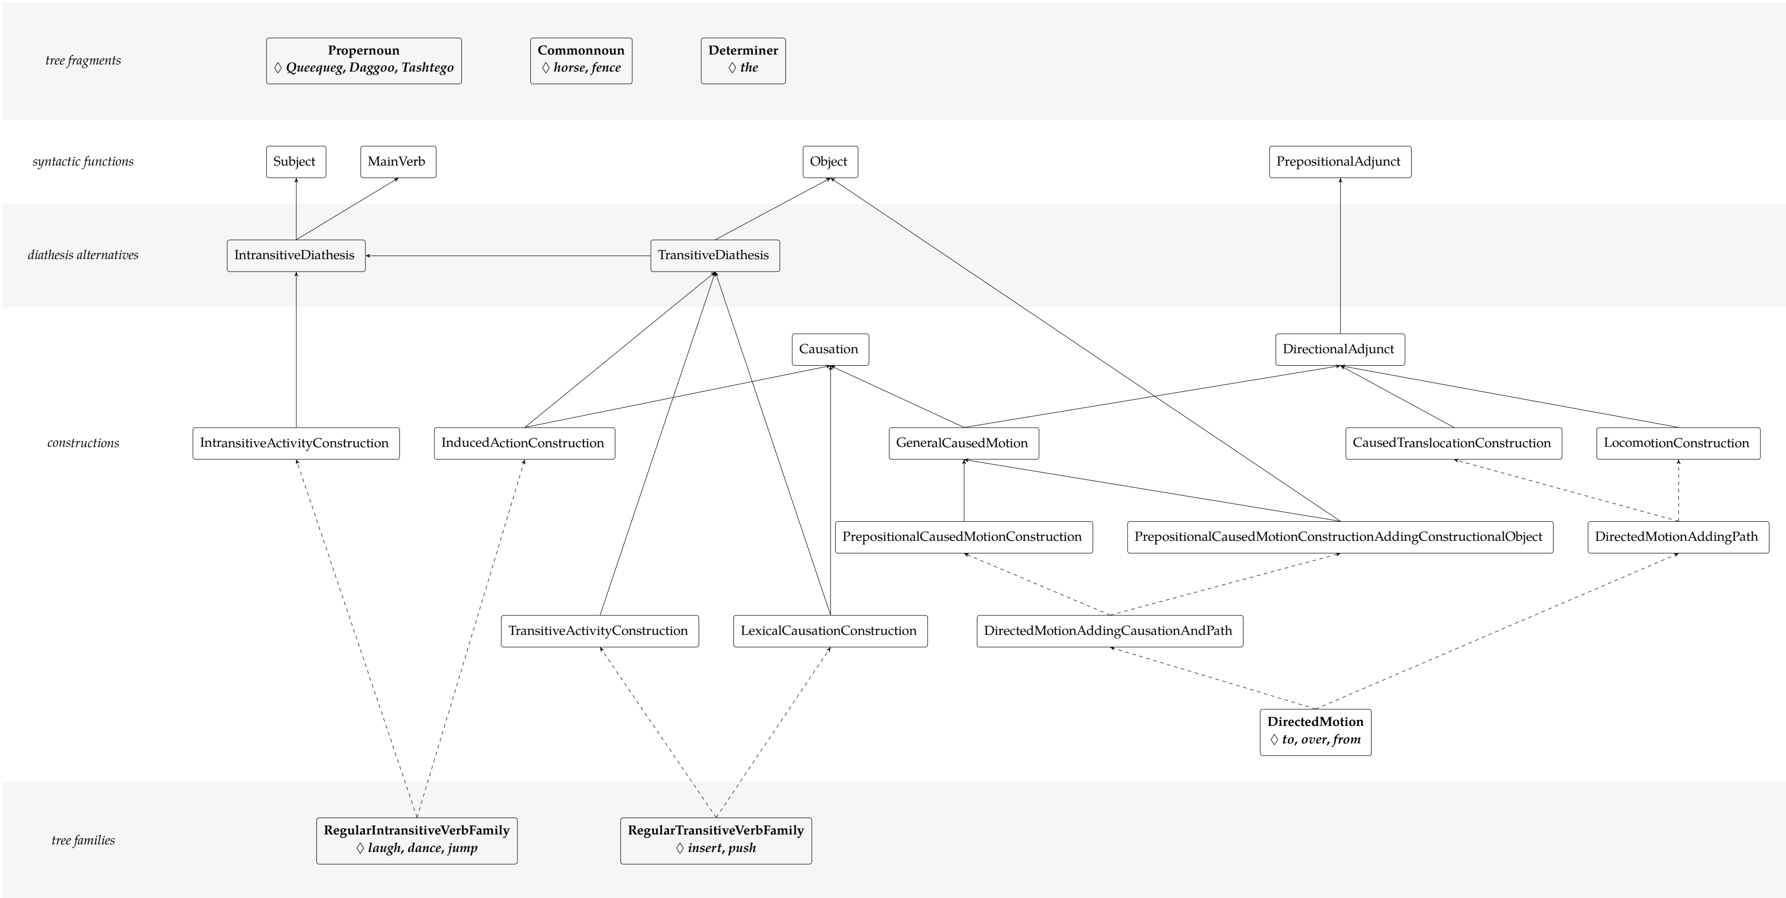


Figure 40: Overview of the class inheritance hierarchy described in file `syn_dimension.mg` of the metagrammar for CMC/IAA. Solid arrows denote structure inheritance. Dashed arrows denote disjunction.

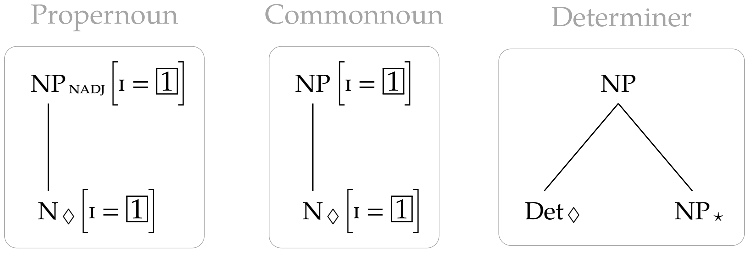


Figure 41: Trees for the *Propernoun* class (left), the *Commonnoun* class (middle) and the *Determiner* class (right) in the CMC/IAA metagrammar.

### 6.6.1 *Tree fragments at the first level of the hierarchy*

At the first level, the metagrammar contains three tree fragments called *Propernoun*, *Commonnoun*, and *Determiner*. Each of them can be anchored by a subset of the lexicon entries found in the `lemma.mg` file; anchor nodes are marked with a $\Diamond$ symbol. The syntactic descriptions contained in these classes are given in Figure 41.

In this metagrammar, determiners do not contribute any additional meaning to trees where they adjoin, so they are not associated with any lexical frame. For each of the other two fragments, the lexical contribution of the element anchoring these tree fragments is associated with the anchor node, for instance, the `N` node in the *Propernoun* fragment, via the `<iface>` dimension. This allows trees which are combined with these fragments via substitution to access these semantic contributions and incorporate them into their larger structures. For the name *Queequeg*, for instance, the interface field `I` will then contain the frame stored in the `frame_dimension.mg` file for the lexicon entry *Queequeg*.

The distinction between common nouns and proper nouns is only sketched here. The *Determiner* class is an auxiliary class that can adjoin to any `NP` node with which it is compatible, which enables the parser to accept sentences like *The horse jumped*. The adjunction of the *Determiner* class to the *Propernoun* class is prevented by an `nadj` mark at the `NP` node of the *Propernoun* class. This prevents any adjunction to that node. This is a simplification that prevents the parser from accepting sentences like **The Queequeg danced*. Further differences between these classes could be added to their descriptions, but this is not the focus of this thesis, so it will not be elaborated on further.

### 6.6.2 *Syntactic functions at the second level of the hierarchy*

At the next level of the hierarchy, the grammar contains four classes named *Subject*, *MainVerb*, *Object* and *PrepositionalAdjunct*. They represent syntactic functions that can appear as constituents in sentences. At this stage, no description is required in the `<frame>` dimension. The descriptions for these classes are given in Figure 42 on page 121.

Substitution nodes are marked with a $\downarrow$ symbol. These are the sites at which the tree fragments discussed previously can be inserted via substitution, given that the unification of the substitution site and the substituting tree's root node is possible without conflicts. This is also why nodes that are not involved in any class-internal structure sharing still receive an index number, such as the $\boxed{1}$ in the *Subject* class in Figure 42. In this case, the function of the tag is to identify the contribution from the tree that will

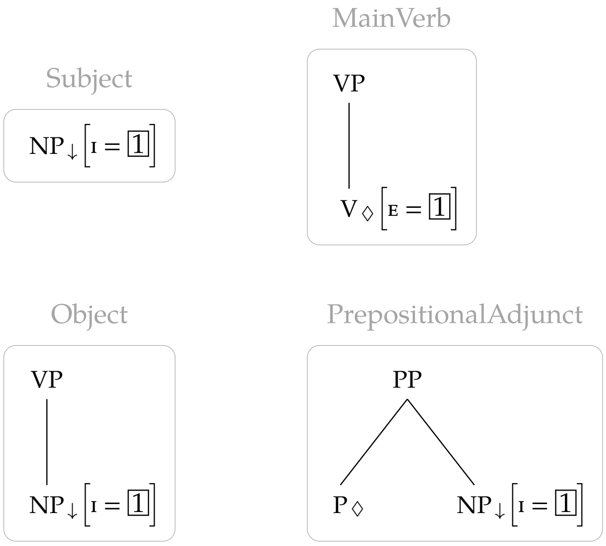


Figure 42: Trees for the *Subject* class (top left), the *MainVerb* class (top right), the *Object* class (bottom left) and the *PrepositionalAdjunct* class (bottom right) in the CMC/IAA metagrammar.

be substituted at the `NP` node with the `NP` node itself, and to enable structure sharing between this `NP` node and whatever node in another class it will itself substitute to.

The `V` node in the *MainVerb* class is marked as an anchor node. When a verb anchors this tree, the semantic contribution from the verb's lexicon entry is unified with the `E` interface field in the class, which makes the verb's lexical frame available to the *MainVerb* class and all classes it is used in.

Note that the described grammatical functions are not linked to particular semantic roles at this point; the classes at this level do not contain any descriptions in the `<frame>` dimension. A higher level of the hierarchy will be responsible for linking the semantics of each argument with the appropriate slot in the overall semantic frame for the input sentence. This separation is necessary because the various constructions covered in the metagrammar can assign different semantic roles to the same syntactic position.

The classes on this level of the metagrammar each export a number of variables that appear in them. This allows other classes to refer to these nodes specifically in order to unify values. For instance, the classes discussed in the next section import the classes on the current level of the hierarchy and add node ordering constraints. Thanks to the export statements, the nodes can be accessed directly from each importing class.

The `VP` node in the *MainVerb* class does not have a value for its `E` feature. This is because some constructions will assign the same frame to the `VP` that is contributed by the anchoring verb's lexicon entry, while others will assign different frames to the `V` and `VP` nodes. These differences will be pointed out again in the section on construction classes.

### 6.6.3 *Diathesis alternatives at the third level of the hierarchy*

The next level of the hierarchy contains the classes *IntransitiveDiathesis* and *TransitiveDiathesis*, which describe the structure of intransitive and transitive basic sentence patterns on a purely syntactic level. The *IntransitiveDiathesis* class imports *Subject* and

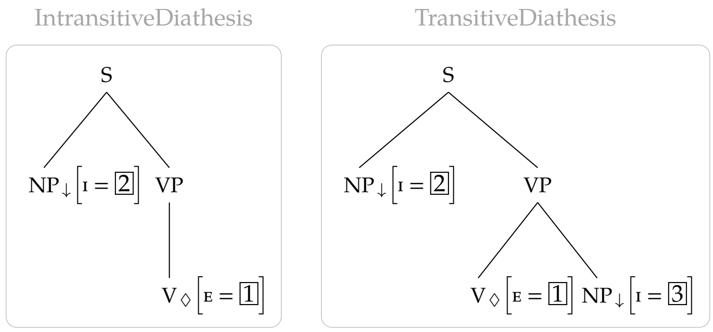


Figure 43: Trees for the *IntransitiveDiathesis* class (left) and the *TransitiveDiathesis* class (right) in the CMC/IAA metagrammar.

*MainVerb* and adds node precedence and dominance constraints. The *TransitiveDiathesis* class imports *IntransitiveDiathesis* and *Object* and also adds node ordering constraints. The descriptions for these classes are given in Figure 43.

These metagrammar classes illustrate how XMG's inheritance mechanism works. While *IntransitiveDiathesis* imports both *Subject* and *MainVerb*, it is not a subconcept of them. Instead, by importing them, it inherits the partial syntactic descriptions provided by those classes. *Subject* provides information on an NP node that can function as a subject. *MainVerb* provides information on a VP node dominating a V node. These pieces of information are now also available to *IntransitiveDiathesis*, which additionally contains information about an S node dominating the imported NP node (from *Subject*) and the imported VP node (from *MainVerb*). *IntransitiveDiathesis* also encodes node precedence between the NP and VP node.

Each of the diathesis classes specifies the order and relative position of the different nodes in the trees they describe. Crabbé (2005a: 93) suggests that this level of the hierarchy can be used to allow different surface realizations, for instance in terms of various permutations of argument order; however, in the context of this thesis, one realization of each relevant construction will be taken into account, so no such argument order permutations are included in the metagrammar. The metagrammar can easily be extended if desired so that it covers more different sentence types.

### 6.6.4 *Construction classes at the fourth level of the hierarchy*

The fourth level of the hierarchy is the level that pairs syntactic forms with different semantic interpretations to describe construction-specific metagrammar classes. Constructions that are anchored by verbs import diathesis alternative classes from the third level of the hierarchy. The metagrammar also contains constructions that are anchored by prepositions, which inherit their syntactic descriptions from the *PrepositionalAdjunct* class on the second level of the hierarchy instead. When a class inherits its syntactic structure from a diathesis alternative class or from the *PrepositionalAdjunct* class, it can optionally add more nodes or constraints to the inherited structure. All constructions that are anchored by prepositions constitute auxiliary classes that can adjoin to certain nodes in other classes during derivation. There are no instances of obligatory adjunction in this metagrammar. All prepositional phrases are treated as adjuncts instead of arguments.

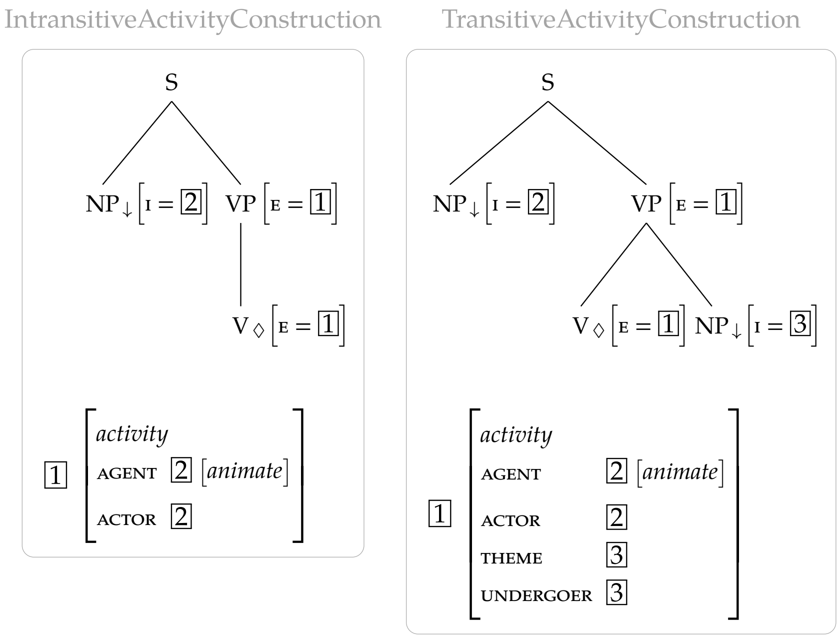


Figure 44: Trees and frames for the *IntransitiveActivityConstruction* class (left) and the *TransitiveActivityConstruction* class (right) in the CMC/IAA metagrammar.

The following sections will present the constructions in three categories. The first category covers constructions reflecting basic activity verb uses. The second category covers causative uses of verbs. The third category covers all constructions anchored by prepositions.

#### 6.6.4.1 *Constructions for regular intransitive and transitive activity verb uses*

The classes describing regular uses of activity verbs that do not instantiate the caused-motion construction or the induced action construction are called *IntransitiveActivityConstruction* and *TransitiveActivityConstruction*. The *IntransitiveActivityConstruction* class covers sentences like (39a) *Queequeg danced*, and the *TransitiveActivityConstruction* class covers sentences like (41c) *Queequeg pushed Daggoo*. The descriptions for these constructions are given in Figure 44.

These classes come with a simple semantic frame that will unify with the frame contributed by the lexicon entry for the verb that anchors the class. In both constructions, the `VP` node and the `V` node share the same value for the interface feature `E`, which links the syntactic structure to the entire semantic frame for the construction.

Both frames include a role slot for an agent, which is identified with the actor macrorole, and the transitive class additionally includes a role slot for a theme, which is identified with the undergoer macrorole. In these classes, the value of the actor attribute is identified with the semantic contribution from the syntactic subject, and the value of the undergoer attribute in the transitive case is identified with the semantic contribution from the syntactic object. A general constraint in the type hierarchy states that the agent role in the frames must be filled by some entity that is compatible with the type *animate*, which includes persons and animals but excludes objects.

The frames associated with these two constructions are of the type *activity*, which is a subtype of *simple_event*. This prevents verbs like *insert*, which is a subtype of *com-*

*plex_event*, from instantiating these constructions, since *simple_event* and *complex_event* are mutually incompatible. The construction that is available for the basic use of verbs like *insert* will be presented in the next subsection.

Activity verbs may appear with additional directional expressions in the form of adjoined prepositional phrases. The constructions that are responsible for directional prepositional phrases will be discussed in Section 6.6.4.3. Sentences like (39b) *Queequeg danced to Daggoo* will thus be covered by a combination of the *IntransitiveActivityConstruction* and a directional phrase construction.

Note that the *TransitiveActivityConstruction* is available exclusively to lexically transitive activity verbs, like *push*. Only these verbs are specified to be able to anchor this class in their lexicon entries. Transitive uses of lexically intransitive verbs like *jump* or *dance* are always causative, and will thus instantiate one of the causative constructions that will be described in the following section.

#### 6.6.4.2 *Constructions for causative verb uses*

Some verbs lexically express causation, which is considered a complex event involving a causing subevent and an effect subevent. This category of verbs is here illustrated with the verb *insert*, which lexically expresses caused motion of a theme into some location or position. The construction that is instantiated by sentences like (40c) *Queequeg inserted the key* is called *LexicalCausationConstruction*.

The construction responsible for modeling the induced action alternation is called *InducedActionConstruction*. This construction is instantiated by sentences like (39d) *Queequeg danced Daggoo to Tashtego* or (42c) *Queequeg jumped the horse*.

The descriptions for these two constructions are given in Figure 45 on page 125.

The multiple inheritance mechanism provided by XMG makes it possible for construction classes to incorporate constraints from separate classes for their syntactic description and their semantic description. Both causative constructions discussed in this section import the *TransitiveDiathesis* class (for its syntactic description) and the *Causation* class (for its semantic description), and add node ordering constraints and additional type constraints for the frame. In particular, the semantic contribution of the anchoring verb's lexicon entry is incorporated into the frame in different ways.

In the *LexicalCausationConstruction* class, the semantic contribution from the anchoring verb's lexicon entry is unified with the entire *causation* frame. In other words, only verbs that are lexically causative have access to this construction. The EFFECT attribute in the frame for the construction has no description of its own because it will be unified with the lexically-specified EFFECT that is associated with the anchoring verb. Similarly, while the ACTOR macrorole in the *causation* frame is identified with the entity in the syntactic subject position, the construction does not assign a *causation*-level UNDERGOER, as the identity of that entity may depend on the specific EFFECT subframe in the lexical entry of the anchoring verb. A syntactic object is required by this construction, and the semantic contribution from that entity is unified with the UNDERGOER role in the CAUSE subframe.

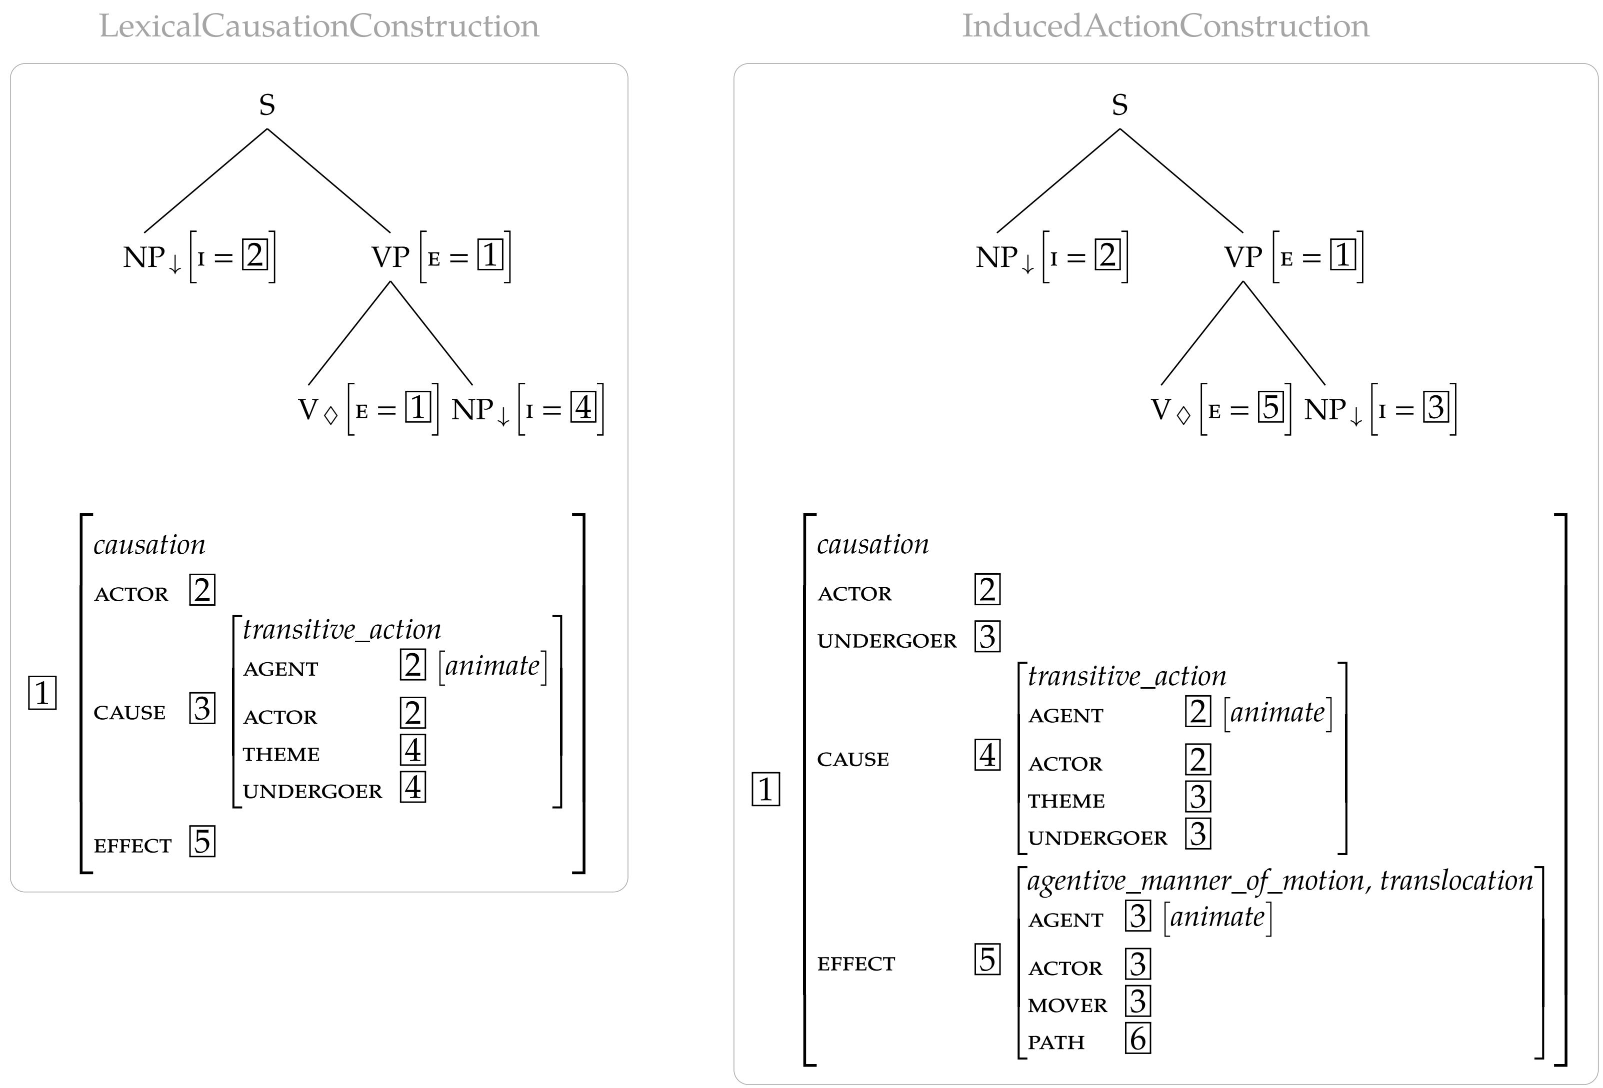


Figure 45: Trees and frames for the *LexicalCausationConstruction* class (left) and *InducedActionConstruction* class (right) in the CMC/IAA metagrammar.

In the *InducedActionConstruction* class, the contribution from the anchoring verb is incorporated into the *causation* frame as the value of its EFFECT attribute: the AGENT in the causing subevent, which is the ACTOR of both the CAUSE subframe and the *causation* frame, is performing some underspecified action that results in the UNDERGOER in the CAUSE subframe performing the action described by the verb (as the AGENT and ACTOR in the EFFECT subframe). Thus, sentence (42c) *Queequeg jumped the horse* is interpreted as expressing that *Queequeg* caused *the horse* to *jump*. The CAUSE subframe is not specified in detail by either the verb or the construction; it merely contains the information that an animate AGENT performs some action impacting an animate UNDERGOER. Since Cruse (1972: 521) calls the UNDERGOER an "obedient, but independent agent", the EFFECT subframe assigns both the MOVER and the AGENT role to this entity, which also licenses the assignment of the ACTOR macrorole to that entity in this subframe. In other words, the movement of the entity expressed in the syntactic object position, while triggered by something that is done by the entity in the syntactic subject position, is still performed as an independent action by the entity in the syntactic object.

The difference between the constructions is apparent when comparing the values of the interface feature `E` for the `VP` and `V` nodes. In the *LexicalCausationConstruction* class, both the `VP` and the `V` are linked to the entire *causation* frame. In the *InducedActionConstruction* class, the `VP` is linked to the entire *causation* frame, but the `V` is linked to the EFFECT subframe. The syntactic descriptions of the two constructions are otherwise identical.

Because the induced action construction is here considered to always involve motion, a PATH attribute is included in the construction's EFFECT subframe. Its value will either remain unspecified, or be unified with a *path* contributed by a preposition-anchored adjoining construction.

This is the only class in this metagrammar where the verb specifies an action performed by the entity expressed in the syntactic object position. Only verbs that participate in the induced action alternation will be able to anchor this construction. The semantic requirement for the verb to be compatible with the *agentive_manner_of_motion* type, which is motivated by the analysis by Levin & Rappaport Hovav (1994) that has been discussed in Section 6.1.2, prevents all verbs outside this category from instantiating this construction. This construction therefore requires the addition of constraints specifying the (in-)compatibility of the existing types in the frame type hierarchy with the *agentive_manner_of_motion* type.

The abstract class *Causation* is imported by both causative constructions discussed in this section. In the imported class, the type of the EFFECT subframe is not specified. The *InducedActionConstruction* class imposes its own type requirements here. In the *LexicalCausationConstruction* class, the type of the EFFECT subframe is not specified by the construction, but will instead be sourced from the lexicon entry of instantiating verbs.

Concerning sentences like (42d) *Queequeg jumped the horse over the fence*, they involve a realization of the *DirectedMotion* class adjoining to the *InducedActionConstruction* at the `VP` node to add a prepositional phrase subtree and its corresponding *path* frame. The different realizations of the *DirectedMotion* class will be presented in the following subsection.

*to*

$$\boxed{1}\begin{bmatrix} \textit{translocation} & \\ \text{PATH} & \boxed{2}\begin{bmatrix}\textit{path} & \\ \text{TRAJECTORY} & \textit{to}\\ \text{LANDMARK} & \boxed{3}\end{bmatrix}\end{bmatrix}$$

*over*

$$\boxed{1}\begin{bmatrix} \textit{translocation} & \\ \text{PATH} & \boxed{2}\begin{bmatrix}\textit{path} & \\ \text{TRAJECTORY} & \textit{over}\\ \text{LANDMARK} & \boxed{3}\end{bmatrix}\end{bmatrix}$$

*from*

$$\boxed{1}\begin{bmatrix} \textit{translocation} & \\ \text{PATH} & \boxed{2}\begin{bmatrix}\textit{path} & \\ \text{TRAJECTORY} & \textit{from}\\ \text{LANDMARK} & \boxed{3}\end{bmatrix}\end{bmatrix}$$

*into*

$$\boxed{1}\begin{bmatrix} \textit{translocation} & \\ \text{PATH} & \boxed{2}\begin{bmatrix}\textit{path} & \\ \text{TRAJECTORY} & \textit{into}\\ \text{LANDMARK} & \boxed{3}\end{bmatrix}\end{bmatrix}$$

Figure 46: Lexical frames for the prepositions *to*, *over*, *from* and *into* in the CMC/IAA metagrammar.

#### 6.6.4.3 *Constructions for prepositional phrases*

Prepositions are specified in the lexicon of this metagrammar to be able to anchor an auxiliary class called *DirectedMotion*. This class is defined as a disjunction of two classes of constructions. The first variant, *DirectedMotionAddingCausationAndPath*, covers instances of the caused-motion construction, which provide a *causation* frame into which the semantic contribution from the verb is embedded. The second variant, *DirectedMotionAddingPath*, covers constructions whose entire semantic frames unify with the frame of the VP node to which they adjoin.

Each preposition can anchor each of these different constructions. The observed syntactic context and the types associated with the observed arguments will determine which of the alternatives are valid for a given input sequence. In all cases, a *path* will be provided to the construction to which the directed motion construction adjoins.

Kallmeyer & Osswald (2013: 302) propose lexical frames for prepositions with the type *event*. The specific directional meaning of the preposition is provided in the value of the PATH attribute of this *event* frame. The lexicon entries for prepositions in the metagrammar developed in this chapter are designed in a similar fashion. In the context of the sentence types that will be parsed with this metagrammar, all frames containing a PATH attribute have the type *translocation*. The lexical frames for the prepositions *to*, *over*, *from* and *into* are shown in Figure 46.

Modeling the exact nature of different types of paths that can be specified by prepositional phrases is outside the scope of the metagrammar to be developed here. Kallmeyer & Osswald (2013: 302) model paths in more detail and are concerned with, among other things, the interaction between the motion-related meaning contributed by the prepositional phrase and the potentially motion-related meaning contributed by the verb. The main focus of this chapter is on differences in the behavior of the verbs under investigation, while the motion-related meaning is almost exclusively contributed by the prepositional phrases.

For the sake of simplicity, paths are not represented at a high level of detail here. Instead, paths are specified as instances of the type *path*, with a LANDMARK attribute specifying the entity in relation to which the path is characterized, and a TRAJECTORY attribute defining the direction of the moving entity in relation to the LANDMARK. The value of the TRAJECTORY attribute is a constant that is contributed by the lexicon entry of prepositions like *to*, *over*, *into* or *from*. This simplified representation of prepositional meaning circumvents the necessity of dealing with the high degree of polysemy of prepositions (see e.g. Tyler & Evans 2001), but does not detract from the validity of the metagrammar with respect to the verbs and their behavior.

The metagrammar classes covering prepositional phrases make use of the class inheritance and factorization mechanisms provided by XMG. Several unanchored helper classes are defined in the metagrammar from which multiple other classes inherit certain partial syntactic or semantic descriptions. The different realizations of the *DirectedMotion* class will be discussed separately in the following.

All realizations of the *DirectedMotion* class inherit from a metagrammar class called *DirectionalAdjunct*, which is responsible for encoding the semantic contribution of directional prepositional phrases in the form of a frame of the type *translocation* whose value for its PATH attribute is determined by the anchoring preposition. Its value for the LANDMARK attribute comes from the prepositional object that substitutes into the construction.

The two realizations of the *DirectedMotionAddingPath* class differ in their semantic descriptions. The *LocomotionConstruction* class provides a *translocation* frame that unifies with the frame at the VP node to which this auxiliary construction adjoins. The *CausedTranslocationConstruction* class, on the other hand, provides a *causation* frame, which means that it can only adjoin to a VP node whose frame is compatible with the *causation* type. The two classes are shown in Figure 47 on page 129.

The *LocomotionConstruction* class can adjoin to any VP node whose semantic frame is compatible with *translocation*. According to the type hierarchy discussed in Section 6.5, *transitive_action* and its subtypes are incompatible with *translocation*. Individual verbs of the *intransitive_action* type, such as *laugh*, may also be incompatible with *translocation*. This is how the parser is prevented from deriving an analysis for sentences like (38b) **Queequeg laughed to Daggoo*. If a verb is compatible with *translocation*, the attributes and attribute values contributed by the *LocomotionConstruction* class are added to or unified with the existing attributes and attribute values in the frame of the construction anchored by the verb.

The *LocomotionConstruction* class is used for sentences like (39b) *Queequeg danced to Daggoo*. In such cases, the ACTOR is the AGENT, the entity expressed in the subject position, and the same entity also fills the MOVER role which is contributed by the *LocomotionConstruction* class. The constraint that the AGENT must be *animate* prevents the parser from accepting sentences like **The door jumped over the fence*.

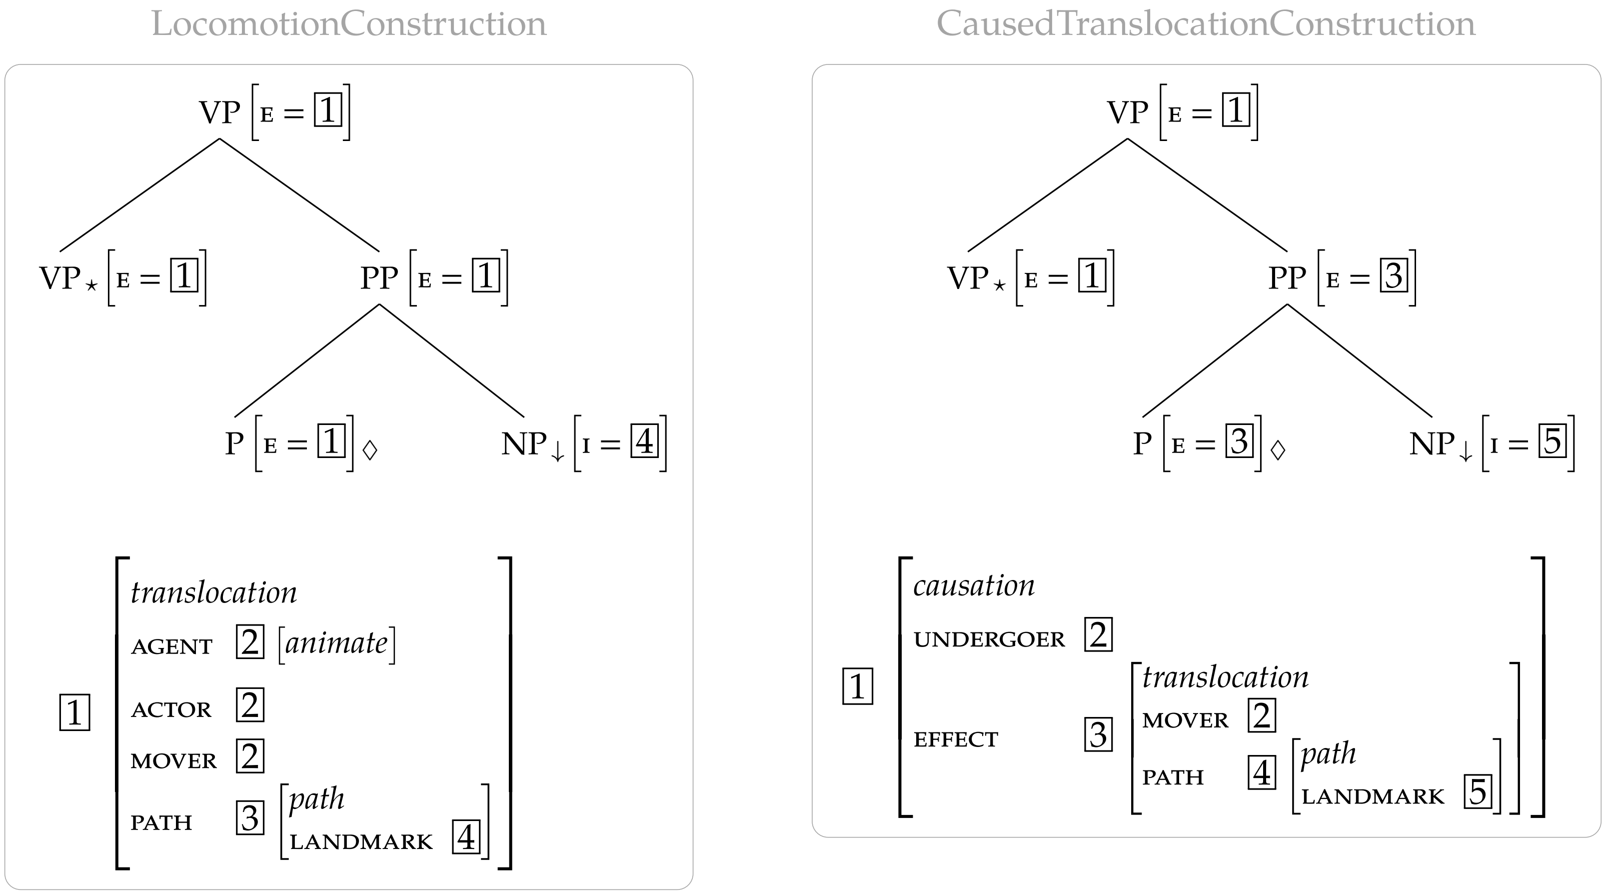


Figure 47: Trees and frames for the *LocomotionConstruction* and *CausedTranslocationConstruction* classes in the CMC/IAA metagrammar.

The *CausedTranslocationConstruction* realization of the *DirectedMotionAddingPath* class can adjoin to VP nodes whose semantic frame is compatible with the type *causation*. The EFFECT subframe of the frame contributed by *CausedTranslocationConstruction* will either be added to the adjunction site's frame, or unify with its EFFECT subframe in case it already exists. This construction is used for sentences like (40d) *Queequeg inserted the key into the lock*. In this case, the lexicon entry for *insert* also specifies a particular PATH that is required (see Figure 39 on page 116), so that directional prepositions other than *into* will not be accepted for this verb by the parser. (Of course, the metagrammar only covers a small number of tokens and constructions, and could be extended to allow for more possibilities if desired. For instance, the preposition *in* may be a valid choice in combination with *insert*, but *in* does not currently feature in the lexicon of this metagrammar.)

The remaining realizations of the *DirectedMotion* class are grouped under the class *DirectedMotionAddingCausationAndPath*. These constructions each contribute a *causation* frame into which the semantic frame of the adjunction site is embedded as the CAUSE subframe. The descriptions for the two classes *PrepositionalCausedMotionConstruction* and *PrepositionalCausedMotionConstructionAddingConstructionalObject* are shown in Figure 48 on page 131.

Both of these constructions import another class called *GeneralCausedMotion*, which in turn imports the classes *DirectionalAdjunct* (providing the PP subtree and a *translocation* frame with a value for its PATH attribute) and *Causation* (providing additional nodes in the VP subtree and the *causation* frame structure).

Each of the two classes shown in Figure 48 contains the PP substructure inherited from *DirectionalAdjunct*, including its semantic description. Inheriting that partial description from a dedicated class prevents redundancy and makes it possible to, for instance, update the metagrammar to make changes specifically to the semantic description of paths or to implement alternative realizations of prepositional phrases.

The *PrepositionalCausedMotionConstruction* class is the one enabling the parser to accept sentences like (41d) *Queequeg pushed Daggoo to Tashtego*. The non-directed variant of this sentence, (41c) *Queequeg pushed Daggoo*, is not associated with a *causation*-type frame, as discussed in Section 6.5. The construction adjoins to a VP node of the tree representing the non-directed sentence, adds a directional phrase and extends the semantics of the original sentence with a caused-motion meaning. The construction is special in a sense because its semantics are not embedded into the frame of the adjunction site, but the embedding actually happens the other way around: the construction is associated with a *causation* event, and the frame at the adjunction site is incorporated into the CAUSE subframe of the *causation* frame.

PrepositionalCausedMotionConstruction

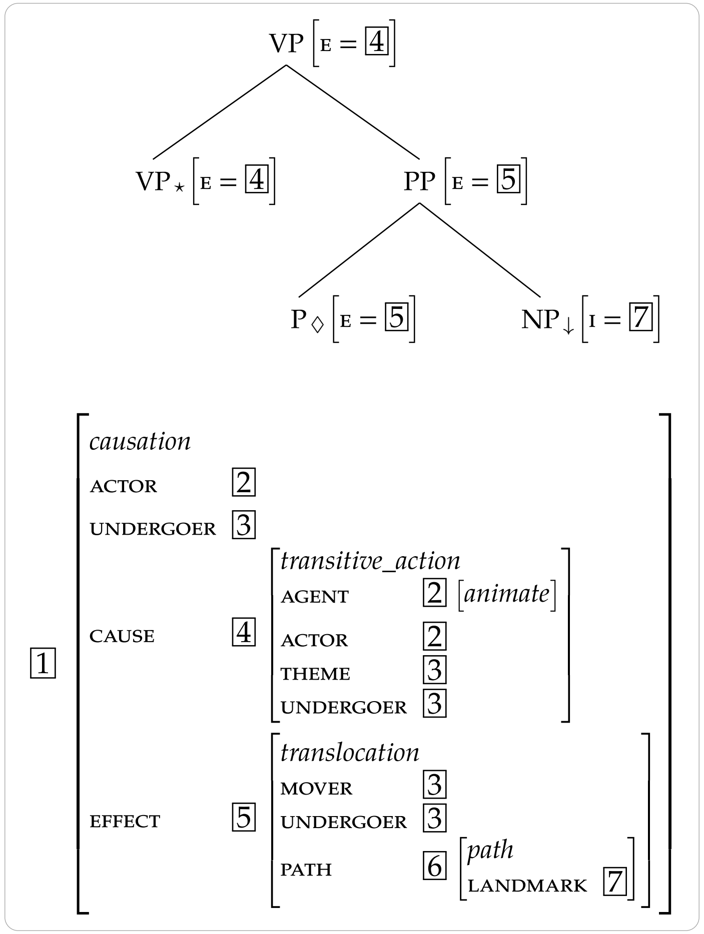


PrepositionalCausedMotionConstructionAddingConstructionalObject

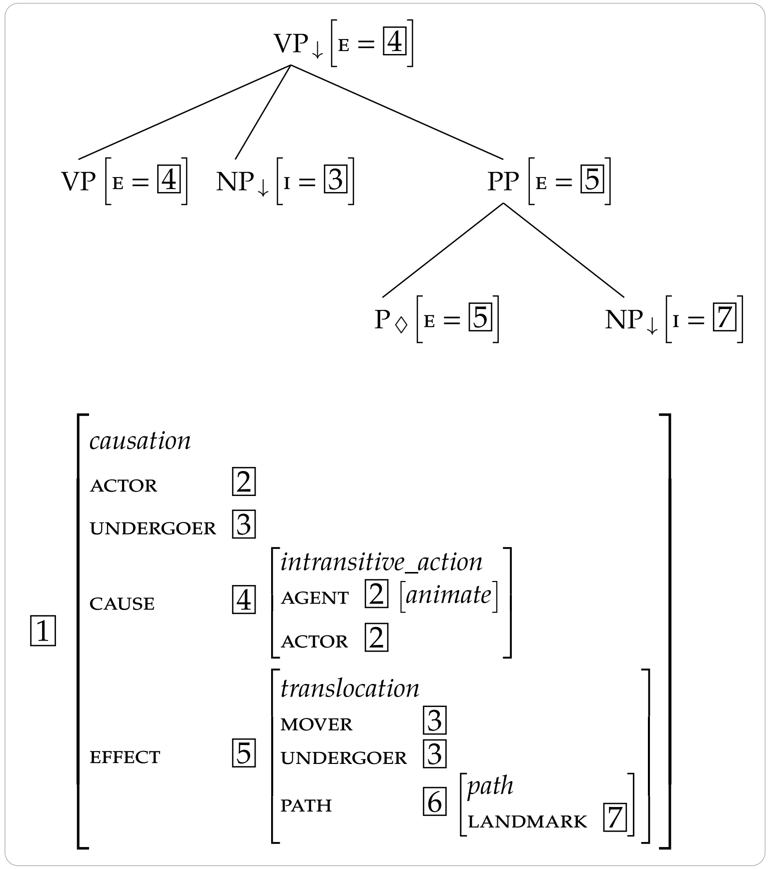


Figure 48: Tree and frame for the *PrepositionalCausedMotionConstruction* class (left) and the *PrepositionalCausedMotionConstructionAddingConstructionalObject* class (right) in the CMC/IAA metagrammar.

The other class, *PrepositionalCausedMotionConstructionAddingConstructionalObject*, is applicable to sentences like (38d) *Queequeg laughed Daggoo to Tashtego*. This sentence has no valid non-directed variant **Queequeg laughed Daggoo*; the object *Daggoo* is licensed exclusively by the caused-motion construction. This is why this metagrammar class adds an `NP` node in the syntactic object position, in addition to the syntactic and semantic contributions associated with the other caused motion class. The added `NP` subtree is contributed by the *Object* class from the second level of the class hierarchy, which is imported by *PrepositionalCausedMotionConstructionAddingConstructionalObject*. The entity expressed in the object position is assigned to the MOVER attribute of the EFFECT subframe. Of the two variants of this construction, one is applicable to verbs that are semantically transitive, while the other is applicable to verbs that are semantically intransitive. This is expressed by the required types for the value in the respective CAUSE subframes in the descriptions of the two variants.

Goldberg (Goldberg 1995: 43, Goldberg 2002: 342) distinguishes between participant roles, which are contributed by verbs, and argument roles, which are contributed by the constructions that are being instantiated. Both realizations of the caused-motion construction appearing in this metagrammar license a MOVER and a PATH. In the construction that is available to verbs like *push*, the MOVER is identified with the UNDERGOER licensed by the verb. In contrast to this, the construction that is available to verbs like *laugh* needs to add the syntactic description that allows the MOVER to appear in the object position, because the verb itself is typically intransitive. Thus, there is a mismatch in profiled roles for the construction as it applies to verbs like *laugh*, but there is no such mismatch in the alternative realization of the construction that applies to verbs like *push*.

The structure of the frames associated with the two caused motion classes is nearly identical, with the exception of the inclusion or exclusion of an UNDERGOER attribute in the CAUSE subframe. For both classes, the CAUSE subframe is unified with the frame representing the event described by the verb during parsing.

In sentences like (41d) *Queequeg pushed Daggoo to Tashtego*, the UNDERGOER, *Daggoo*, is licensed by the basic transitive use of the verb *push*, so this role will appear in the subframe that unifies with the frame representing the *pushing* event. On the other hand, in sentences like (38d) *Queequeg laughed Daggoo to Tashtego*, it is not clear that *Daggoo* truly is an UNDERGOER of the *laughing* event. This is why the CAUSE subframe of the *PrepositionalCausedMotionConstructionAddingConstructionalObject* class only accepts verbs that are compatible with the type *intransitive_action*, and the entity *Daggoo* is only assigned to the MOVER attribute in the EFFECT subframe and does not appear in the CAUSE subframe at all. One may argue that *Daggoo* must be impacted in some direct way by the action performed by *Queequeg* in order for the *causation* event to occur – however, this is not an inherent part of the *laughing* event, and the grammar will not license the addition of the UNDERGOER role to this type of event. This design decision also allows the construction to apply to sentences like Goldberg's *Frank sneezed the tissue off the table*, where the description of a *sneezing* event cannot meaningfully involve an UNDERGOER in the object position.

In both metagrammar classes, there is no type requirement for the entity that is assigned to the MOVER role in the EFFECT subframe. This entity is not required to be, for instance, animate; sentences like *Frank sneezed the napkin off the table* involve an inanimate MOVER, while sentences like (41d) *Queequeg pushed Daggoo to Tashtego* have an animate MOVER.

Defining the caused-motion construction as an auxiliary class reflects the productivity of the construction (see Goldberg 1995: 164): instead of adding this construction as a lexically-motivated realization for individual verbs, as part of one or more tree families, the metagrammar allows it to adjoin to any VP node with which it is compatible during parsing. This is the mechanism by which even verbs that are typically regarded as obligatorily intransitive and non-directed, like *sneeze* or *laugh*, can instantiate this construction.

The approach taken in this chapter concerning directional phrases differs slightly from that of Kallmeyer & Osswald (2013) concerning intransitive and transitive sentences with or without a directional prepositional phrase. While Kallmeyer & Osswald (2013: 276, 311–312) also define classes for basic intransitive or transitive uses of verbs (*n0V* and *n0Vn1*), and another class that is responsible for the tree and frame contributed by the prepositional phrase (*DirPrepObj*), they then combine these classes into additional tree families *n0Vpp*(*dir*) and *n0Vn1pp*(*dir*). Their metagrammar requires each verb to be lexically specified to anchor a particular subset of these four top-level tree families. In contrast to that, the metagrammar developed in this chapter allows verbs to anchor either the class corresponding to Kallmeyer & Osswald's *n0V* class, or the one corresponding to their *n0Vn1* class, and the appearance of directional prepositional phrases is licensed by adjoining auxiliary classes.

The compatibility of a verb with directional phrases can thus either be specified semantically as implemented here, via type constraints and the definition of the construction as an auxiliary class, or as implemented by Kallmeyer & Osswald, by allowing individual verbs to anchor tree families involving directional phrases. Both strategies can be used to distinguish between directional phrases that function as arguments and directional phrases that function as adjuncts. (While the metagrammar developed in this thesis is based on the assumption that all directional prepositional phrases are adjuncts, Kallmeyer & Osswald 2013: 298–299 also allow for argument prepositional phrases.) The approach chosen here reduces the responsibility of the lexicon and adds importance to the type hierarchy describing the relationships between the types of each lexical item covered by the grammar. This allows for a larger extent of factorization and generalization, and is therefore preferred for the metagrammar developed in this thesis. Both approaches are able to prevent overgeneration by allowing the parser to reject sentences involving certain verbs and directional prepositional phrases: in one scenario, there will simply be no tree family available to analyze that type of sentence, and in the other scenario, a basic intransitive or transitive tree family will be available for the verb, but the adjunction of the caused-motion construction licensing the directional prepositional phrase will be blocked via the type constraints. A verb like *laugh* would presumably be allowed to anchor *n0V* but not *n0Vpp*(*dir*) in the grammar presented by Kallmeyer & Osswald.

Beyond this difference, the approach by Kallmeyer & Osswald (2013: 299) also analyzes intransitive sentences with a directional prepositional phrase as simple *translocation* events (as implemented here with the *LocomotionConstruction* class), and transitive sentences with a directional prepositional phrase as *causation* events involving an ACTOR performing some kind of action that causes a *translocation* event involving a MOVER (as implemented here with the *PrepositionalCausedMotionConstruction* and *PrepositionalCausedMotionConstructionAddingConstructionalObject* classes).

### 6.6.5 *Tree families at the fifth level of the hierarchy*

At the highest level of the hierarchy, this metagrammar contains tree families that group related constructions together via disjunction. The *RegularIntransitiveVerbFamily* class is realized via a disjunction of the *IntransitiveActivityConstruction* class and the *InducedActionConstruction* class, and the *RegularTransitiveVerbFamily* class is realized via a disjunction of the *TransitiveActivityConstruction* class and the *LexicalCausationConstruction* class.

The induced action alternation is not covered by a separate verb family. Instead, the induced action construction is one of multiple possible instantiations of the verb family that is responsible for verbs like *laugh*, *dance* or *jump*. The type requirement in the construction – making the construction unavailable to all verbs that are not compatible with *agentive_manner_of_motion* – ensures that only verbs that participate in the alternation instantiate this construction. Beyond being able to instantiate the induced action construction, verbs in the alternation behave like any other intransitive verb.

Similarly, some, but not all, transitive verbs have a causative lexical meaning. Verbs like *insert* or *push* are in the same verb family, and the causative realization of this family is, again, only available to verbs in the family that have a certain semantic property. In this case, the verb's lexical frame must be compatible with the type *causation* in order for the verb to be able to instantiate the construction.

## 6.7 CONSTRUCTIONALLY-MOTIVATED ADDITIONS TO THE FRAME TYPE HIERARCHY

The semantic descriptions of the constructions discussed in the previous section require additional types and constraints to be added to the type hierarchy. The final type hierarchy will allow the parser to correctly accept or reject individual instantiations of constructions involving particular verbs and arguments.

In the original frame type hierarchy, the types associated with verbs were categorized broadly into intransitive actions, transitive actions, and causation events. The induced action construction requires instantiating verbs to be compatible with the type *agentive_manner_of_ motion*, so the compatibility of each event verb with this additional type must be encoded in the frame type hierarchy. The type *agentive_manner_of_motion* will be added as a subtype of *intransitive_action* (recall that Hale & Keyser 1986: 608 call these verbs "active intransitives"). This will automatically prevent verbs like *push*, which belongs to the type *transitive_action*, or *insert*, which belongs to the type *causation*, from instantiating the induced action construction. An additional constraint is also included specifying that *laugh*, while intransitive, is incompatible with *agentive_manner_of_motion*.

One effect of this treatment of verbs that participate in the induced action alternation is that they are always treated as intransitive, even when they appear in the induced action construction, which includes a direct object. The object is not licensed by a transitive sense of the verb, but is instead exclusively contributed by the construction. Goldberg (1995: 54) calls this a mismatch in the number of roles and illustrates this with the fact that the ditransitive construction profiles three roles, while the verb *kick* pro-

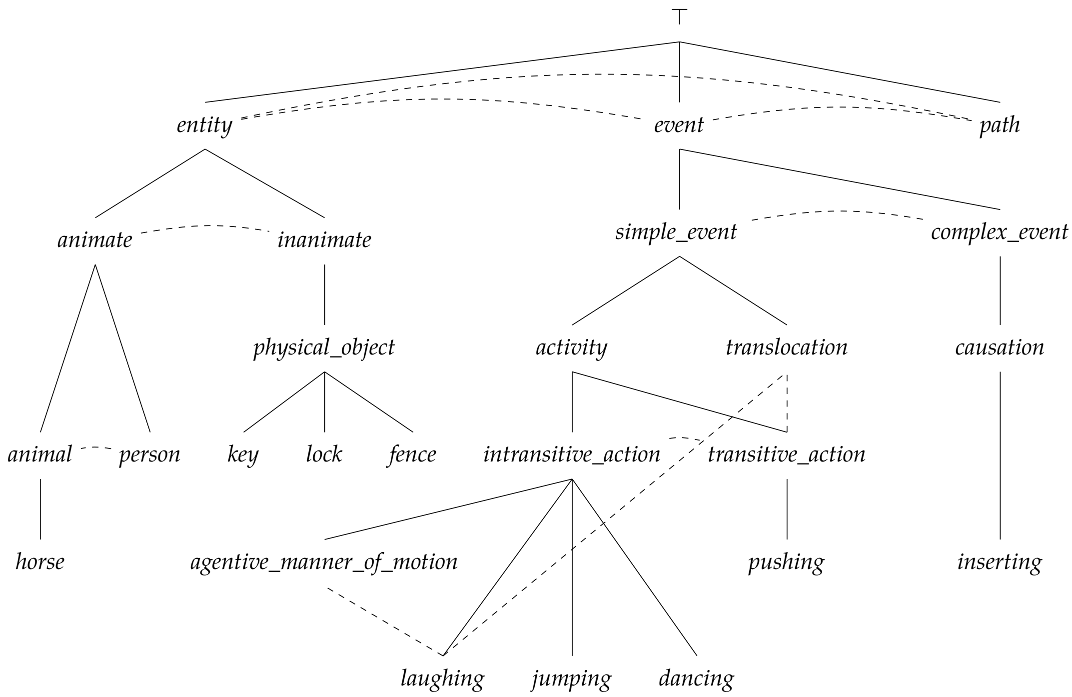


Figure 49: Extended event frame type hierarchy denoting subtype relations and type incompatibilities between the types in the metagrammar for CMC/IAA. Solid lines indicate subtype relations. Dashed lines indicate explicit incompatibility between types.

files two roles but can still appear in that construction. A sentence like (42c) *Queequeg jumped the horse* does not express a transitive *jumping* event performed by *Queequeg*, but instead an intransitive *jumping* event performed by *the horse*, which is triggered by some unnamed action by *Queequeg*.

Additionally, as argued in the previous section, the *translocation* type should be incompatible with *transitive_action* and its subtypes, and also specifically with *laughing*.

The extended frame type hierarchy including these additional constraints is shown in Figure 49.

All constructions discussed in this section impose certain constraints on the entities in their various semantic role slots. Androutsopoulos & Dale (2000) point out that a strict view of selectional restrictions can cause problems in contexts involving linguistic negation, where an entity can seemingly appear in a role slot without fulfilling the constraints associated with that slot. They illustrate this problem with the sentence given in (44).

(44) Tom cannot eat a keyboard.

Modeling linguistic negation in LTAG with frame semantics is, at the time of writing, an open problem. Lichte & Kallmeyer (2006: 84) propose an LTAG model for German negative polarity items, using the semantics framework from Kallmeyer & Romero (2008). In the work by Lichte & Kallmeyer, the presence of a negation is indicated by a global field called NEG, which is ordinarily set to a value of *no* unless a negation adjoins somewhere and switches it to *yes*. Such a field could conceivably be added to the metagrammar presented here as well to handle negation.

In order to add a representation of linguistic negation to the metagrammar, a negation construction class would need to be defined. In FrameNet (Fillmore, Johnson & Petruck 2003, Fillmore & Baker 2009), the *Negation* frame comprises two core frame elements, named *Factual_situation* and *Negated_proposition*. These two frame elements could also be included as attributes in a *NegationConstruction* class in this metagrammar. This would make it possible to assign the syntactic arguments to semantic roles that are separated from the frame representing the event expressed by the verb. This way, selectional restriction violations in the context of negation would not automatically prevent the parser from being able to accept a sentence.

However, negation is a bit more complex than that. Its scope is often ambiguous, so the construction responsible for negation would need to be defined as a disjunction of various available realizations. The acceptability of sentences involving selectional restriction violations under linguistic negation is not always clear. Because the focus of this thesis is on modeling various alternation-specific constructions following the linguistic literature on the alternations under discussion, the implementation of negation is left for future work. Due to the modular nature of XMG, the addition of a *NegationConstruction* class would require no changes or almost no changes in the definition of the classes already presented here.

## 6.8 DERIVATIONS FOR EXAMPLE SENTENCES

This section presents the derivations for a subset of the sentence types covered by the metagrammar developed in this chapter. The metagrammar files are compiled with the XMG-2 compilers for the relevant dimensions. The resulting grammar files are then loaded into the TuLiPA parser so that it can output all valid derivations for each input sentence. The following will only discuss a subset of the example sentences, but the grammar covers the remaining sentences as well. They can be parsed at any time using TuLiPA (Kallmeyer et al. 2008, Arps & Petitjean 2018) and the grammar code provided in the supplementary material for this thesis. The examples show that the metagrammar models the behavior of the verbs discussed here correctly and makes all required distinctions.

For each input sentence that is being parsed, the tree families anchored by each element in the sentence determine which constructions are potentially available. Which derivations are actually valid is based not exclusively on syntactic categories, but also on the type hierarchy in which type constraints for verbs and arguments are specified. Only derivations consisting of a valid tree with the symbol `S` as its root and a valid semantic frame are presented as possible analyses of the given input sentence. Each sentence may receive zero, one, or more valid analyses.

The derivations that will be discussed in the following include sentences instantiating various combinations of the constructions covered in the metagrammar.

First, sentence (39b) *Queequeg danced to Daggoo* will be parsed, instantiating the *IntransitiveActivityConstruction* and the *LocomotionConstruction* variant of the *DirectedMotion* class.

Second, sentence (40c) *Queequeg inserted the key* will be parsed, instantiating the *LexicalCausationConstruction*.

Third, sentence (38d) *Queequeg laughed Daggoo to Tashtego* will be parsed, instantiating the *IntransitiveActivityConstruction* and the *PrepositionalCausedMotionConstruction-AddingConstructionalObject* variant of the *DirectedMotion* class.

Finally, sentence (42d) *Queequeg jumped the horse over the fence* will be parsed, which has two different readings, one instantiating the *PrepositionalCausedMotionConstruction* variant of the *DirectedMotion* class, and another one instantiating the *InducedActionConstruction*.

### 6.8.1 *Derivation for basic intransitive directed sentences*

In the metagrammar, intransitive verbs like *dance* are lexically specified to anchor the *RegularIntransitiveVerbFamily* class. The derivation licensed by the metagrammar for the input sentence (39b) *Queequeg danced to Daggoo* is illustrated in Figure 50 on page 138. The final derived tree and frame for this sentence are shown in Figure 51 on page 139. In the interest of space, this is the only sentence for which the derived tree and frame are presented in addition to the derivation.

In this sentence, the syntactic subject, the agent, is identified with the actor macrorole in the semantic frame representing the meaning of the whole sentence. The directed motion construction that is instantiated by the directional prepositional phrase adjoins to the `VP` node of the construction anchored by the verb. The frames associated with each construction are unified. The resulting frame is of the conjunctive type *dancing, translocation*.

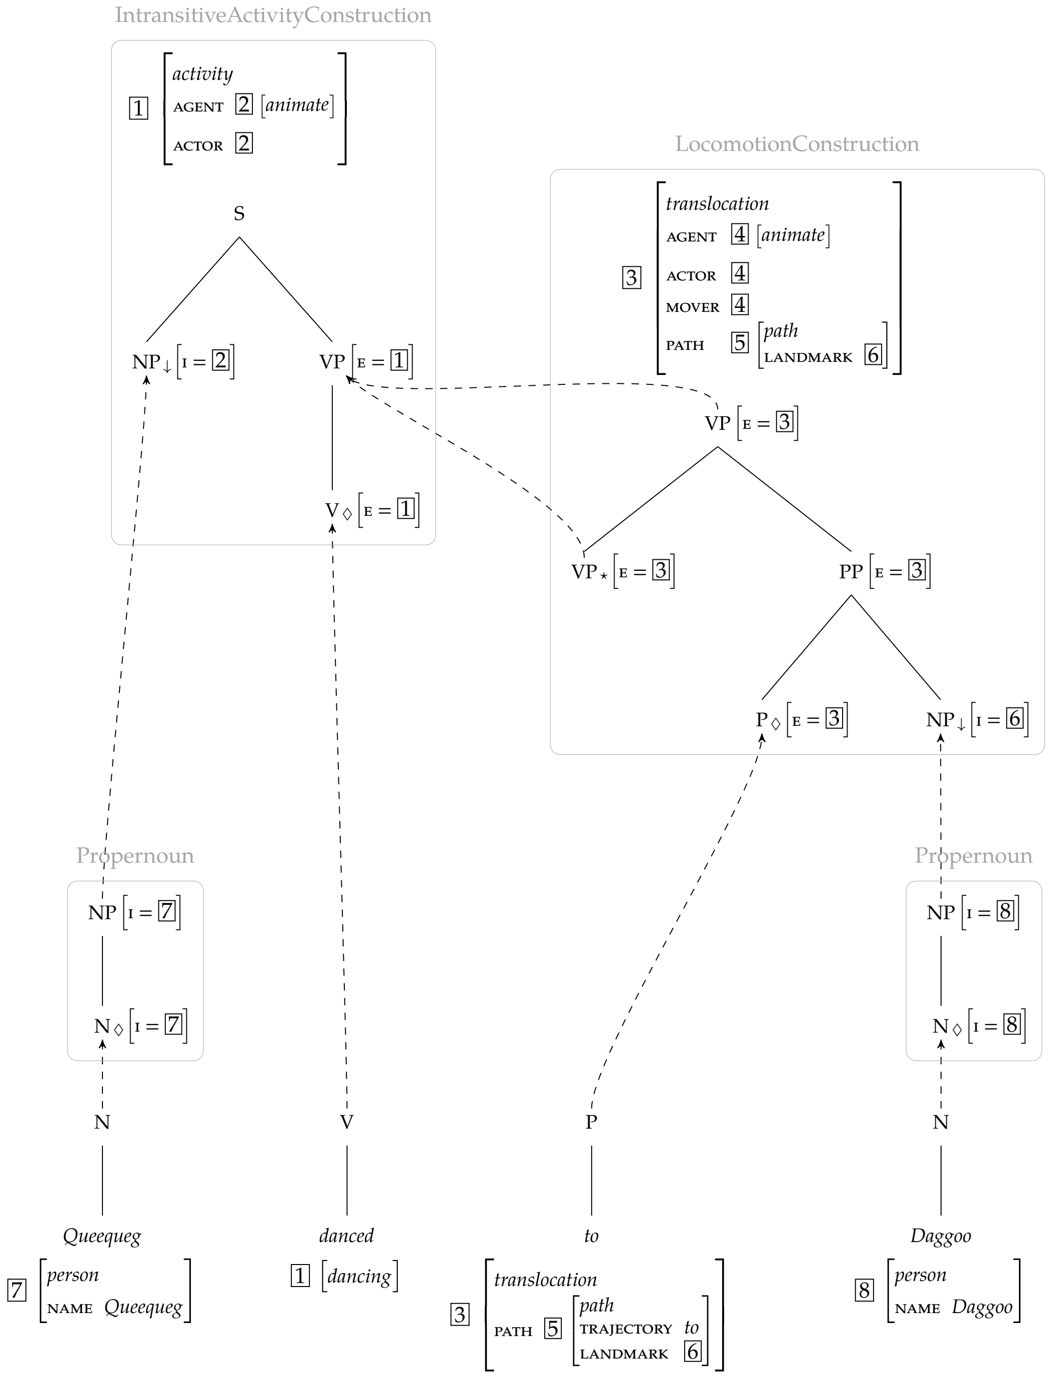


Figure 50: TAG derivation for sentences like (39b) *Queequeg danced to Daggoo*.

Derived tree and frame

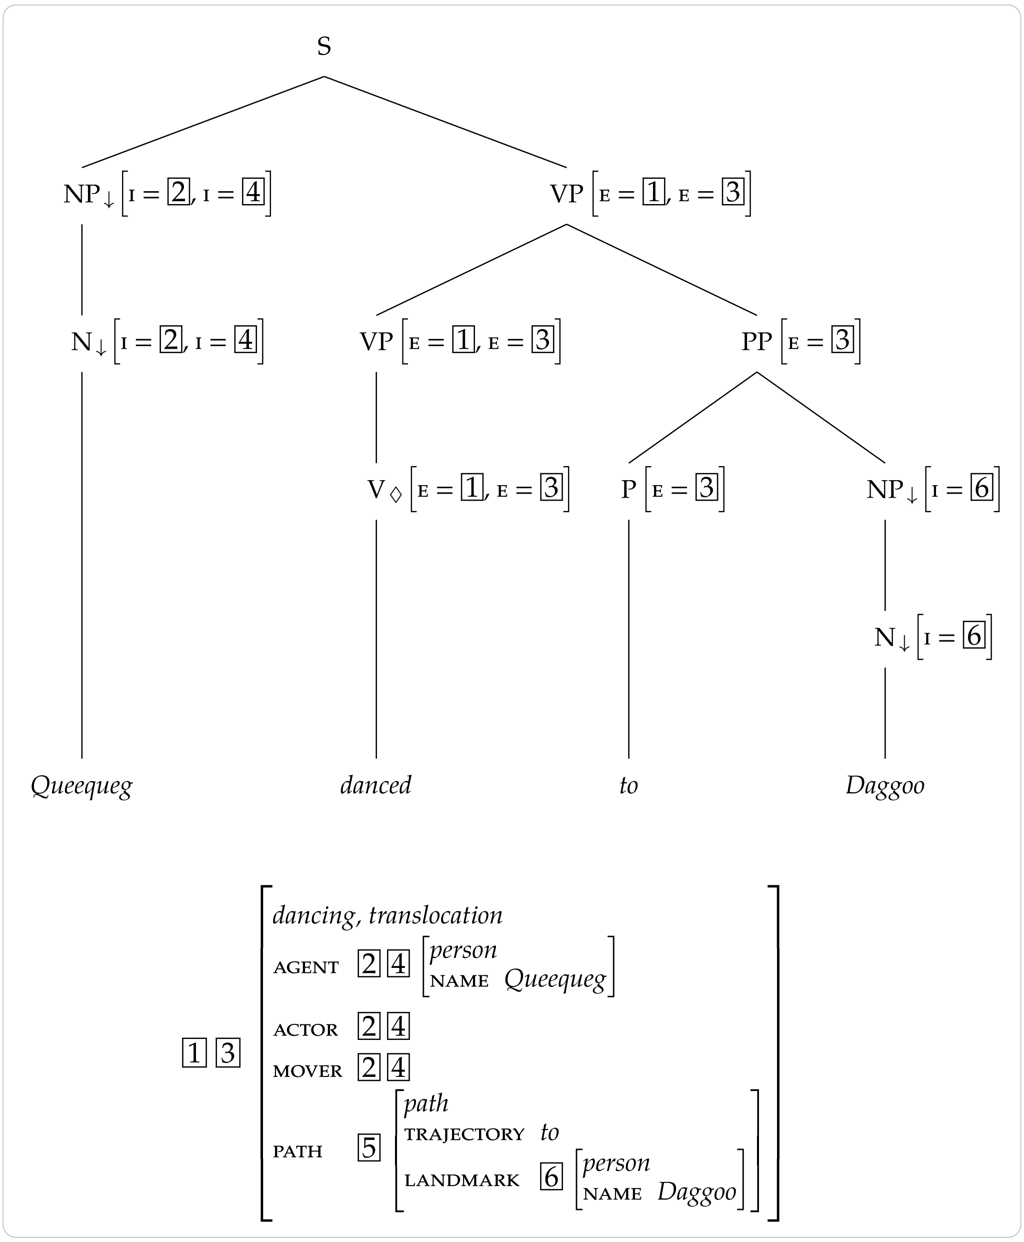


Figure 51: Derived tree and frame for sentence (39b) *Queequeg danced to Daggoo*.

### 6.8.2 *Derivation for transitive sentences with lexically causative verbs*

Verbs like *insert* can instantiate the *LexicalCausationConstruction*. The *causation* frame contributed by the construction is unified with the *causation* frame contributed by the verb's lexicon entry. The classes anchored by arguments of the verb unify with the nodes at which they are substituted into the construction tree. The derivation licensed by the metagrammar for the input sentence (40c) *Queequeg inserted the key* is shown in Figure 52 on page 141.

In the derivation, a PATH for the caused motion is included in the semantic frame for the sentence, but its LANDMARK is not filled with a specific value. The sentence describes an event in which *Queequeg* inserts *the key* into something, but no additional information as to the location is given. Sentences like (40d) *Queequeg inserted the key into the lock* would lead to a similar derivation, with the addition of an adjoining tree adding the directional prepositional phrase whose PATH subframe would unify with the value of the existing PATH attribute in the construction.

### 6.8.3 *Derivations for sentences instantiating the caused-motion construction*

The caused-motion construction can be instantiated with intransitive or transitive verbs that are not lexically causative. If a syntactic object is not contributed by the verb, it can be contributed by the construction. This is the case for sentences like (38d) *Queequeg laughed Daggoo to Tashtego*. The derivation for this sentence is shown in Figure 53 on page 142. The adjunction of the caused-motion construction at the VP node of the construction anchored by the verb leads to an overall sentence frame of the type *causation*. The frame contributed by the verb's lexicon entry unifies with the CAUSE subframe of the *causation* event.

This derivation illustrates the versatility of possible interactions between constructions and verbs. While the adjoining construction shown in Figure 51 merely adds a PATH attribute and its value to the frame provided by the verb's lexicon entry, the adjoining construction in Figure 53 constitutes an embedding of subframes in the other direction: here, the frame contributed by the verb is embedded as a subframe into the causative frame contributed by the construction. Contrasting these examples demonstrates that meanings contributed by verbs, arguments and constructions can combine in a number of different ways that can lead to resulting frame representations of varying complexity.

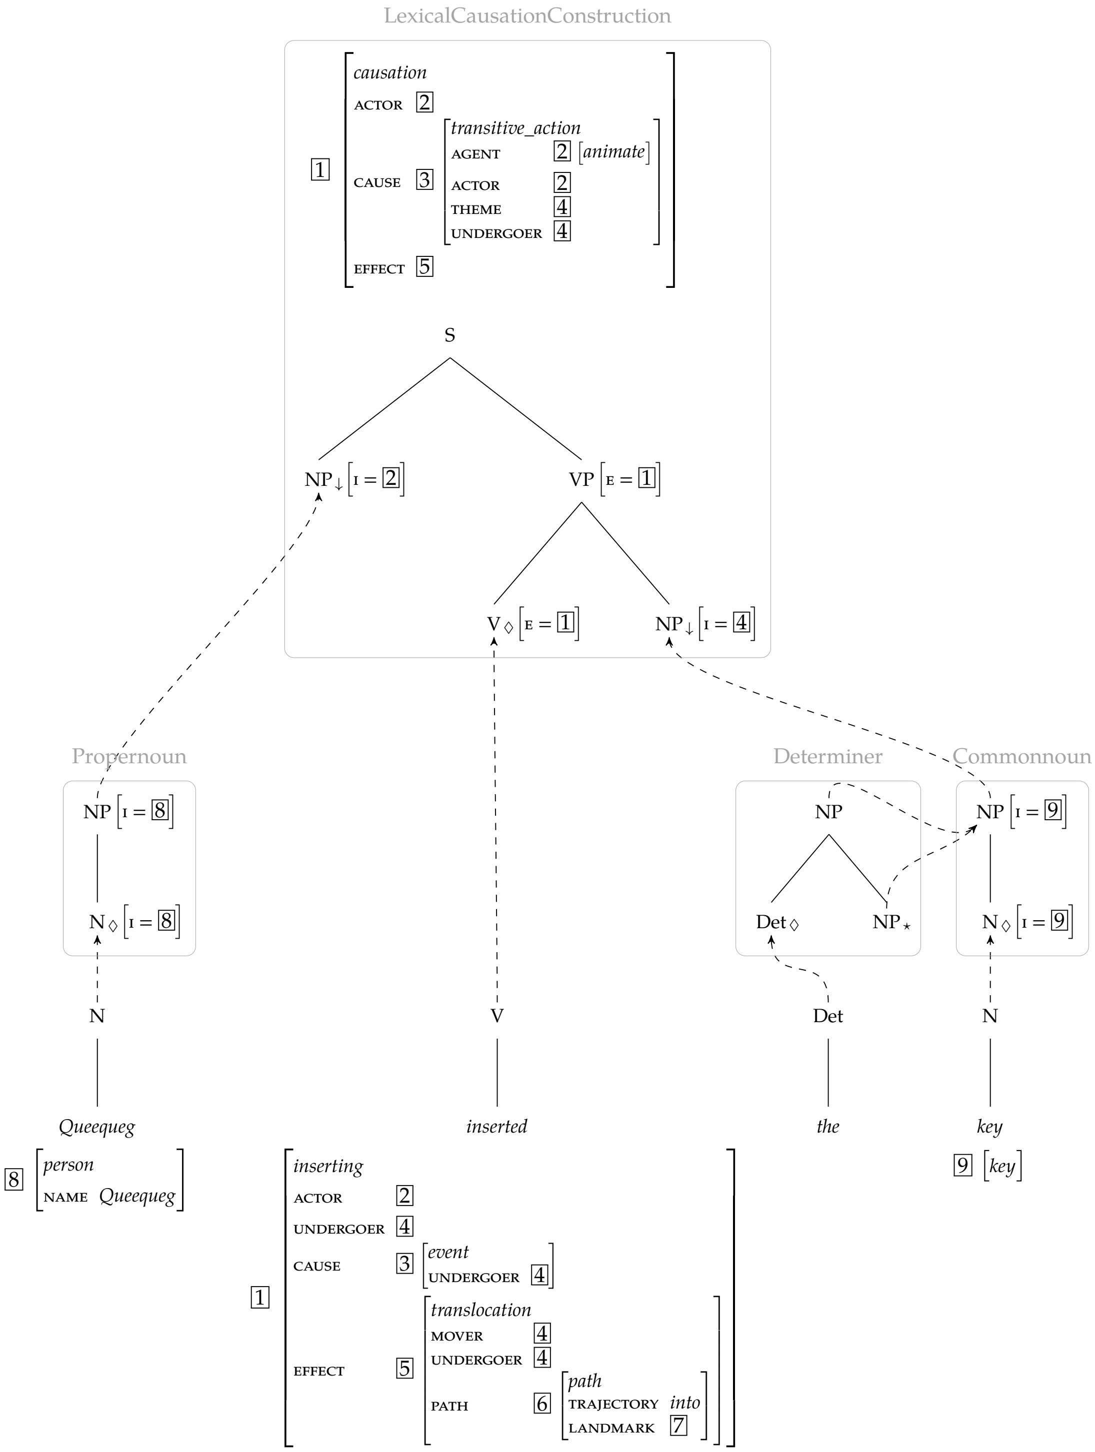


Figure 52: TAG derivation for sentences like (40c) *Queequeg inserted the key*.

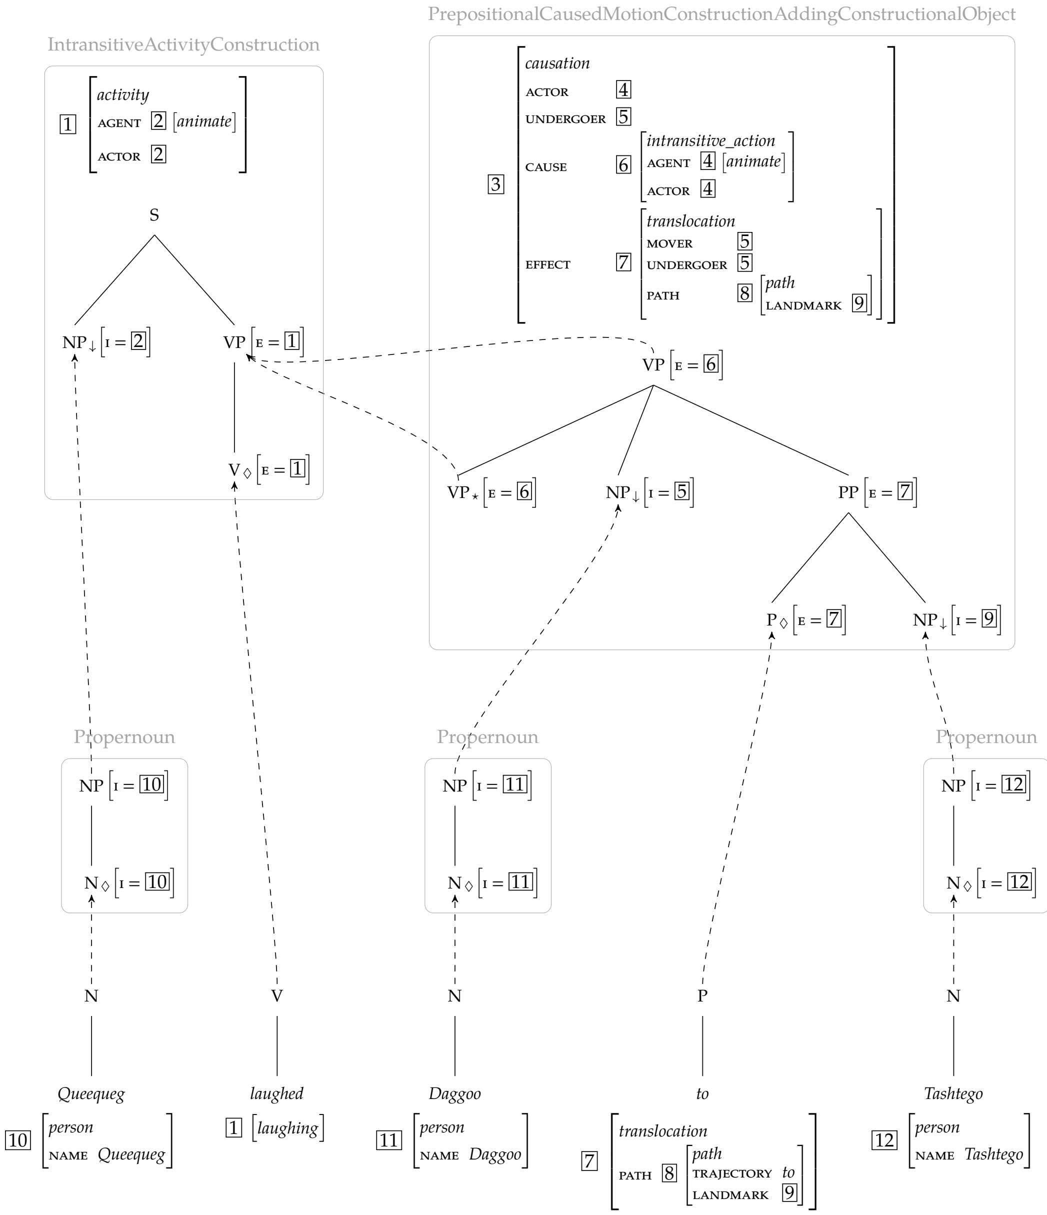


Figure 53: TAG derivation for sentences like (38d) *Queequeg laughed Daggoo to Tashtego*.

### 6.8.4 *Derivations for sentences with IAA-alternating verbs*

Sentences whose verb participates in the induced action alternation can be ambiguous. The induced action construction is available to verbs that are compatible with the *agentive_manner_of_motion* type, and the caused-motion constructions are typically also available to these verbs.

The caused-motion derivation for Sentence (42d) *Queequeg jumped the horse over the fence* is shown in Figure 54 on page 143. The derivation involves the same construction classes as the derivation shown in Figure 53.

An induced action derivation for Sentence (42d) *Queequeg jumped the horse over the fence* is also available. This derivation is shown in Figure 55 on page 144.

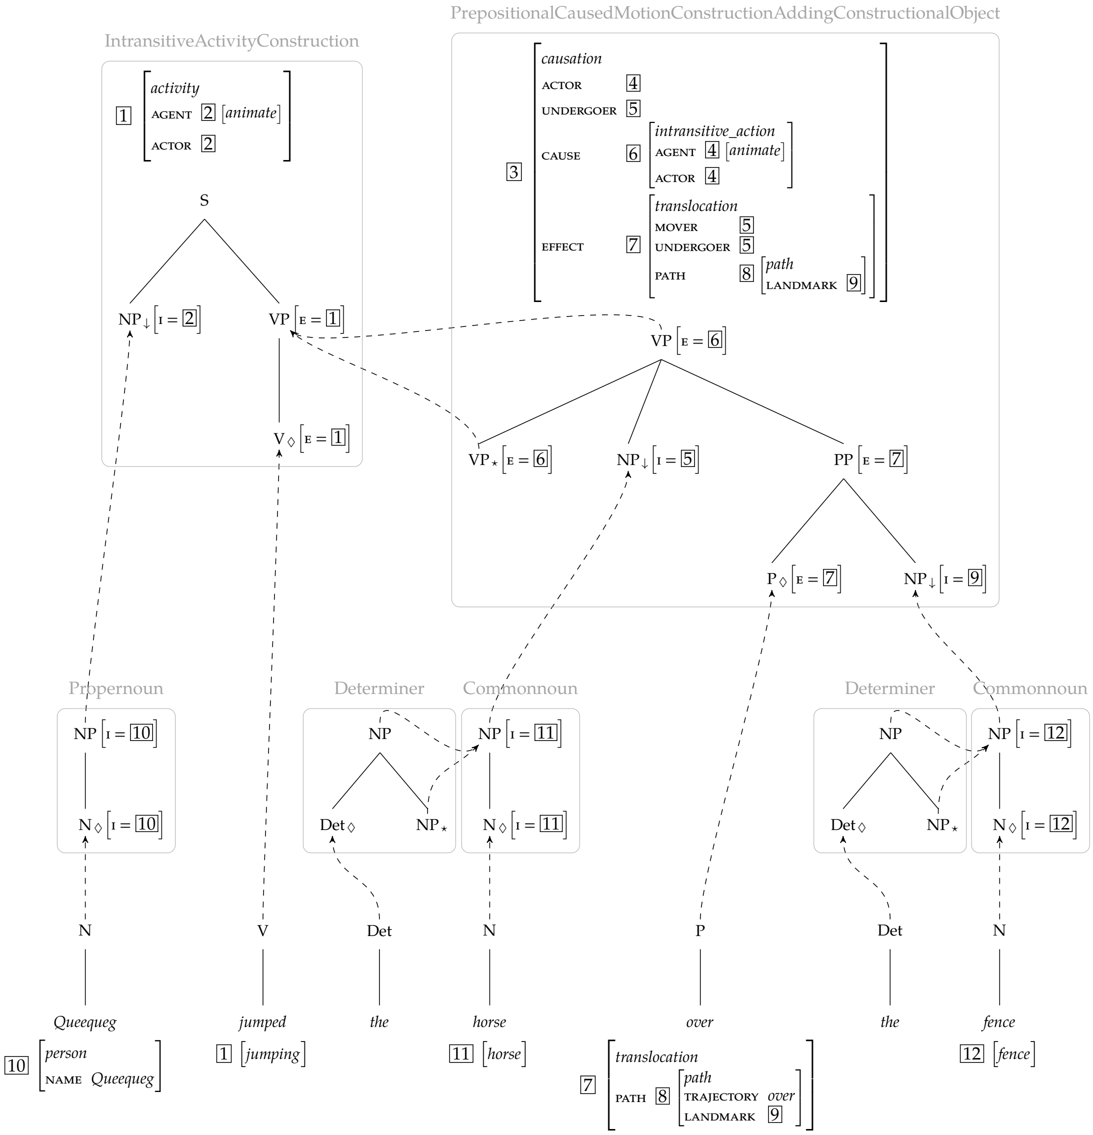


Figure 54: TAG derivation for sentences like (42d) *Queequeg jumped the horse over the fence* (CMC interpretation).

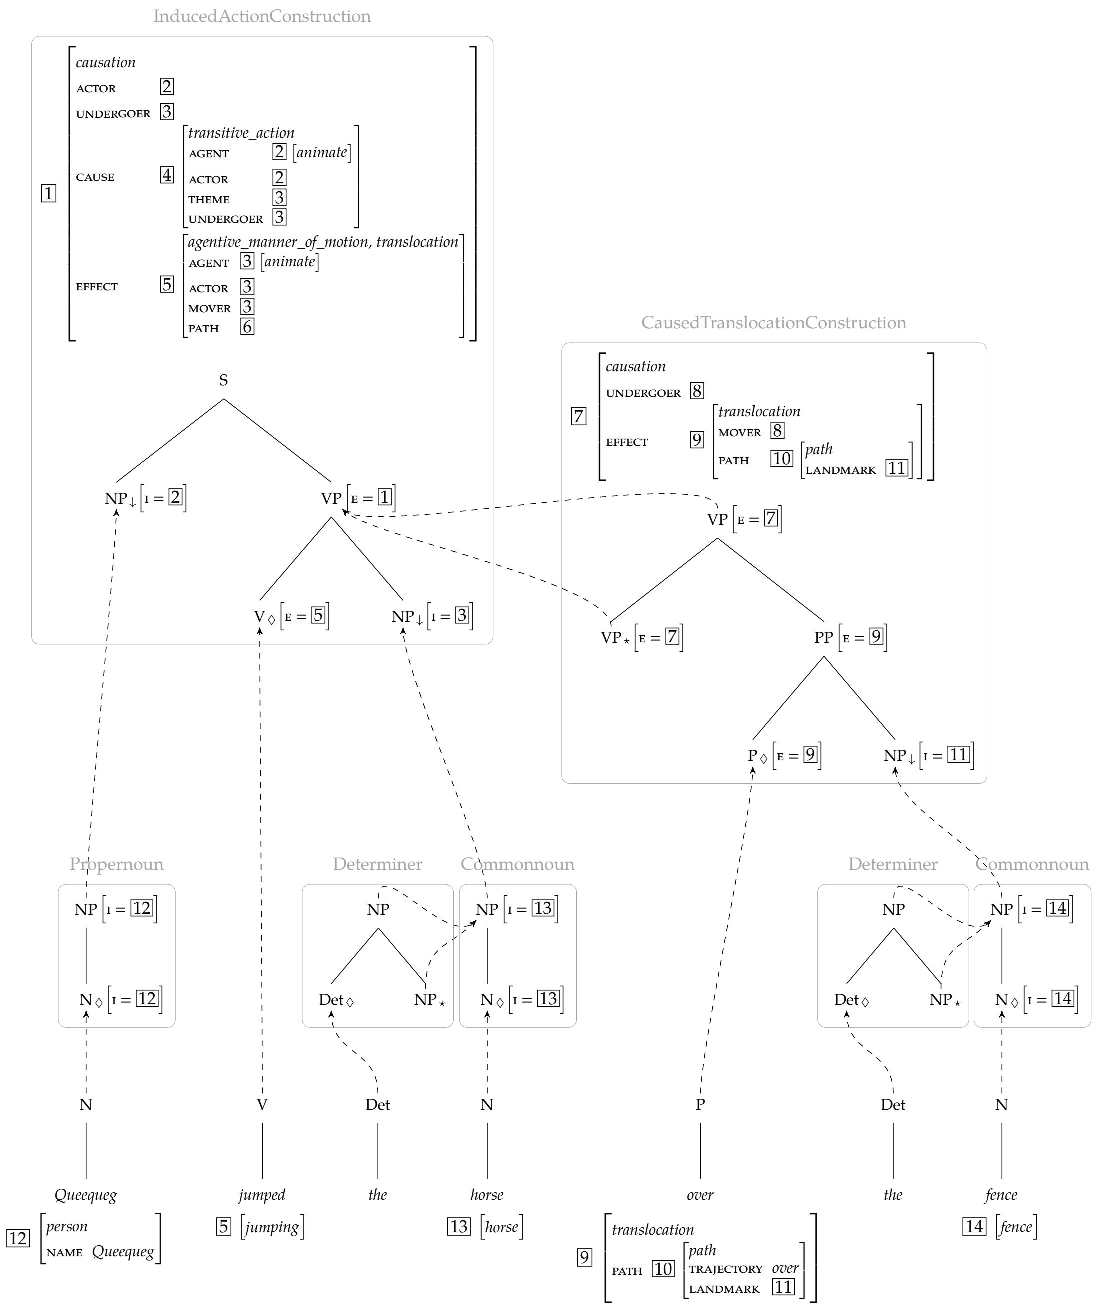


Figure 55: TAG derivation for sentences like (42d) *Queequeg jumped the horse over the fence* (IAA interpretation).

In the caused-motion reading (Figure 54), the *jumping* frame contributed by the sentence's verb is unified with the CAUSE subframe of the *causation* frame. This means that the syntactic subject (here: *Queequeg*) is the entity performing the *jumping* action. In contrast to this, in the induced action reading (Figure 55), the *jumping* frame contributed by the verb is unified with the EFFECT subframe. In this scenario, it is the syntactic object, *the horse*, that is performing the *jumping* action.

According to Cruse (1972: 521), the MOVER in a sentence instantiating the induced action alternation has to be an "obedient but independent" agent. This is encoded in the metagrammar with an animacy constraint for the moving entity in the induced action construction: only entities that are of a type compatible with *animate* can take on the AGENT role in the EFFECT subframe. For sentences involving an inanimate direct object, like *Queequeg blasted the door over the fence*, the induced action derivation is therefore not available. This type of sentence is only compatible with the caused-motion derivation, analogous to the derivation shown in Figure 54.

## 6.9 CONCLUSION

The metagrammar presented in this chapter focuses on sentences instantiating the caused-motion construction and the constructions involved in the induced action alternation. Both intransitive and transitive verbs can instantiate one of the variants of the caused-motion construction. The induced action construction can only be instantiated by agentive manner of motion verbs, which are sometimes referred to as active intransitives. Sentences involving the same syntactic structure can receive different semantic frame structures, depending on which category of verb is observed in the sentence.

The caused-motion construction is a productive phenomenon in English. Relatively general frame type constraints are used to encode in the relevant classes in the metagrammar which types of entities are allowed for certain semantic roles. Type constraints determine the availability of certain entity types as arguments in specific positions, as well as the (in)compatibility of each construction with particular sets of verbs.

Requirements imposed by verbs on their arguments can appear in the verb's lexicon entries, as in the case of *insert* here, which can only instantiate the caused-motion construction with directional phrases headed by *into* (see Figure 39 on page 116). This prevents the parser from accepting sentences like (40e) **Queequeg inserted the key from the lock*, while accepting similar sentences like (40d) *Queequeg inserted the key into the lock*. The verb *push* does not have such requirements, so that any directional phrase can be combined with it in instances of the caused-motion construction.

Note that there is an alternative reading of Sentence (40e) in which the prepositional phrase attaches to the `NP` instead of the `VP`. However, prepositional phrase attachment ambiguity is not the focus of this thesis, and has therefore not been modeled in the metagrammar. If desired, the existing classes in the metagrammar can be extended at a later time to also cover scenarios in which prepositional phrases adjoin to nominal phrases.

The availability of the induced action construction depends on the verb in a given input sentence. If the verb is compatible with the type *agentive_manner_of_motion* (Halliday 1967: 42, Hale & Keyser 1986: 608, Levin & Rappaport Hovav 1994: 40, Van Valin 2005: 34), it can instantiate the induced action construction, and the parser can attempt to derive a tree and frame structure for the sentence.

The frames associated with the caused-motion construction and the induced action construction differ with respect to the entity that is performing the action expressed by the verb. In caused-motion sentences, the sole AGENT, encoded in the syntactic subject position and functioning as the ACTOR, performs that action. In induced action sentences, the MOVER performs it as a secondary AGENT, and the exact action performed by the primary AGENT is not specified.

Sentences that instantiate the induced action construction, but also fulfill the type requirements of the caused-motion construction, lead to multiple derivations from the parser, since they are truly ambiguous.

The metagrammar developed to model these phenomena is inspired by the work by Kallmeyer & Osswald (2013), who model directed motion expressions with a focus on *translocation* and *caused-motion* events. Their metagrammar does not take the induced action alternation into account. In their treatment of directional prepositional phrases, the semantic frames representing paths are described at a higher level of detail than the frames for paths in the metagrammar presented in this chapter. Here, paths are represented in a simplified way, since the focus is on modeling the differences in the behavior of verbs belonging to different categories. However, the path descriptions from Kallmeyer & Osswald (2013) could easily be incorporated into the metagrammar developed here, without requiring any updates to the rest of the metagrammar. This is due to the modular design and the hierarchical structure of the metagrammar.

Among other things, Kallmeyer & Osswald (2013) are concerned with the question of how different pieces of information about a directed motion event are contributed by the verb and by the directional prepositional phrase. They model this with examples of verbs that describe the manner of a movement, but not the direction; the derivations for sentences involving such verbs then combine the partial *path* frames from the verb and from the prepositional phrase to represent the whole event. In the caused-motion construction modeled in this chapter, the motion event and its path can be expressed purely by the directional prepositional phrase, even when the verb does not express motion lexically. In these cases, the *translocation* event arises from the construction itself, and its path is described exclusively by the directional prepositional phrase.

Goldberg (1995: 161) lists several possible realizations of the caused-motion construction, each associated with different implications with respect to the nature of the causation or of the motion. The metagrammar presented in this chapter does not make such a distinction. However, different subtypes of constructions could be modeled with the tools employed here. For instance, Goldberg lists a sense of the construction that expresses the prevention of movement of the syntactic object in a certain direction. This sense of the construction is associated with verbs like *lock*, *keep* or *barricade*. The prevented-motion frame could either be added to the verbs' lexicon entries, similar to the way that direction requirements for the verb *insert* have been implemented here, or they could be encoded via additional constraints in the frame type hierarchy. Then, the prevented-motion construction could be specified to only be instantiatable by verbs that are compatible with the prevented-motion semantic type.

The derivations shown in this chapter based on the developed metagrammar illustrate how constructions and verbs interact to form the meanings of sentences. In some cases, the meaning of an adjoining construction adds a subframe to the frame provided by the verb's lexicon entry. In other cases, adjoining constructions provide a larger frame into which the lexical contribution by the verb is embedded as a subframe. These processes are in line with the idea formulated by Goldberg (1995, 2002, 2013) that constructions should be viewed as independent meaning-carrying elements whose contributions are as important as the contributions from instantiating verbs and their arguments.

# 7 METAGRAMMAR: THE CAUSATIVE-INCHOATIVE ALTERNATION AND THE INSTRUMENT SUBJECT ALTERNATION

This chapter is concerned with developing a metagrammar modeling the behavior of verbs instantiating the causative-inchoative alternation (CIA) and the instrument subject alternation (ISA). As in the previous chapter, the two phenomena modeled here have some overlap in their syntactic patterns, but the constructions at play differ in the semantics they assign to the relevant syntactic patterns. The metagrammar will allow a parser to derive the appropriate semantic frames jointly with the syntactic trees for each of the constructions covered.

For the purposes of this chapter, the phenomena will be modeled in a separate metagrammar. The relevant basic tree fragments from the previous chapter are reused, and construction-specific metagrammar classes are added. Chapter 8 describes the grammar resulting from combining and compiling the two metagrammars discussed here and in the previous section.

## 7.1 INTRODUCTION

This section will briefly describe the constructions covered by the metagrammar presented in this chapter, and discuss the differences between them that need to be modeled in order for the parser to correctly analyze input sentences instantiating them.

### 7.1.1 *The causative-inchoative alternation*

The causative-inchoative alternation (Halliday 1968: 184, Fillmore 1970a: 253, Vendler 1972: 210, Hale & Keyser 1986: 607, Croft 1986: 231, Levin 1993: 27) allows verbs to appear either in a causative construction or in an inchoative construction. Transitive uses of these verbs are typically causative, while intransitive uses are typically inchoative. Participating verbs typically express some change of state of a THEME. In the inchoative construction, the UNDERGOER is expressed as the syntactic subject; no external cause of the change of state is expressed in these sentences. In the causative construction, the syntactic subject is the position of the ACTOR that causes the change of state in the UNDERGOER, which is expressed in the syntactic object position. Examples for the causative and inchoative constructions are given in (45).

(45) a. *Ahab's compass* broke. (subject is UNDERGOER)
b. *Tashtego* broke *Ahab's compass*. (subject is ACTOR)

Sentences instantiating the inchoative construction, like (45a), express a change of state of the UNDERGOER, expressed in the syntactic subject position, with the exact nature of the change of state being described by the meaning of the verb. Sentences instantiating the causative construction, like (45b), express a caused change of state. They will be analyzed here as complex *causation* events, similar to the sentences instantiating the caused-motion construction and the induced action construction discussed in the previous chapter.

Whether an entity is a valid candidate for the UNDERGOER role of a change of state as expressed by a specific verb participating in the causative-inchoative alternation depends on the properties of that entity. For instance, the alternating verb *dry* requires that its UNDERGOER can have a "wetness" property whose intensity changes over the course of the *drying* event. For a detailed discussion of such state changes and possible ways to represent them in semantic frames, see Osswald & Van Valin (2014: 141).

Some selectional constraints determining alternation behavior will be expressed in the frame type hierarchy in the metagrammar, but they are not meant to be comprehensive. For instance, the alternating verb *empty* will only be able to take THEME arguments whose type is compatible with a type *container*. For more on these constraints, see Section 7.4.

### 7.1.2 *The instrument subject alternation*

The instrument subject alternation (Fillmore 1970b: 126, Levin & Rappaport 1988: 1071) allows verbs to appear either with an explicitly stated AGENT (functioning as the ACTOR) with or without an INSTRUMENT, or with the INSTRUMENT in the syntactic subject position (functioning as the ACTOR), and no explicitly stated AGENT. Examples of the agent subject construction and the instrument subject construction are given in (46).

(46) a. *Starbuck* scratched their backs with *his lance*. (subject is AGENT)
b. *His lance* scratched their backs. (subject is INSTRUMENT)

Most verbs in this alternation do not lexically require the inclusion of an INSTRUMENT argument in all usage contexts – in fact, the INSTRUMENT role is often optional, both for verbs participating in this alternation and those not participating in it. Not all verbs that can take an instrument allow the alternation, as illustrated in (47).

(47) a. *Stubb* eats the steak with *a fork*. (subject is AGENT)
b. * *A fork* eats the steak. (INSTRUMENT cannot be moved to the subject position)

According to Van Hooste (2018: 187), the availability of this alternation is based on requirements of the predicate – here, the verb *eat* – on the instrument in terms of the instrument's position on an actionality scale, which is a two-dimensional ranking taking into account both autonomy and animacy. A *spoon*, as in the example sentences above, would be placed fairly low on this scale, since it is not animate (animacy dimension) and cannot move without being controlled by some external entity (autonomy dimension). Van Hooste motivates the difference in acceptability between sentences like (46b) and sentences like (47b) with the different instrument actionality requirements of *scratch* and *eat*. In this perspective, alternation participation is not controlled exclusively by the verb's lexical properties, but also by the interaction of these lexical properties and requirements with the specific arguments in a given sentence.

In order to categorize different sorts of instruments, Van Hooste (2018: 189) proposes a frame type hierarchy similar to the ones developed in this thesis part. His frame type hierarchy expresses the properties located at different points of the two dimensions of the actionality scale as individual types, with which each example argument type can be compatible or incompatible. For instance, a *stick* is assigned to the

*Artifact* and + *Movable* types. Other types include *Specifically tailored*, *Para-autonomous*, *Semi-autonomous*, *Autonomous* (*proper*), ± *Animate*, ± *Sentient*, and ± *Organization*. The hierarchy proposed by Van Hooste is preliminary with no claim to completeness, but illustrates the level of detail at which entity types would have to be described in order to make their availability as instrument subjects for certain verbs completely predictable. In an XMG implementation, this fine-grained distinction of instrument types would need to be accompanied by a set of attribute constraints on the events denoted by verbs.

The entities that can fill the INSTRUMENT role of verbs participating in the instrument subject alternation are not limited to a fixed set of prototypical instruments for the event described by the verb. For instance, a *door* would typically be opened with a *key* or a *remote control*, but it is also possible to use a *crowbar* or a *credit card* as an INSTRUMENT to open a door. What seems to matter are not so much fixed entity categories, and more the affordances of objects, which can be contextual as in the scenario of a *credit card* opening a door.

In the metagrammar to be developed in this thesis chapter, these different sets of requirements will merely be sketched. The affordance requirements for instruments for each event verb are encoded via attribute constraints in the type hierarchy. The actionality requirements for each verb will not be spelled out, but instead, a single type *autonomous* will be created with which certain entities will be compatible and others incompatible. The metagrammar class describing the instrument subject construction will then require the instrument subjects substituting into it during derivation to be compatible with *autonomous*, while also fulfilling the affordance requirements imposed by the anchoring verb. Objects like *spoon* are typically discussed as unlikely instrument subjects in the literature on the instrument subject alternation; however, a corpus analysis may yield a handful of attestations of such objects in that position. As the metagrammars developed in this thesis are meant to be a first step towards representing alternations in terms of constructions as an LTAG model with semantic frames, there is no claim to completeness. Instead, the model reflects the tendencies identified by the linguistic work referenced throughout this thesis. A corpus-driven extension and refinement of the type hierarchy with respect to instrument subject candidates like spoons is a promising direction for future work.

Sentences instantiating the instrument subject alternation's agent subject construction typically express the INSTRUMENT, if it is present, in a prepositional phrase headed by *with*. *With* is a highly polysemous preposition and can express a range of meanings beyond the one associated with instrument use; however, since the phenomena under investigation here do not involve other senses of *with*, the polysemy of this preposition will not be discussed further. For the same reason, issues of preposition attachment ambiguity will not be taken into account. For more on possible issues arising from the assumption that *with* encodes INSTRUMENTS, see Section 9.4 of this thesis.

Alternatively, instruments can be expressed in an adjunct phrase headed by *using*. The metagrammar to be developed here will not take this variant of the agent subject construction with an added instrument into account. In Part IV of this thesis, *using* sentences from corpora will be used as possible indicators for the participation of individual verbs in the instrument subject alternation.

$$\begin{bmatrix} \textit{becoming\_broken} & \\ \text{UNDERGOER} & \boxed{1}\ \textit{the bucket} \\ \text{THEME} & \boxed{1} \\ \text{RESULT} & \begin{bmatrix} \textit{being\_broken} & \\ \text{THEME} & \boxed{1} \end{bmatrix} \end{bmatrix} \qquad \begin{bmatrix} \textit{causation} & \\ \text{ACTOR} & \boxed{1}\ \textit{Tashtego} \\ \text{UNDERGOER} & \boxed{2}\ \textit{the bucket} \\ \text{CAUSE} & \begin{bmatrix} \textit{activity} & \\ \text{AGENT} & \boxed{1} \\ \text{ACTOR} & \boxed{1} \\ \text{THEME} & \boxed{2} \\ \text{UNDERGOER} & \boxed{2} \end{bmatrix} \\ \text{EFFECT} & \begin{bmatrix} \textit{becoming\_broken} & \\ \text{UNDERGOER} & \boxed{2} \\ \text{THEME} & \boxed{2} \\ \text{RESULT} & \begin{bmatrix} \textit{being\_broken} & \\ \text{THEME} & \boxed{2} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Figure 56: Semantic frames for sentences instantiating the inchoative construction (48a) *The bucket breaks* (left) and the causative construction (48b) *Tashtego breaks the bucket* (right) of the causative-inchoative alternation.

Note that verbs can participate in more than one alternation. For instance, the verb *break* can be used with the constructions of both the causative-inchoative alternation and the instrument subject alternation, as illustrated in (48).

(48)
a. The bucket breaks. (*inchoative*)
b. Tashtego breaks the bucket. (*causative, agent subject*)
c. Tashtego breaks the bucket with a hammer. (*causative, agent subject*)
d. The hammer breaks the bucket. (*instrument subject*)

As illustrated in these sentences, both alternations under investigation here allow different semantic roles to be associated with the syntactic subject position, provided some selectional preferences of the observed verb for the relevant role are met.

### 7.1.3 *Modeling differences between the causative-inchoative alternation and the instrument subject alternation*

Verbs that participate in the causative-inchoative alternation express either their ACTOR or their UNDERGOER in the subject position. The frame contributed by the inchoative construction is unified with the frame contributed by the verb instantiating it. In contrast to this, the causative construction contributes a *causation* frame with whose EFFECT subframe the instantiating verb's lexical frame is unified. Figure 56 illustrates the different representations. For more on possible alternative ways to represent these meanings in semantic frames, see Osswald & Van Valin (2014: 140) and Seyffarth (2018).

Figure 57 on page 151 shows the semantic frames for two example sentences instantiating the two constructions associated with the instrument subject alternation. Here, both frames are of the type *causation*. The agent subject construction with an added prepositional phrase expresses both the AGENT and the INSTRUMENT roles, while the instrument subject construction only expresses the INSTRUMENT and leaves the AGENT unspecified.

$$\left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{causation}} \\
\textsc{actor} & \boxed{1}\ \textit{Tashtego} \\
\textsc{undergoer} & \boxed{2}\ \textit{the bucket} \\
\textsc{cause} & \left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{activity}} \\
\textsc{agent} & \boxed{1} \\
\textsc{actor} & \boxed{1} \\
\textsc{theme} & \boxed{2} \\
\textsc{undergoer} & \boxed{2} \\
\textsc{instrument} & \textit{the hammer}
\end{array}\right] \\
\textsc{effect} & \left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{becoming_broken}} \\
\textsc{undergoer} & \boxed{2} \\
\textsc{theme} & \boxed{2} \\
\textsc{result} & \left[\begin{array}{l}
\textit{being_broken} \\
\textsc{theme}\ \boxed{2}
\end{array}\right]
\end{array}\right]
\end{array}\right]
\qquad
\left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{causation}} \\
\textsc{actor} & \boxed{1}\ \textit{the hammer} \\
\textsc{undergoer} & \boxed{2}\ \textit{the bucket} \\
\textsc{cause} & \left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{activity}} \\
\textsc{instrument} & \boxed{1} \\
\textsc{actor} & \boxed{1} \\
\textsc{theme} & \boxed{2} \\
\textsc{undergoer} & \boxed{2}
\end{array}\right] \\
\textsc{effect} & \left[\begin{array}{ll}
\multicolumn{2}{l}{\textit{becoming_broken}} \\
\textsc{undergoer} & \boxed{2} \\
\textsc{theme} & \boxed{2} \\
\textsc{result} & \left[\begin{array}{l}
\textit{being_broken} \\
\textsc{theme}\ \boxed{2}
\end{array}\right]
\end{array}\right]
\end{array}\right]$$

Figure 57: Semantic frames for sentences instantiating the agent subject construction (48c) *Tashtego breaks the bucket with a hammer* (left) and the instrument subject construction (48d) *The hammer breaks the bucket* (right) of the instrument subject alternation.

The inchoation construction and the instrument subject construction are each specific to one of the two alternations, and are unavailable to verbs that do not participate in the respective alternation. The agent subject construction with a *causation* meaning can be instantiated by all verbs that are compatible with a causative construction, provided that the arguments meet the requirements from the verb's lexicon entry and from the construction.

While there is an overlap in the sets of verbs associated with the causative-inchoative alternation and the instrument subject alternation according to VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007), it is not necessarily the case that any verb that participates in the causative-inchoative alternation can also be used in all constructions of the instrument subject alternation, and vice versa. A metagrammar that models these phenomena should take those differences into account and be able to derive all possible analyses for verbs that participate in none, either, or both of these two alternations.

As in the previous chapter, the metagrammar to be developed here will cover the constructions discussed above as well as some basic syntactic environments in which the featured verbs can appear. The sentence types that will be covered are presented in the following section.

## 7.2 SENTENCE TYPES COVERED BY THIS METAGRAMMAR

The metagrammar presented in this chapter will focus on the constructions that characterize the causative-inchoative alternation and the instrument subject alternation. The grammar covers verbs that participate in the causative-inchoative alternation, verbs that participate in the instrument subject alternation, verbs that participate in both alternations, and also non-alternating verbs.

The alternations under investigation in this chapter license sentences that are intransitive or transitive, with an optional *with* phrase (in the case of the causative-inchoative alternation), or transitive sentences with or without a *with* phrase (in the case of the instrument subject alternation). In the previous chapter, it was possible for sentences to simultaneously instantiate both the caused-motion construction and the induced action alternation, licensing two competing derivations for them. The examples to be covered in the present chapter also involve a category of verbs that participate in both the causative-inchoative alternation and the instrument subject alternation. However, this does not lead to competing derivations for individual sentences. Instead, all constructions associated with all relevant alternations will be available to such verbs. The verb *break* will exemplify this by licensing all sentence patterns associated with the constructions of the causative-inchoative alternation, as well as all patterns associated with the constructions of the instrument subject alternation.

One main focus of this metagrammar is the issue which semantic role is expressed in the syntactic subject position of a given sentence. For the causative-inchoative alternation, either the ACTOR or the UNDERGOER can appear in the subject position, while the semantic role that can appear in the subject position for verbs participating in the instrument subject alternation is always the ACTOR, which can be identified with either the AGENT or the INSTRUMENT. The grammar will include type constraints in the relevant constructions that determine which entities can be associated with these different roles. The non-alternating verbs that will be covered in this metagrammar exclusively encode the AGENT (functioning as the ACTOR) in the subject position.

Some example sentences with verbs participating in the causative-inchoative alternation, verbs participating in the instrument subject alternation, and nonalternating verbs are given in (49) through (52).

The verb *eat*, which features in the sentences given in (49), participates in the unexpressed object alternation. This allows it to appear either in transitive or intransitive environments; the effects of this will be discussed later. Only the AGENT can be expressed in the syntactic subject position. An INSTRUMENT can optionally be expressed in a prepositional phrase headed by *with*.

(49) a. Queequeg eats.
b. * The soup eats.
c. Queequeg eats the soup.
d. Queequeg eats the soup with the spoon.
e. * The spoon eats the soup.
f. Queequeg eats with the spoon.
g. * The spoon eats.

Example sentences for verbs that participate in the causative-inchoative alternation are given in (50). Either the ACTOR or the UNDERGOER can be expressed in the subject position. Transitive sentences involving these verbs express a *causation* event whose EFFECT subframe corresponds to the *inchoation* event expressed by intransitive sentences involving the same verb and UNDERGOER (see Osswald & Van Valin 2014: 140 and Section 4.2 of this thesis).

(50) a. ? Queequeg empties.
b. The bucket empties.
c. Queequeg empties the bucket.
d. Queequeg empties the bucket with the spoon.
e. * The spoon empties the bucket.
f. * Queequeg empties with the spoon.

g. * The spoon empties.
h. Queequeg empties the bucket with the pump.
i. The pump empties the bucket.

Sentence (50a) is not valid unless the person named *Queequeg* is regarded as the UNDERGOER of a *becoming_empty* event, which is not an expected role for this type of entity. Sentence (50f) is invalid because the ACTOR is expressed in the syntactic subject position, which makes a direct object necessary in which the UNDERGOER is expressed.

An instrument can appear with these verbs in the form of a *with* phrase, but the instrument *the spoon* cannot be expressed in the syntactic subject position of the verb *empty* as in (50e) or (50g). A non-autonomous entity like *spoon* cannot perform the causative action of *emptying something* without being controlled by an agentive entity. In contrast to this, sentence (50i) is accepted: a *pump* is a more autonomous entity that can change the state of another entity to be *empty* even without being controlled by an animate entity. Note, again, that these distinctions are simplified, and that Van Hooste (2018: 189) proposes a number of additional properties that distinguish entities like *spoon*, *pump*, and other conceivable instruments for *becoming_empty* and similar events.

A set of examples for verbs that participate in the instrument subject alternation is given in (51). Only transitive uses of this verb are accepted. Either the AGENT or the INSTRUMENT can appear in the subject position and function as the ACTOR.

(51) a. * Queequeg slices.
b. * The coconut slices.
c. Queequeg slices the coconut.
d. Queequeg slices the coconut with the machete.
e. The machete slices the coconut.
f. * Queequeg slices with the machete.
g. * The machete slices.

Finally, the behavior of verbs that participate both in the causative-inchoative alternation and in the instrument subject alternation is illustrated in (52). Intransitive uses of such verbs are allowed, but only if the UNDERGOER is expressed in the subject position. Transitive uses can express either the AGENT or the INSTRUMENT in the subject position.

(52) a. * Queequeg breaks. (*disallow inchoative reading*)
b. The window breaks.
c. Queequeg breaks the window.
d. Queequeg breaks the window with the hammer.
e. The hammer breaks the window.
f. * Queequeg breaks with the hammer.
g. ? The hammer breaks. (*disallow causative reading*)

## 7.3 REUSING PARTS OF THE CMC/IAA METAGRAMMAR

Thanks to the modular fashion in which the previous metagrammar was designed, the majority of metagrammar classes defined there can be reused in this metagrammar. All tree fragments on the first level of the hierarchy, all syntactic functions on the second level of the hierarchy, and all diathesis alternatives on the third level of the hierarchy are applicable in the current metagrammar in the same form as in the previous one. A number of construction classes and tree families on the fourth and fifth levels of the hierarchy will be defined specifically to cover the sentence types relevant in the current chapter.

The descriptions of lexical items in the `<morpho>`, `<lemma>` and `<frame>` dimensions follow the same principles as the descriptions in the CMC/IAA metagrammar (see Section 6.4), and will not be discussed in detail here.

The preposition *with* will be defined similarly as, but not identically to, the prepositions featured in the previous metagrammar. Directional prepositions such as *to* were described with a lexical frame of the type *translocation* with a PATH attribute whose nature depends on the individual preposition, while the preposition *with* will be represented with a frame of the type *activity*, with an INSTRUMENT role slot that can be filled by the prepositional object during parsing. Directional prepositions were able to anchor the class *DirectedMotion*, whereas *with* will anchor a class called *InstrumentWithConstruction*. The metagrammar classes containing the descriptions required for the correct handling of *with* are shown in Listing 9.

```
% For the synframe compiler:
class FrameWith
declare ?W ?X0
{
  <frame>{
    ?W[activity,
    instrument: ?X0]
  };
  <iface>{
    [e=?W]
  }
}
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% For the lex compiler:
class LemmaWith
{
  <lemma> {
    entry <- "with";
    sem <- FrameWith;
    cat <- p;
    fam <- InstrumentWithConstruction
  }
}
```

Listing 9: Lexical description for *with* in the CIA/ISA metagrammar

The remainder of this chapter will focus on the frame type hierarchy for the CIA/ISA metagrammar and on the descriptions of the constructions that are involved in the sentence types covered by the grammar.

## 7.4 THE LEXICAL FRAME TYPE INHERITANCE HIERARCHY FOR THE CIA/ISA METAGRAMMAR

The lexical frames for the lemmas appearing in this metagrammar are ordered in the type hierarchy as shown in Figure 58 on page 156. More general types are located at the top, more specific types are located at the bottom. Incompatibility between types is indicated with dashed lines. Sibling types at the leaf level are always mutually exclusive; the incompatibility relations between these types are not shown in the figure for the sake of readability. The same is true for the types *entity*, *event* and *state*, which are also mutually exclusive.

At the highest level, the types in this hierarchy fall into the three categories *entity*, *event* and *state*. The *entity* type contains all entities that can be arguments of verbs in the example sentences to be covered by the metagrammar. The *event* type covers all events that can be expressed by verbs and constructions in the metagrammar. Finally, the *state* type covers all states of entities that can come about as a result of one of the state-changing verbs in the metagrammar.

The *inanimate* type has one subtype, *physical_object*. The *physical_object* type has the three subtypes *inedible*, *edible* and *solid*; the former two are mutually exclusive, while the *solid* type is compatible with either of them. Events like *eating* require an UNDERGOER that is compatible with *edible*, while events like *becoming_broken* require their UNDERGOER to be compatible with the *solid* type, without a distinction between edible and inedible entities. The UNDERGOER role for *becoming_empty* events can only be filled by entities that are compatible with *container*. The type *bucket* is a subtype of *container*, while all other types at the leaf level under *inedible* are explicitly incompatible with *container*.

Additional subtypes are available under *inedible* to cover specific selectional constraints on instruments. Events of the type *slicing* require their INSTRUMENT role to be filled by an entity that is compatible with *sharp_object*. Causative events whose EFFECT subevent is of the type *becoming_empty* require the INSTRUMENT role in the CAUSE subframe to be filled by an entity that is compatible with *can_move_substances*. Events of the type *eating* require their INSTRUMENT role to be filled by an entity that is compatible with *can_manipulate_food*. The latter two instrument type requirements are rough approximations of the required affordances of instruments in these events.

Concerning the verbs in this metagrammar, three of them fall into the type *simple_event*, while one, *slicing*, lexically expresses a *causation*, which is a subtype of *complex_event*.

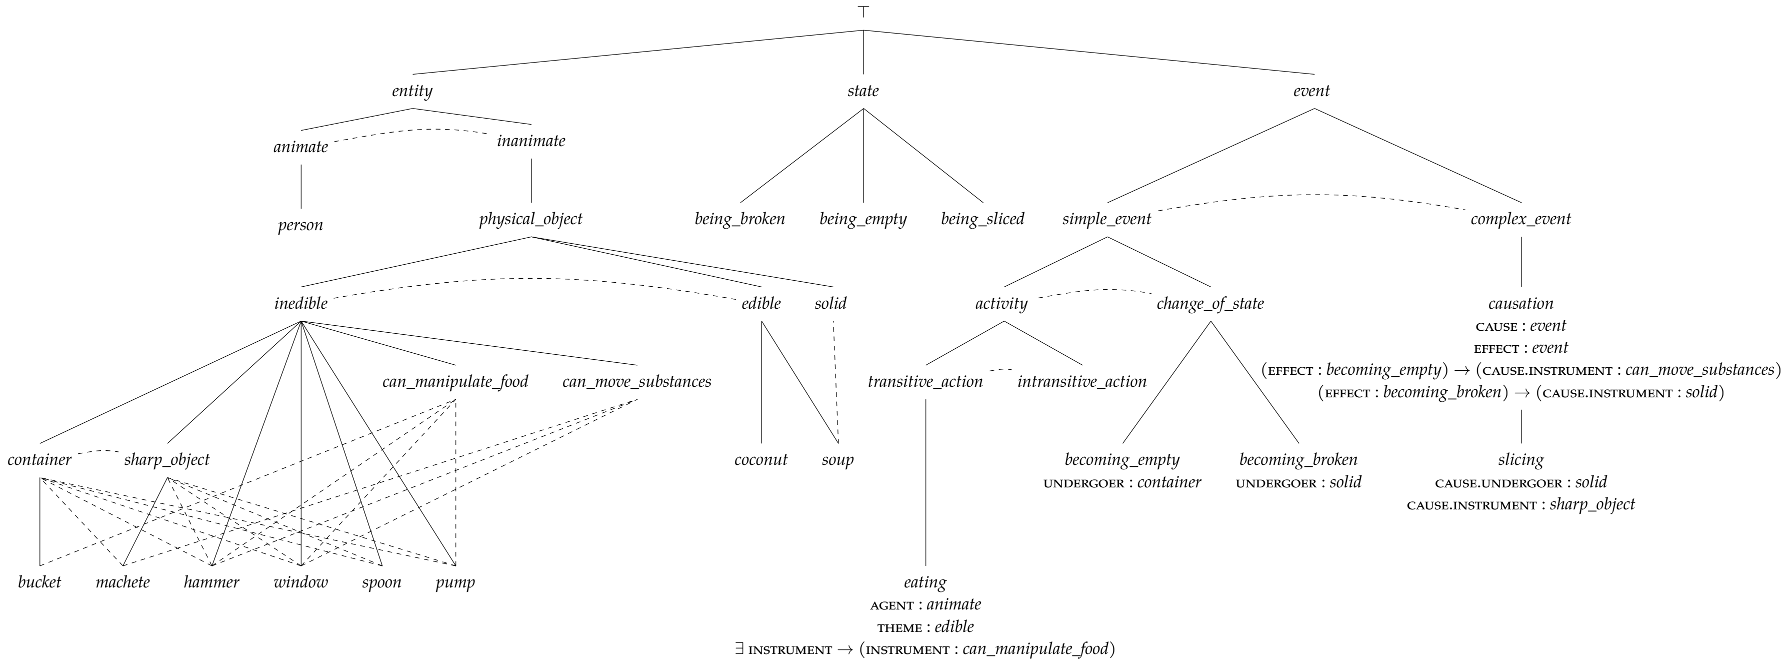


Figure 58: Lexical frame type hierarchy denoting subtype relations and type incompatibilities between the frames assigned to the lexicon entries in the metagrammar for CIA/ISA. Solid lines indicate subtype relations. Dashed lines indicate explicit incompatibility between types.

Among the three simple event verbs, one is an *activity* verb, and the others are subtypes of *change_of_state*. In this implementation, the *activity* verb *eating* is always regarded as transitive, since a THEME may be unexpressed but is always understood for *eating* events (see e.g. Kallmeyer & Osswald 2013: 292). Note that the subtype relationship between *eating* and *transitive_action* describes these events on a semantic level, while it will still be possible for the verb *eat* to appear in a syntactically intransitive environment. The semantic frame for such a sentence would then still include a THEME element in the *eating* frame, whose description merely contains all conditions imposed on it by the lexicon entry for *eating*, without being linked to any specific entity expressed on the surface.

When *change_of_state* verbs appear in a sentence, the frame representing this change will also include the result state of the UNDERGOER. The specific result states associated with each change of state verb are part of the verbs' lexicon entries. Osswald & Van Valin (2014: 141) go into more detail and provide examples that explicitly spell out different changes of state with respect to the state of the UNDERGOER at the beginning of the event and its state at the end or after the end of the event. In the context of this thesis, the focus is not on the specific progression of the state changes in these types of events, so the frames do not need to express these details. Due to the modular nature of metagrammars developed in XMG, the implementation presented here can be a basis for a more detailed model of state changes in the future.

Attribute constraints for certain types are included in the hierarchy graph. For instance, the UNDERGOER of a *becoming_empty* event needs to be compatible with the *container* type. Some verb-specific requirements are also associated with the *causation* type: for causative sentences whose EFFECT is of the type *becoming_broken*, the INSTRUMENT role, which is located inside the CAUSE subframe, must be compatible with the type *solid*.

Attribute constraints are either conditioned on existence (such as the INSTRUMENT in *eating* events) or not (such as the INSTRUMENT in *slicing* events). In other words, constraints on the type of an *eating* event's INSTRUMENT only matter in contexts where an INSTRUMENT is explicitly given. In contrast to this, constraints on the type of a *slicing* event's INSTRUMENT will hold in all contexts. This is because *eating* events do not necessarily involve an INSTRUMENT, while it is central to the lexical semantics of *slice* that a sharp object is used to cut something into pieces of a certain shape. Levin & Rappaport Hovav (2013: 67) describe the semantics of *slice* as an event that "is brought about through a well-defined use of a specialized instrument", and proceed to refer to that specialized instrument as "knife-like". This strong association of *slicing* events with a certain type of instrument is why the INSTRUMENT role will always be included in semantic frames for *slicing* events, even when an instrument is not explicitly mentioned, as in sentence (51c) *Queequeg slices the coconut*.

## 7.5 construction descriptions and tree families

The inheritance hierarchy of the metagrammar classes responsible for modeling all elementary trees and their frames in this metagrammar is shown in Figure 59 on page 158. All classes on the first three levels of the hierarchy are defined exactly as described in Section 6.6 of this thesis and will not be discussed here in detail.

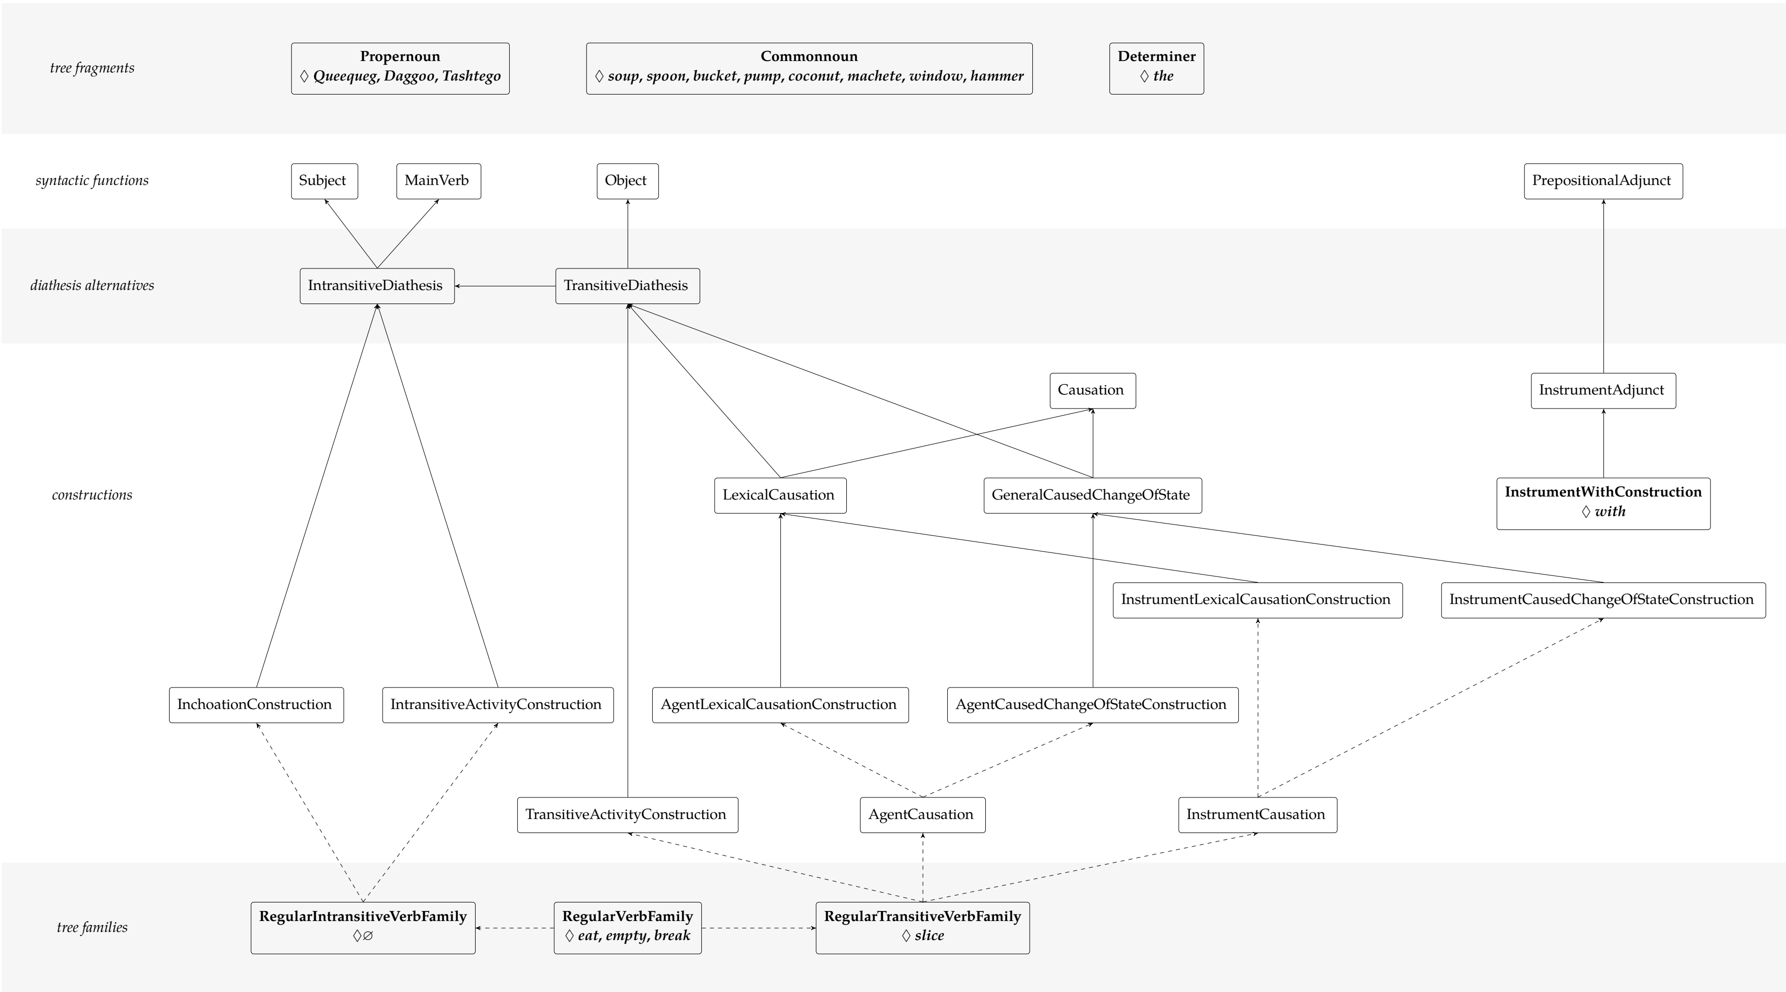


Figure 59: Overview of the class inheritance hierarchy described in file `syn_dimension.mg` of the metagrammar for CIA/ISA. Solid arrows denote structure inheritance. Dashed arrows denote disjunction.

On the fifth level of the hierarchy, the *RegularIntransitiveVerbFamily* and *RegularTransitiveVerbFamily* classes differ slightly from their corresponding classes in the CMC/IAA grammar. Here, the *RegularIntransitiveVerbFamily* class is not directly anchored by any lemmas, but is instead one of the available realizations of the added *RegularVerbFamily* class, covering verbs that can either appear in intransitive or transitive environments, such as *eat*, *empty* or *break*.

The *RegularTransitiveVerbFamily* class has three alternative variants with which it can be realized. These cover the basic transitive construction, *TransitiveActivityConstruction*, as well as the causative classes *AgentCausation* and *InstrumentCausation*. Similarly, the *RegularIntransitiveVerbFamily* class can be realized with the basic *IntransitiveActivityConstruction* class or with the *InchoationConstruction*, which is specific to the causative-inchoative alternation.

This means that all constructions involving a transitive syntactic environment will be associated with the same tree family, and all intransitive constructions are also grouped into one tree family. The different semantics of these constructions – for instance, the inchoative construction's assignment of the UNDERGOER role to the entity in the subject position – will be compatible or incompatible with an input sentence based on the types of the observed verb and its arguments. For example, in order for a sentence to instantiate the *InchoationConstruction* class, the verb's lexical frame must be compatible with the frame contributed by the construction with which it will unify, and the subject argument's frame must be compatible with the type constraints on the UNDERGOER attribute of the verb's frame. The type constraints will be discussed for each construction in the following. The construction definitions and the type requirements contained in them will require an extension of the type hierarchy that will be summarized in Section 7.6.

The following sections will present the constructions in this metagrammar in three categories. The first category covers the inchoation construction. The second category covers the different causation constructions. The third category covers the construction that adds an INSTRUMENT to an event via a prepositional phrase headed by *with*. The constructions covering basic activity verb uses, *IntransitiveActivityConstruction* and *TransitiveActivityConstruction*, are defined identically to the classes with the same names in the previous chapter, and will not be discussed again.

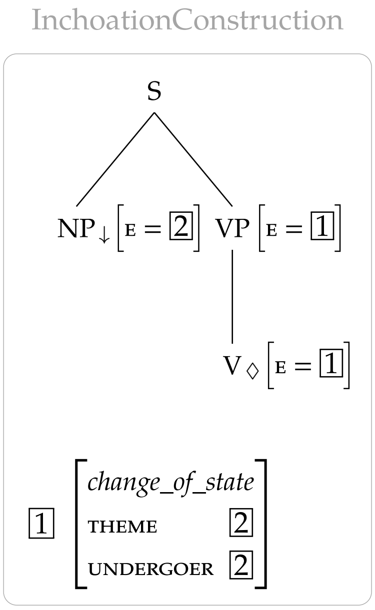


Figure 60: Tree and frame for the *InchoationConstruction* class in the CIA/ISA metagrammar.

### 7.5.1 *The inchoation construction*

The description of the *InchoationConstruction* class is given in Figure 60. This class identifies the semantic frame contributed by the subject with the THEME and UNDERGOER attributes of the event described by the verb. Since the frame contributed by the verb's lexicon entry is unified with the entire frame contributed by the construction, the lexical type of the verb must be compatible with the type *change_of_state*.

The *InchoationConstruction* class covers sentences like (50b) *The bucket empties*. The prevention of sentences like (50a) **Queequeg empties* is ensured by existing attribute constraints in the lexical frame type hierarchy specifying type requirements for the values of the UNDERGOER attributes of the verb frames.

### 7.5.2 *The causation constructions*

The two causation construction families are called *AgentCausation* and *InstrumentCausation*. These constructions differ minimally in their semantic representations: the *AgentCausation* class contains an AGENT attribute in the CAUSE subframe, while the *InstrumentCausation* class contains an INSTRUMENT attribute in the CAUSE subframe. In both constructions, the role that is expressed in the syntactic subject position is identified with the ACTOR semantic macrorole. The descriptions for the two construction classes are shown in Figure 61 on page 161.

Each of these two causative constructions has two variants: one that directly unifies with verbs whose semantic frame is of the type *causation*, and another one that embeds non-causative verbs in the EFFECT subframe. Internally, both constructions inherit from a more general, abstract class called *GeneralCausedChangeOfState*, which in turn inherits its syntactic properties from the *TransitiveDiathesis* class and its semantic properties from the *Causation* class that already appeared in the metagrammar presented in the previous chapter.

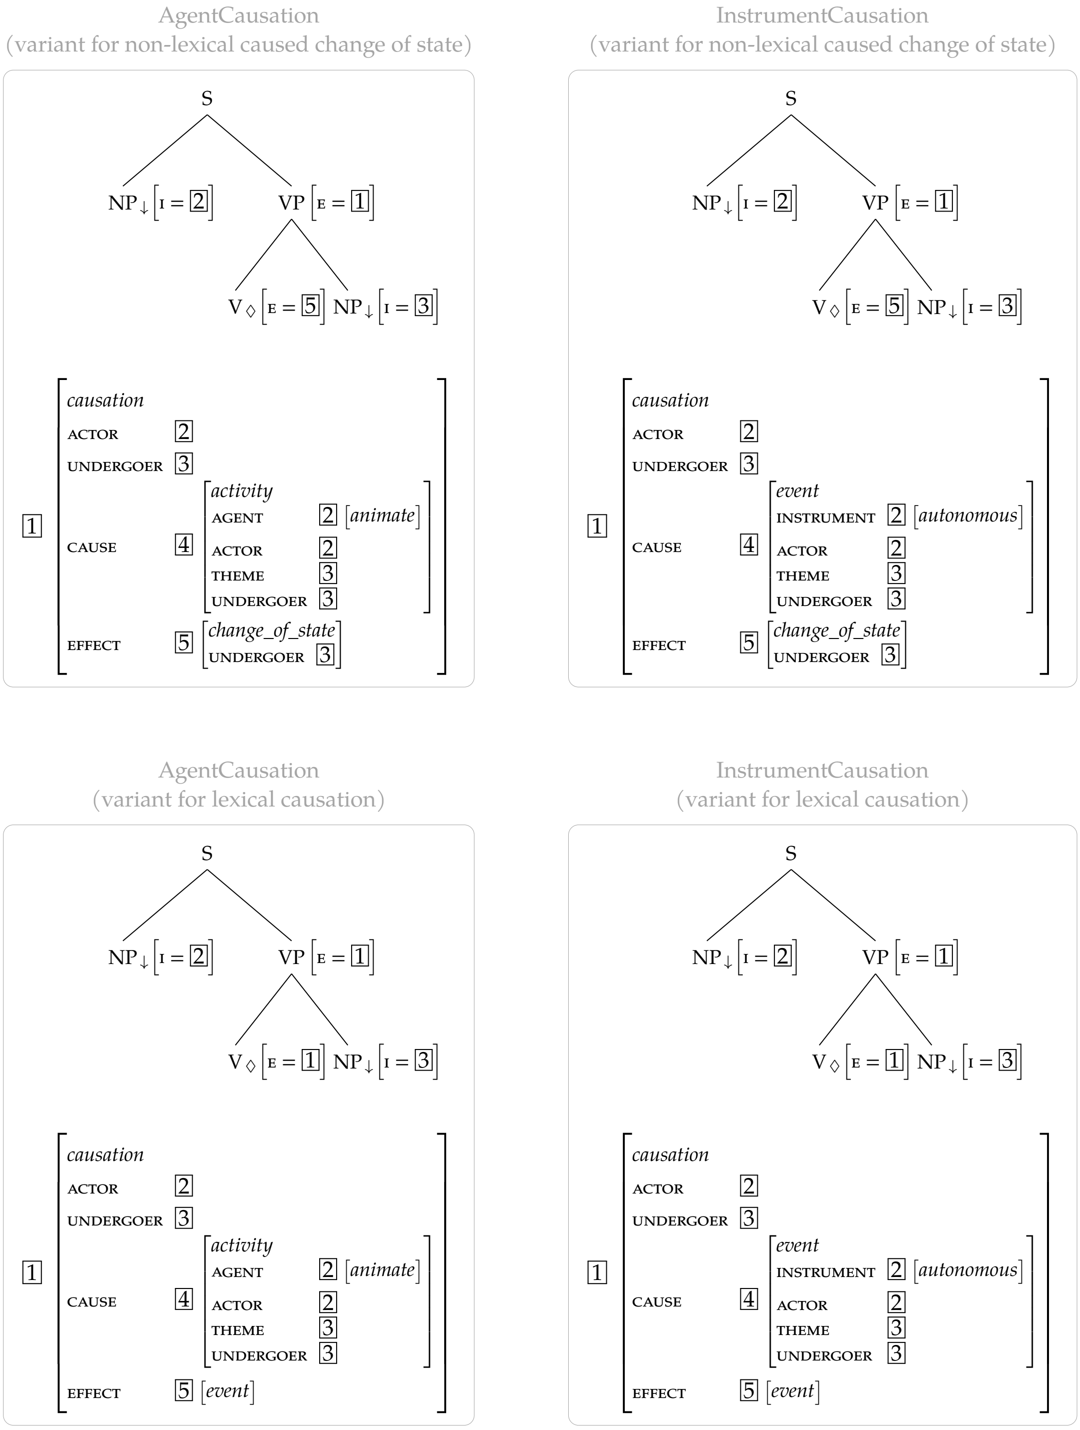


Figure 61: Trees and frames for the *AgentCausation* class (left) and the *InstrumentCausation* class (right) in the CIA/ISA metagrammar.

For change of state verbs like *empty*, the constructions incorporate the verb's lexical contribution in the EFFECT subframe (nodes and frames tagged with $\boxed{5}$ in the two top constructions in Figure 61). For lexically causative verbs, like *slice*, the verb's semantic contribution is instead unified with the entire *causation* frame of the construction (nodes and frames tagged with $\boxed{1}$ in the two bottom constructions in Figure 61).

Transitive uses of *change_of_state* verbs can instantiate the non-lexically causative variant of the *AgentCausation* class if the entity in the subject position is compatible with the type *animate*. No further requirements are imposed on the types of the verb or the subject argument. The construction covers sentences like (50c) *Queequeg empties the bucket*. Constraints on the type of the object argument – here, the UNDERGOER – are already contained in the lexical type hierarchy.

If the entity in the subject position of a *change_of_state* verb in a transitive environment is compatible with the type *autonomous*, the sentence may instantiate the non-lexically causative variant of the *InstrumentCausation* class. In these cases, the subject is identified with the INSTRUMENT role in the CAUSE subframe. This argument is then the most agent-like argument and is thus assigned the ACTOR semantic macrorole. The frame type hierarchy includes constraints on the nature of instruments for different events. For instance, an INSTRUMENT for a causative event whose EFFECT is of the type *becoming_broken* must be compatible with the type *physical_object*. The *autonomous* requirement in the *InstrumentCausation* class ensures that sentence (50i) *The pump empties the bucket* is accepted by the parser, while the similar sentence (50e) **The spoon empties the bucket* is rejected due to the non-autonomous instrument *spoon*.

In the previous chapter, the class named *LexicalCausationConstruction* provided a construction for lexically-causative verbs like *insert*. That chapter did not distinguish between agent subjects and instrument subjects. The distinctions made in the present metagrammar add detail to the representation of constructions involving lexically causative verbs. A class named *LexicalCausation* in the present metagrammar functions as a helper class whose partial descriptions are inherited by other constructions. When the two metagrammars are merged into one combined model, the *LexicalCausationConstruction* class from the previous chapter will be removed in favor of the more detailed version presented here, where the *LexicalCausation* class is inherited by the variants of *AgentCausation* and *InstrumentCausation* which are responsible for lexically causative verbs.

### 7.5.3 *The instrument adjunct construction*

Instruments can also be added to sentences via a prepositional adjunct phrase headed by *with*. The metagrammar class responsible for such phrases is called *InstrumentWithConstruction*. This class has two alternative realizations that can adjoin to different types of sentences. The descriptions of the two variants are shown in Figure 62 on page 163.

The first realization of this construction can adjoin to a VP node of a tree whose frame is compatible with the type *activity*. This variant of the construction allows the parser to derive the correct tree and frame for sentences like (49d) *Queequeg eats the soup with the spoon*. The frame of the construction unifies with the complete frame of the VP node at which it adjoins, adding an INSTRUMENT attribute and a value for it to the frame representing the sentence.

The syntactic structure of the second variant of this construction is nearly identical to the structure of the first variant, with one difference: the values for the interface feature E assigned to the PP node and the tree's root VP node are identical in the non-causative variant, but they differ in the causative variant. This is necessary because when the construction adjoins to a sentence instantiating a causative construction, the INSTRUMENT will not be added to the frame contributed by the verb. Instead, the INSTRU-

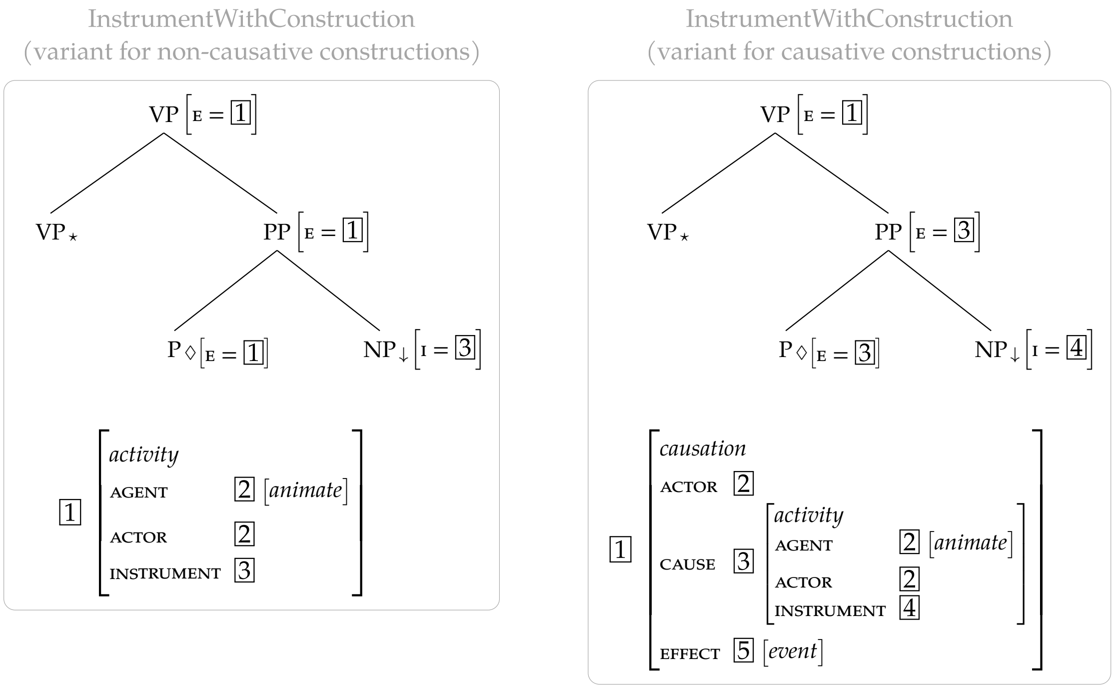


Figure 62: Trees and frames for the non-causative variant of the *InstrumentWithConstruction* class (left) and the causative variant of the same class (right) in the CIA/ISA metagrammar.

$$\textit{with}\quad \boxed{1}\begin{bmatrix} \textit{activity} & \\ \textsc{instrument} & \boxed{2} \end{bmatrix}$$

Figure 63: Lexical frame for the preposition *with* in the CIA/ISA metagrammar.

MENT will appear in the CAUSE subframe, and the frame contributed by the verb will be unified either with the EFFECT subframe or with the entire *causation* frame. This variant of the construction allows the parser to derive the correct tree and frame for sentences like (51d) *Queequeg slices the coconut with the machete*.

At this point, no distinction between lexically causative and lexically non-causative verbs is necessary. Since the causative variant of the *InstrumentWithConstruction* class can only ever adjoin to causative constructions, it will always interact with a frame of the type *causation*, which may or may not be contributed by the verb in addition to the construction.

The metagrammar presented in the previous chapter was concerned with directional prepositional phrases, and the lexicon entries for all prepositions were specified with frames of the type *translocation*, where the TRAJECTORY of the PATH attribute in the *translocation* event was contributed by the individual preposition. In the current metagrammar, the only preposition that appears in the lexicon is *with*. The lexical frame for *with* is shown in Figure 63. It describes an event of the type *activity* which includes an INSTRUMENT attribute. The identity of the INSTRUMENT is determined during parsing based on unification with semantic contributions from constructions such as the ones shown in Figure 62. Note that although the full structure of the prepositional phrase is shown in each of the variants presented in Figure 62, the PP is not actually described redundantly in the metagrammar. Its structure, comprising the PP, P and NP nodes and their dominance and precedence relations, as well as a partial semantic description, is inherited from an unanchored helper class called *InstrumentAdjunct* (see overview in Figure 59 on page 158).

As discussed on page 149, issues of preposition attachment ambiguity are not taken into account here because they are outside the scope of this work. In everyday use, *with*-phrases can attach not only to verb phrases, but also, for instance, to nouns denoting objects or events, as in *the removal of the screw with a screwdriver*. Because nominalizations are not in focus in this thesis, the *InstrumentWithConstruction* class is defined specifically for the context of VPs. In order to extend the metagrammar to allow *with*-phrases to attach to NPs, alternative realizations of *InstrumentWithConstruction* could be defined. The structure of these additional descriptions would closely resemble the ones given in Figure 62, but their root node and foot node would belong to the category NP instead of VP. The semantic description of the alternative forms would be identical to that of the ones given in the figure, because a noun like *removal* would express a *causation* event, analogously to the corresponding verb *remove*. Concerning the metagrammar class inheritance hierarchy, the NP-adjoining variants and the VP-adjoining variants of *InstrumentWithConstruction* could each be characterized in terms of a description of their root and foot node, in addition to an inheritance link to a more abstract class in which the rest of the relevant syntactic structure as well as the entire semantic description would be encoded. Elements that are shared among all these variants, such as the syntactic structure and frame that are specific to the PP itself, are inherited from *InstrumentAdjunct*, so that there is no redundancy with respect to the description of the PP.

## 7.6 CONSTRUCTIONALLY-MOTIVATED ADDITIONS TO THE FRAME TYPE HIERARCHY

Based on the construction classes discussed in the previous section, one addition to the type hierarchy of this metagrammar is necessary to allow the grammar to correctly predict which verbs can instantiate which of the modeled constructions. The *InstrumentCausation* class requires the entity filling the INSTRUMENT role to be compatible with the type *autonomous*. This new type *autonomous* will be added as a subtype of *inedible* (since edible or animate entities will not be regarded as possible instruments). Objects like *machete*, *hammer* or *pump* fall into this category, while objects like *bucket*, *window* or *spoon* are incompatible with *autonomous*.

The extended frame type hierarchy for this metagrammar is shown in Figure 64.

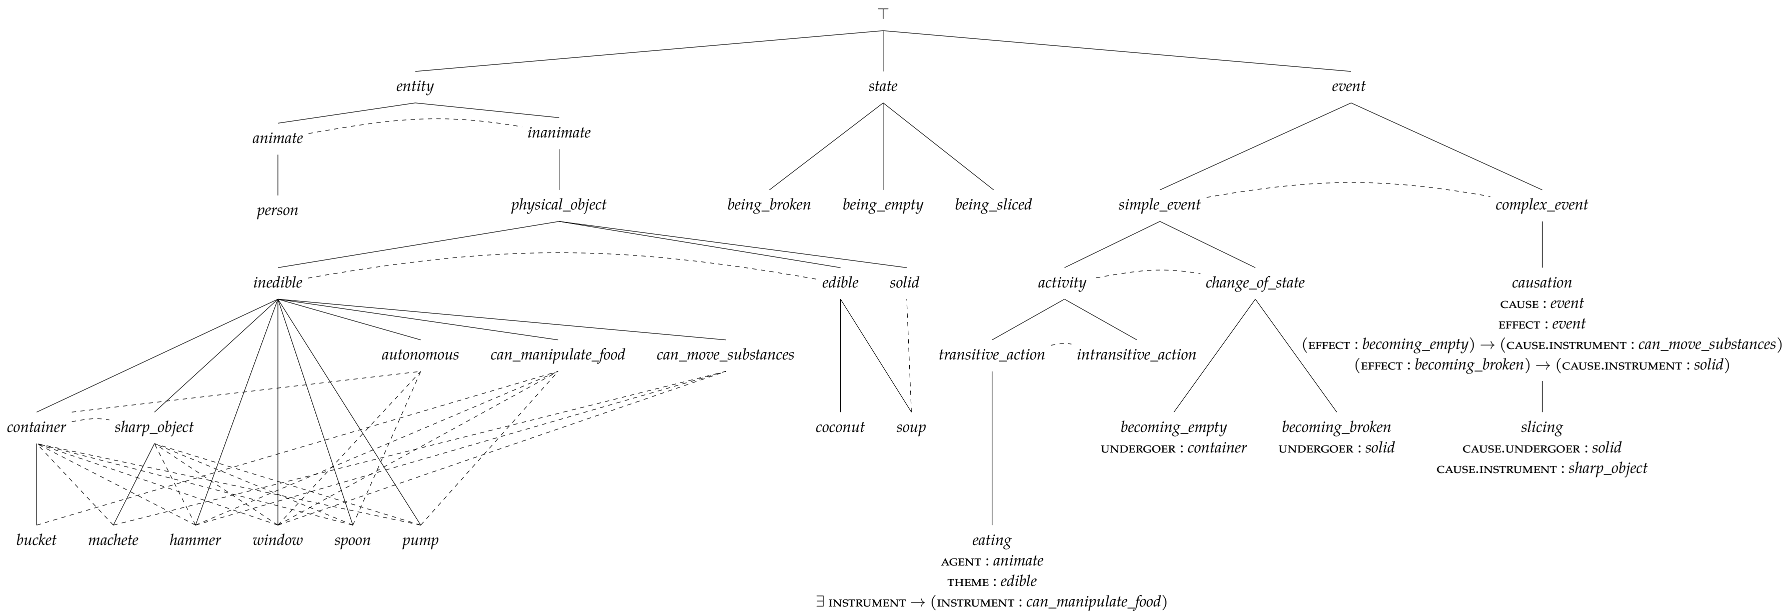


Figure 64: Extended event frame type hierarchy denoting subtype relations and type incompatibilities between the types in the metagrammar for CIA/ISA. Solid lines indicate subtype relations. Dashed lines indicate explicit incompatibility between types.

An anonymous reviewer pointed out that even if we do not typically expect instruments to be edible, it is in fact sometimes possible to form an instrument out of edible material, like a waffle in a specific shape that functions as a jam cup. However, the contexts in which instruments can be crafted out of edible materials and at the same time completely fulfill their purpose as instruments are fairly limited. A spoon that can be eaten is a bit unusual, but a functional lawnmower or power drill that can be eaten are extremely hard to imagine. The type hierarchy presented here does not take into account edge cases like edible jam cups, but in order to include them, the types currently located under *inedible* could directly inherit from *physical_object* instead, and additional incompatibility constraints would need to be added to express the fact that a waffle jam cup would not be an appropriate instrument for actions like *slicing*.

Another extension that could benefit the representation of such edge cases would be one that assigns probabilities to type and attribute constraints, following the proposal for stochastic frames presented by Schuster et al. (2020). The type constraint requiring all instruments to be inedible could then instead be formulated as a high probability for instruments being inedible, leaving room for waffle jam cups and the like. Similarly, probabilities could be introduced to express that while a spoon or a machete are expected to be able to manipulate food, it is also not impossible (albeit fairly unusual) for a hammer or a bucket to be used in conjunction with food. XMG-2 does not currently support probabilities. The actionality scale proposed by Van Hooste (2018: 156) could also be represented with probabilities, such that objects that are regarded as "more autonomous" would be encoded to be more likely instrument subjects, and objects that are "less autonomous", like spoons, would be unlikely but not impossible in that position.

Recall that the type hierarchies presented here are not meant to be exhaustive (see discussion on page 116). Since the metagrammars cover a specific set of constructions, the design of the type hierarchies mainly follows the descriptions of these constructions in the linguistic literature. The modular nature of XMG and the description of the type hierarchy in the form of a set of constraints means that it is always possible to make extensions as needed for special cases.

## 7.7 DERIVATIONS FOR EXAMPLE SENTENCES

This section presents the derivations licensed by the metagrammar for the different sentence types discussed in Section 7.2. Sentences whose derivations are not shown here can be parsed by loading the grammar code provided in the thesis repository into the TuLiPA parser.

The sentences whose derivations will be discussed in the following instantiate various combinations of the constructions covered in this metagrammar.

First, sentence (49d) *Queequeg eats the soup with the spoon* will be parsed, instantiating the *TransitiveActivityConstruction* and the *InstrumentWithConstruction*.

Second, sentence (52b) *The window breaks* will be parsed, instantiating the *Inchoation-Construction*.

Third, sentence (50c) *Queequeg empties the bucket* will be parsed, instantiating the *AgentCausation* class.

Finally, sentence (51e) *The machete slices the coconut* will be parsed, instantiating the *InstrumentCausation* class.

### 7.7.1 *Derivations for sentences instantiating the basic transitive construction with an instrument*

Verbs like *eat* can appear in intransitive or transitive environments, but the lexical frame for *eat* contains an obligatory UNDERGOER that can be unexpressed in a sentence. The *IntransitiveActivityConstruction* and *TransitiveActivityConstruction* classes that can be anchored by this verb are associated with basic frames of the type *activity*. The lexical frame contributed by the verb unifies with the entire constructional frame during derivation. The derivation for Sentence (49d) *Queequeg eats the soup with the spoon* is shown in Figure 65 on page 168.

In this sentence, the prepositional phrase instantiates the *InstrumentWithConstruction*, which adjoins to the VP node of the verb-anchored construction. The frame contributed by the *InstrumentWithConstruction* is unified with the entire frame of the *IntransitiveActivityConstruction*.

### 7.7.2 *Derivations for sentences instantiating the inchoation construction*

Sentence (52b) *The window breaks* is intransitive and instantiates the inchoation construction. The derivation for this sentence is shown in Figure 66 on page 169. The frame contributed by the verb is unified with the entire frame contributed by the construction. The lexical frame contributed by the lemma in the subject position is identified with the value of the UNDERGOER role of the *change_of_state* frame.

### 7.7.3 *Derivations for sentences instantiating the agent-causation construction*

When one of the verbs of the type *change_of_state* appears in a transitive environment, it instantiates the non-lexically causative variant of the *AgentCausation* class. The derivation for Sentence (50c) *Queequeg empties the bucket* is shown in Figure 67 on page 170.

In this derivation, the lexical frame contributed by the verb unifies with the EFFECT subframe of the *causation* frame contributed by the construction. The entity in the subject position is identified with the value of the AGENT and ACTOR attributes in the construction's CAUSE subframe.

### 7.7.4 *Derivations for sentences instantiating the instrument-causation construction*

Sentence (51e) *The machete slices the coconut* instantiates the instrument-causation construction. The distinction between instances of this construction and the agent-causation construction is made purely based on the semantic types of entities appearing in the subject position. *The machete* belongs to the type *inanimate*, while *Queequeg* in sentence (51c) *Queequeg slices the coconut* belongs to the type *animate*. The derivation for Sentence (51e) is given in Figure 68 on page 171.

This sentence involves the lexically causative verb *slice*, whose entire semantic contribution is therefore unified with the *causation* frame contributed by the construction. Requirements on the nature of the INSTRUMENT in *slicing* events are not contained in the lexicon entry for the verb, but in the frame type hierarchy in the form of an attribute constraint.

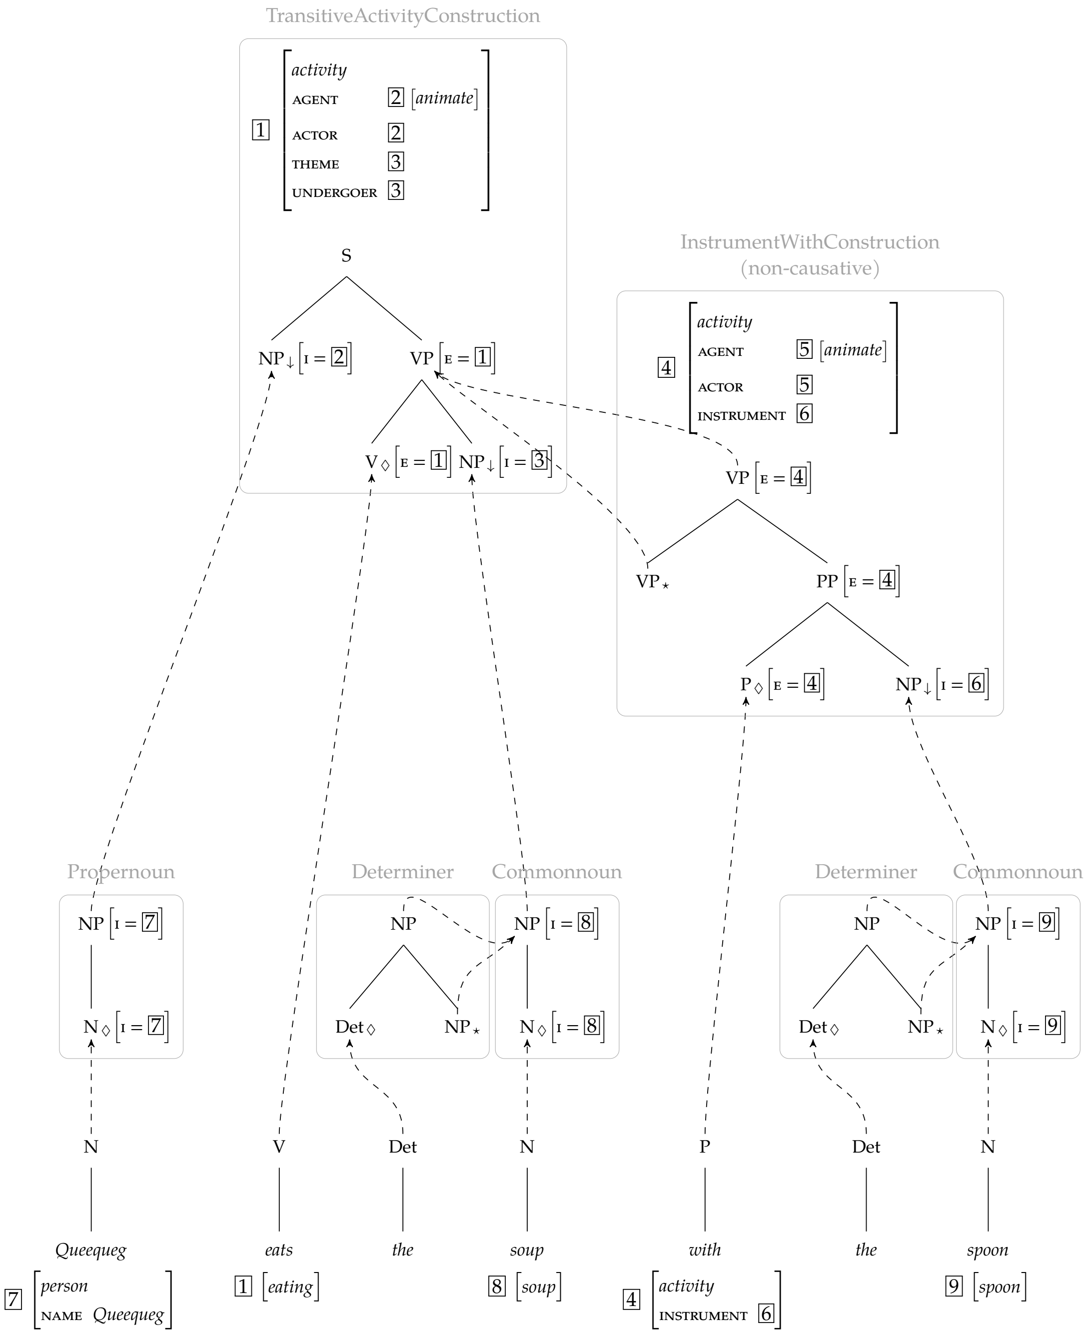


Figure 65: TAG derivation for sentences like (49d) *Queequeg eats the soup with the spoon*.

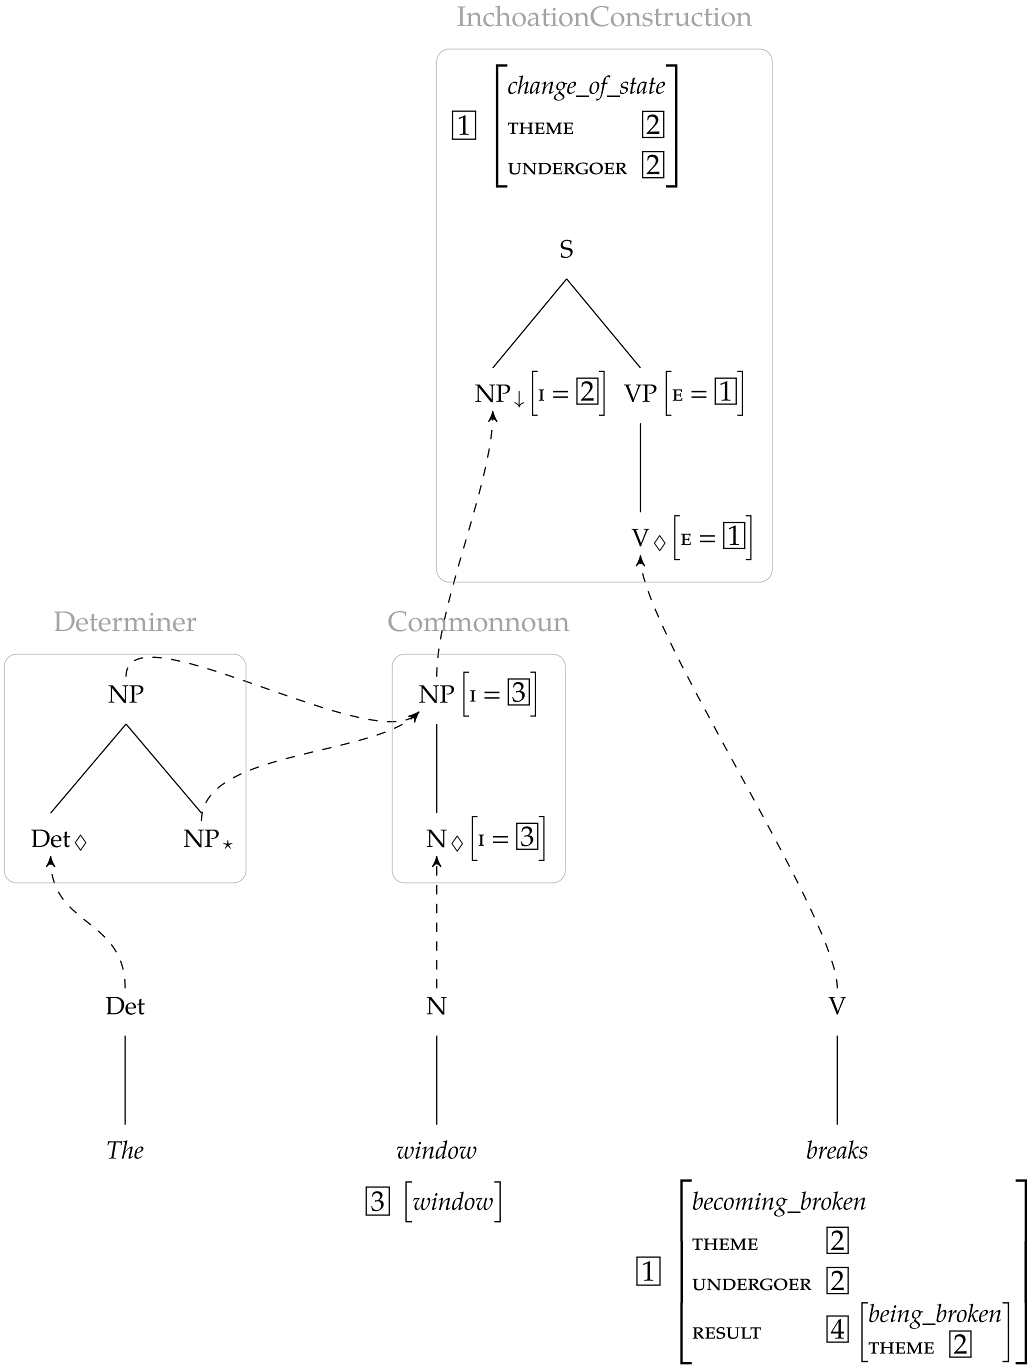


Figure 66: TAG derivation for sentences like (52b) *The window breaks*.

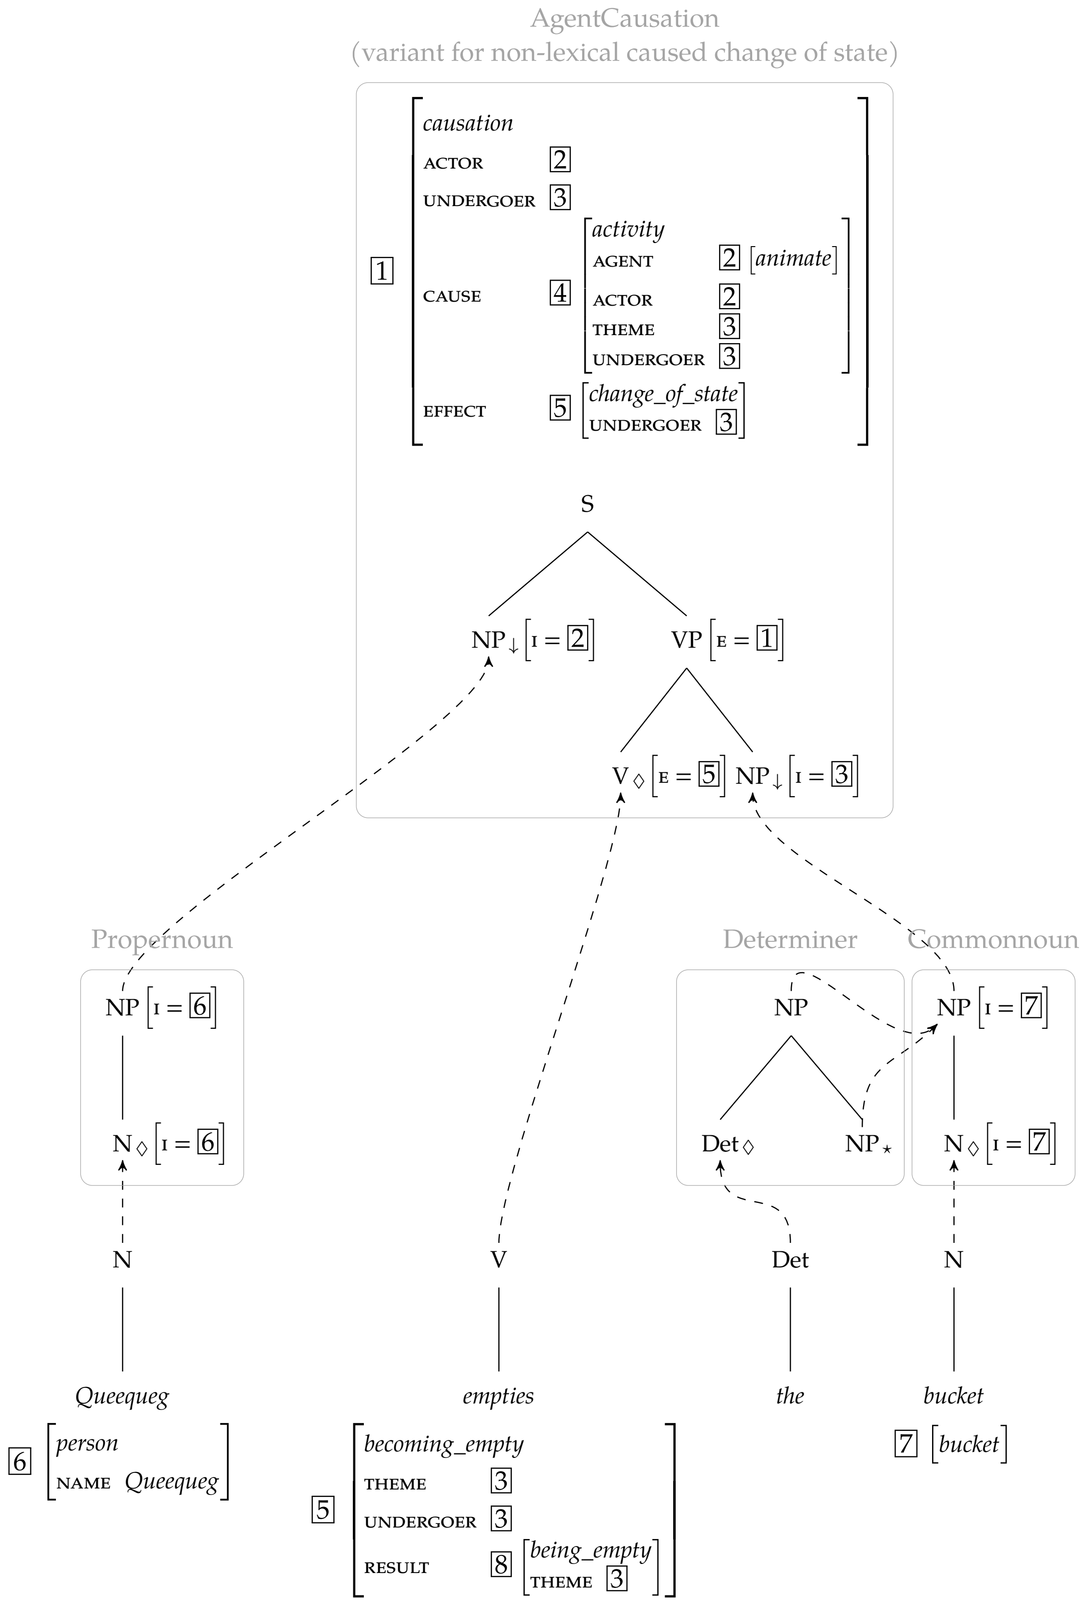


Figure 67: TAG derivation for sentences like (50c) *Queequeg empties the bucket*.

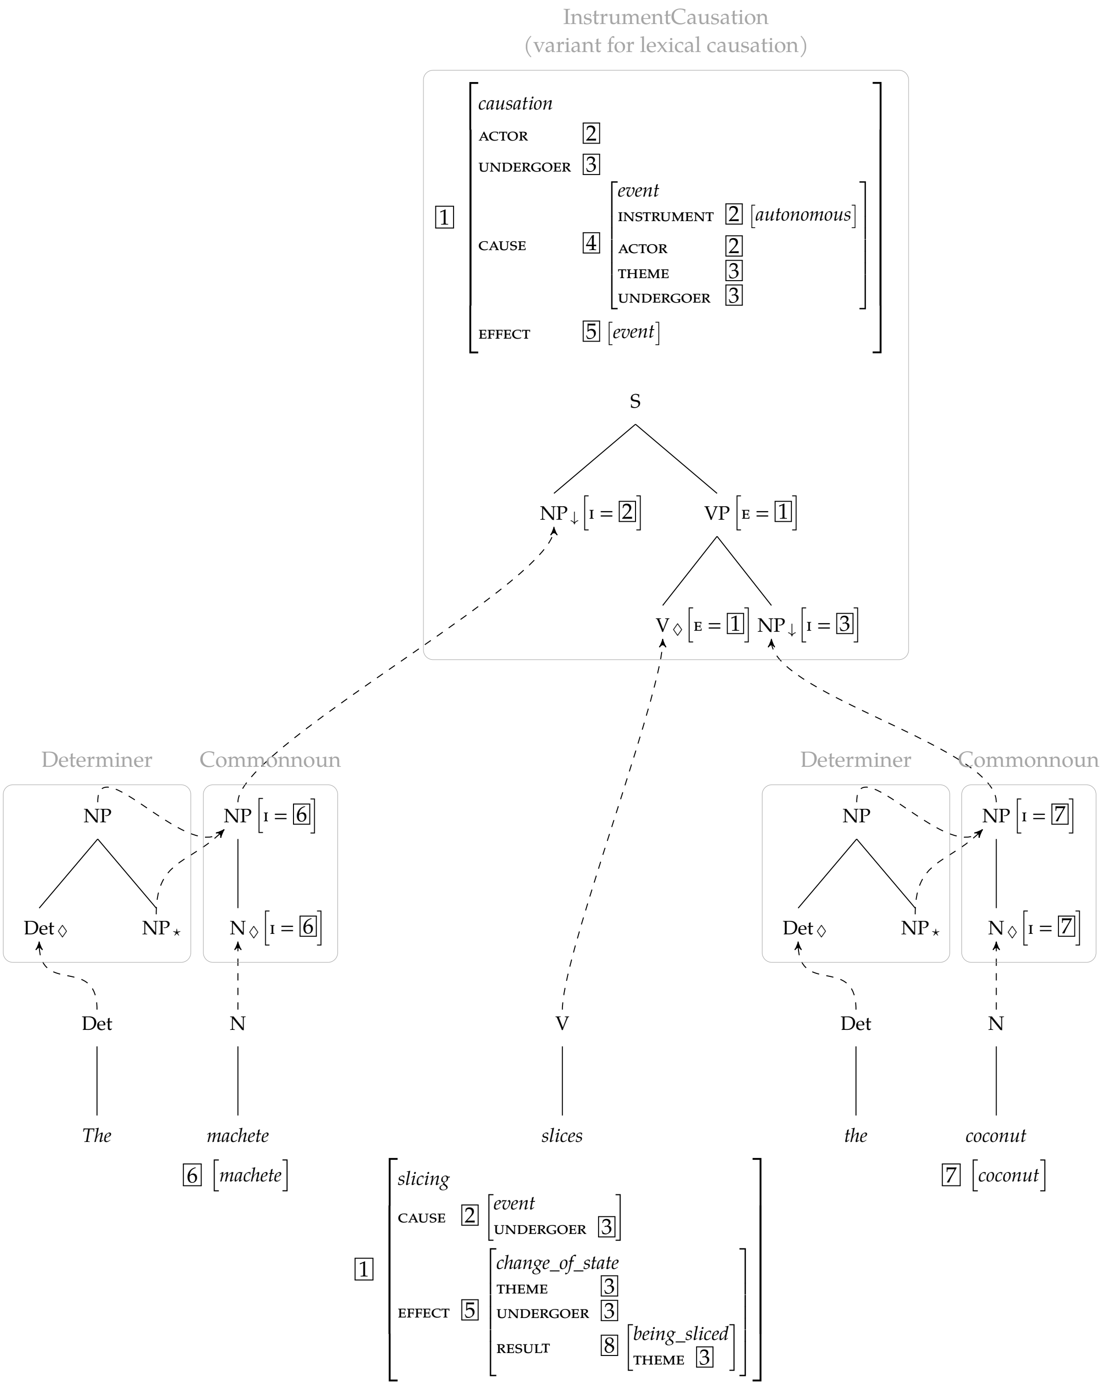


Figure 68: TAG derivation for sentences like (51e) *The machete slices the coconut*.

## 7.8 CONCLUSION

The metagrammar presented in this chapter focuses on modeling the behavior of verbs in the causative-inchoative alternation and verbs in the instrument subject alternation, and contrasts these with verbs that participate in neither or both of these alternations. Each alternation is only available for sentences whose verbs are compatible with the constructions of the alternation. Compatibility is determined based on the type requirements included in the construction descriptions and on type constraints in the frame type hierarchy.

In the metagrammar developed here, the agent causation construction is associated with both the causative-inchoative alternation and the instrument subject alternation. Verbs participating in either of these alternations are modeled correctly by the metagrammar presented here. Verbs like *break* participate in both alternations and can instantiate all constructions associated with either alternation. Attribute constraints for the types that are involved ensure that only valid entities are assigned to each argument slot.

Levin (1993: 4) argues that alternation behavior can be predicted based on certain meaning facets shared by all verbs that participate in a particular alternation. The implementation presented in the previous chapter involved the meaning facet *agentive_manner_of_motion*, which was defined as a type in the frame type hierarchy to allow or disallow the instantiation of the induced action construction by certain verbs. In the present chapter, the meaning facets controlling the availability of alternation-specific constructions are not restricted to the event verbs themselves, but also involve the types of their observed arguments. For instance, if the INSTRUMENT in a causative sentence with an EFFECT of the type *becoming_empty* is compatible with the type (or meaning facet) *autonomous*, then this combination of verb and argument can feature in the instrument subject construction. For instruments that are not compatible with *autonomous*, that construction is unavailable.

The *autonomous* type is a rough approximation of semantic requirements regarding instrument subjects. Van Hooste (2018: 156) argues that each predicate requires its instrument subjects to be located above a certain threshold in a two-dimensional actionality scale that ranks entities based on the two axes autonomy and animacy. Spelling out these verb-specific requirements using type constraints in XMG is possible, but was not pursued for this chapter.

Further semantic requirements on entities in the different role slots for event verbs were also implemented via type constraints. The affordances required for all instruments in events of a certain type were expressed with semantic types like *can_move_substances* or *can_manipulate_food*.

The meaning facets identified here and in the previous chapter as essential properties controlling the availability of various alternations are analogous to the proposal by Levin. She identifies the (in)compatibility of verbs with the meaning facets *motion*, *contact* and *change of state* as relevant indicators for the conative alternation, the body-part possessor ascension alternation, and the middle alternation. The strategy proposed for modeling alternations in this chapter can easily be transferred over to those and other meaning facets (or types) and alternations.

The two alternations modeled in this metagrammar both allow verbs to link their syntactic subject position to different semantic roles. For the causative-inchoative alternation, the subject can be associated with the ACTOR role or the UNDERGOER role. For the instrument subject alternation, the ACTOR is expressed in the syntactic subject position and can be associated with the AGENT role or the INSTRUMENT role.

As in the previous chapter, the example sentences discussed here also involve a type of verb that is lexically causative. Verbs like *slice* have a lexical semantic frame of the type *causation*, and can therefore directly unify with causative frames contributed by constructions. Causative verbs are unable to instantiate non-causative constructions. Non-causative verbs can appear in causative constructions under certain circumstances, and the semantic contribution from the verb is in those cases unified with a subframe of the causative frame contributed by the construction.

The metagrammar presented here also includes verbs that can appear transitively or intransitively, but do not express a change of state and do not participate in the causative-inchoative alternation or the instrument subject alternation. For these verbs, the presence or absence of a direct object or of a prepositional phrase denoting an INSTRUMENT simply results in the inclusion or omission of the corresponding role in the semantic frame for the sentence. Depending on the lexical entry of the verb, roles may be included in the lexical frame even when they are not expressed in a sentence, such as the implied UNDERGOER of intransitive *eating* events or the implied INSTRUMENT in *slicing* events.

So far, the metagrammars developed in the previous chapter and the current chapter were presented separately, since each of them is concerned with contrasting a particular set of constructions that share some, but not all, of their syntactic and semantic properties. The modular fashion in which the metagrammars have been designed makes it simple to merge them into one larger grammar covering all constructions discussed in this thesis part. Chapter 8 will present the combined metagrammar.

# 8 MERGING THE METAGRAMMARS

## 8.1 INTRODUCTION

This chapter is concerned with merging the two metagrammars developed in Chapters 6 and 7. The resulting metagrammar is still limited to a small number of elementary classes, compared to large-coverage TAG grammars like XTAG (Abeillé et al. 1990). However, the ease with which the two metagrammars developed in this thesis can be combined indicates that more constructions and alternations could effortlessly be added at any time, due to the modular architecture of XMG and the sharing of structural information made possible by class inheritance. The results presented here can thus be viewed as a prototype or proof of concept for the architecture of larger TAG grammars modeling the syntax and semantics of verb alternations.

## 8.2 MERGING LEXICAL DESCRIPTIONS

The lexical descriptions in the `<morpho>`, `<lemma>` and `<frame>` dimensions can simply be unified into a combined lexicon comprising descriptions in these dimensions sourced from the two original metagrammars. Since lexicon entries exist independently of each other, no conflicts can arise at this stage, and no instances of redundancy are to be expected. For lexical items that are defined in both separate metagrammars, the sets of trees they can anchor are merged.

Additional lexicon entries can be added as desired. To add lexicon entries referring to entities whose types are not yet reflected in the frame type hierarchy, extensions to the hierarchy will be required. The addition of such entries will be sketched in Section 8.5.

## 8.3 MERGING CONSTRUCTIONS

The `syn_dimension.mg` file of the combined metagrammar contains all anchored elementary constructions appearing in either one of the original metagrammars, as well as all abstract classes from which the anchored constructions inherit partial descriptions in various dimensions. As mentioned in Chapter 7, the classes at the first through third levels of the hierarchy are already completely shared between the two metagrammars.

Chapter 7 made a distinction between caused events whose causing entity, the ACTOR, is an AGENT, and caused events whose causing entity is an INSTRUMENT. This distinction replaces the less detailed *LexicalCausationConstruction* class developed in Chapter 6. In the combined metagrammar, agents as well as instruments can be the subjects in events involving either lexically causative verbs or lexically non-causative verbs.

The remaining constructions that are specific to one of the two metagrammars can be merged into a combined grammar without conflicts. The resulting class inheritance hierarchy for the combined metagrammar is shown in Figures 69 and 70 on pages 176 and 177. For the sake of readability, Figure 69 only includes initial constructions, while Figure 70 only includes auxiliary constructions.

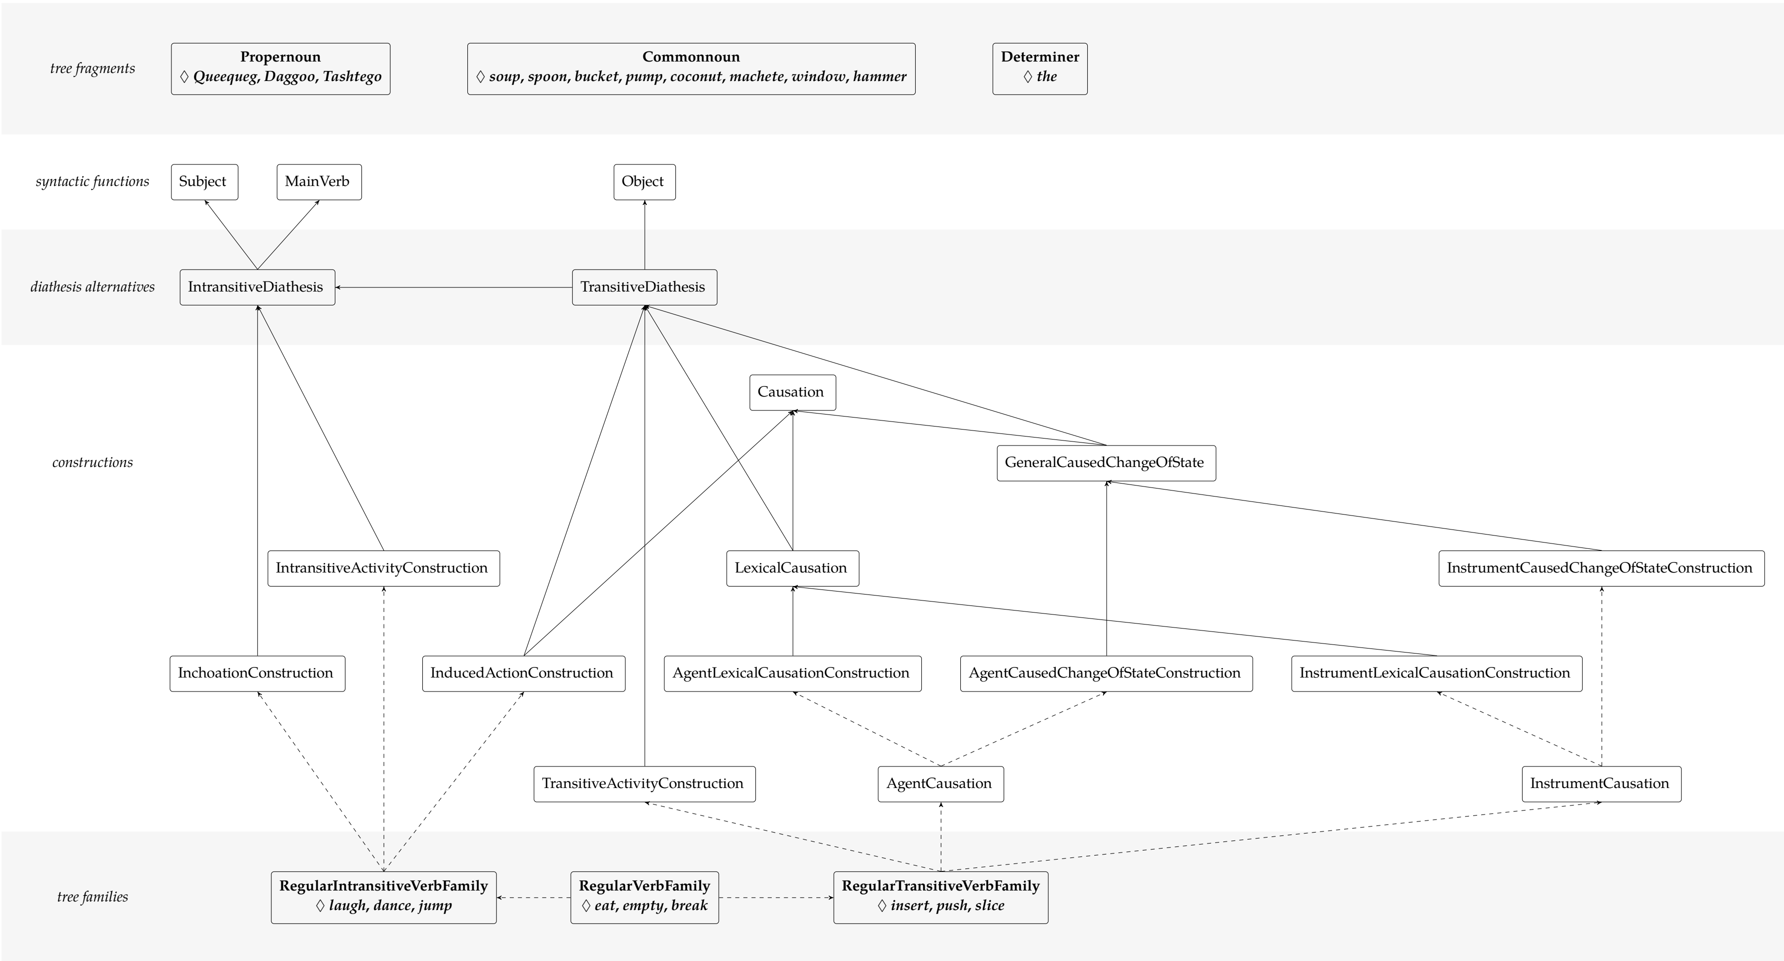


Figure 69: Class inheritance hierarchy of all initial constructions in the combined metagrammar covering CMC, IAA, CIA and ISA. Solid arrows denote structure inheritance. Dashed arrows denote disjunction.

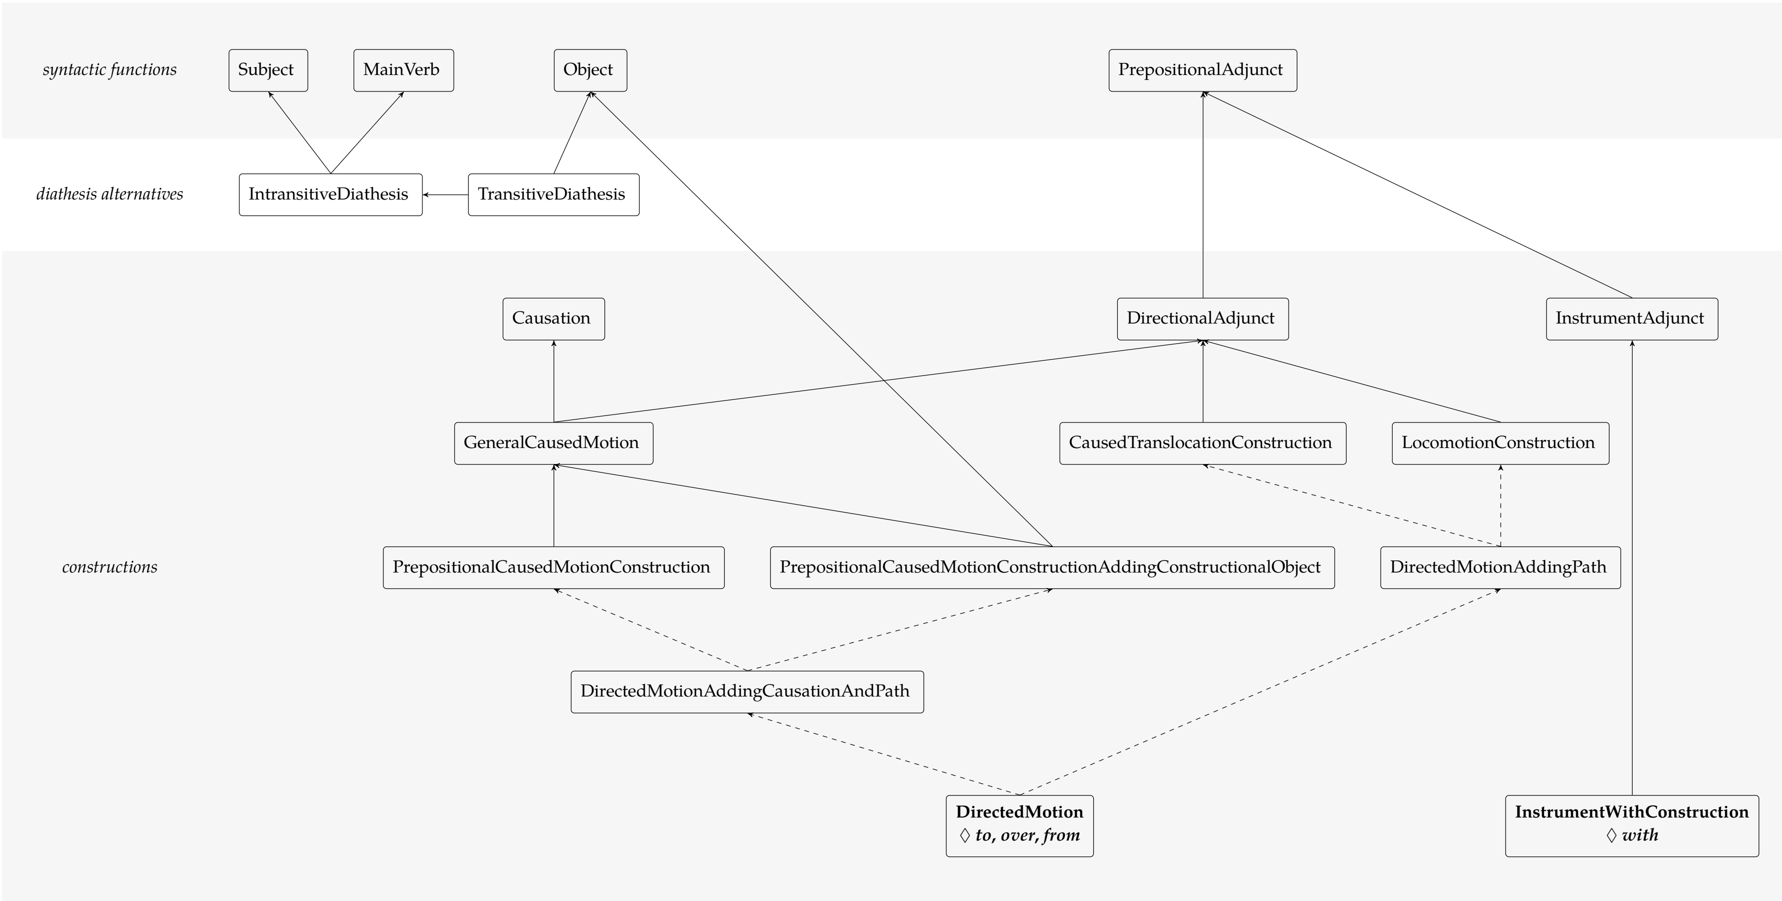


Figure 70: Class inheritance hierarchy of all auxiliary constructions in the combined metagrammar covering CMC, IAA, CIA and ISA. Solid arrows denote structure inheritance. Dashed arrows denote disjunction.

Additions to the set of construction descriptions are simple to implement. For instance, specific causative constructions can be added that inherit from the classes *Causation*, *GeneralCausedMotion* or *GeneralCausedChangeOfState*; the partial descriptions contributed by those classes can then be complemented with more descriptions as required for the particular construction that is being added. Section 8.5 sketches the addition of another verb alternation.

## 8.4 MERGING TYPE HIERARCHIES

Finally, the frame type hierarchies from both metagrammars are merged to form one combined hierarchy. The general structure of the separate type hierarchies is comparable. The combined type hierarchy is shown in Figures 71 and 72 on pages 179 and 180. In the interest of readability, Figure 71 shows the subtypes of *entity* and *state* and hides the subtypes of *event* and *path*, while Figure 72 shows the subtypes of *event* and hides the other subtypes.

In the combined metagrammar, the caused-motion construction is available to all verbs whose lexical frame is of the type *activity*, which includes verbs like *laugh*, or the type *causation* with an *activity*-type CAUSE subframe and an EFFECT subframe that is compatible with *translocation*, which includes verbs like *insert*. If the arguments of a verb are not compatible with the semantic requirements from the caused-motion construction or with the verb's lexical semantic frame, the construction cannot be anchored by that verb with those arguments. The lexically-causative verb *slice* cannot instantiate the caused-motion construction because its EFFECT subframe has the type *change_of_state*, which is incompatible with *translocation*.

The induced action construction is available to verbs that are compatible with the *agentive_manner_of_motion* type, such as *jump* or *dance*.

The causative-inchoative alternation allows verbs to instantiate either the inchoation construction or a construction expressing a caused change of state. This is possible exclusively for verbs whose semantic type is *change_of_state*, such as *empty* or *break*.

Finally, the instrument subject alternation allows verbs to express either their AGENT or their INSTRUMENT in the subject position. The entity in the subject position in both relevant constructions is associated with the ACTOR macrorole. Both constructions involved in this alternation are characterized by a semantic frame of the type *causation*. Verbs can only instantiate these constructions if they are either lexically causative, such as *slice* or *insert*, or if their lexical type is *change_of_state*, as in the case of *empty* or *break*.

The metagrammar modeling the causative-inchoative alternation and the instrument subject alternation makes distinctions at the level of physical objects that do not appear in the CMC/IAA metagrammar: because events of the type *eating* require edible UNDERGOERS, the type for physical objects in the CIA/ISA metagrammar has the subtypes *inedible* and *edible* to allow this distinction. Additionally, a distinction between *solid* entities and entities that are incompatible with that type was introduced. The combined type hierarchy reflects the more detailed structure from the CIA/ISA metagrammar.

As mentioned above, in order to extend the metagrammar to cover more constructions and lexical items, additions to the type hierarchy may be required. The following section details how such extensions can be implemented with little effort.

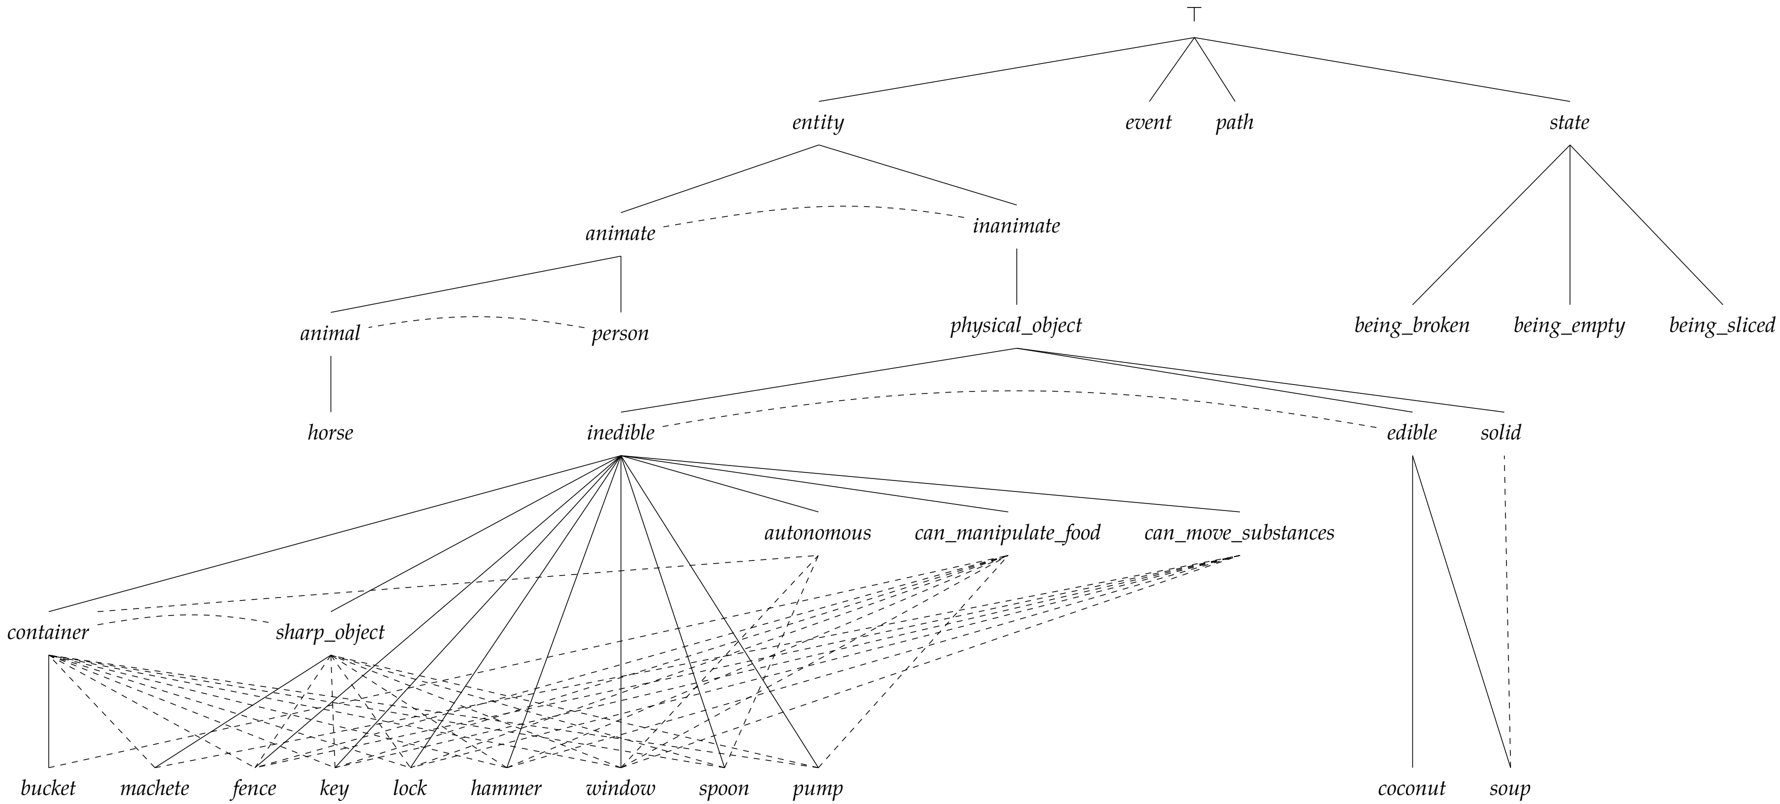


Figure 71: Frame type hierarchy of the combined metagrammar covering CMC, IAA, CIA and ISA, showing subtypes of *entity* and *state*.

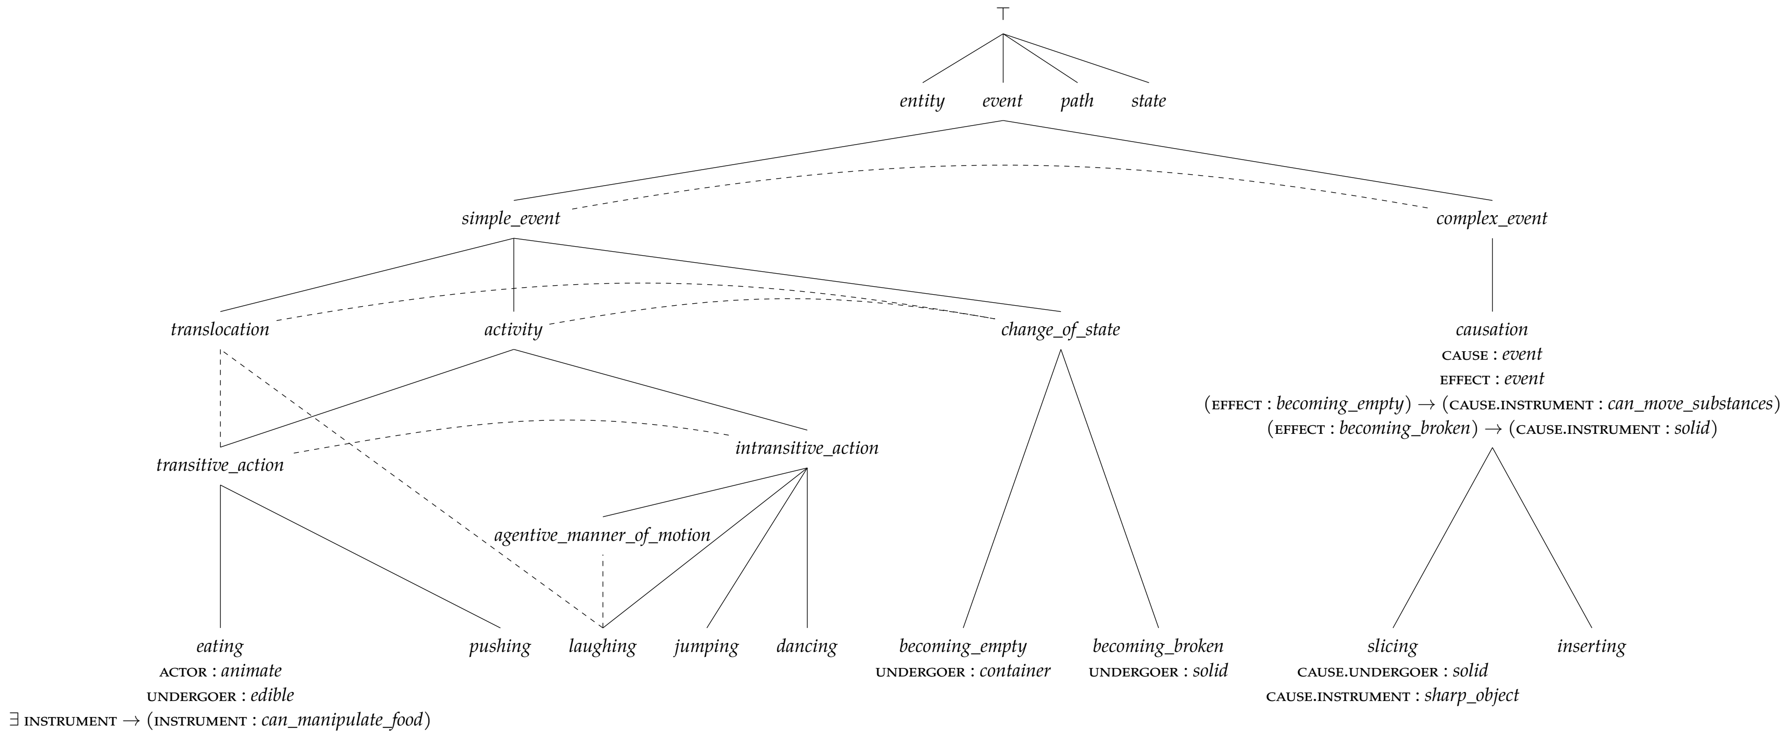


Figure 72: Frame type hierarchy of the combined metagrammar covering CMC, IAA, CIA and ISA, showing subtypes of *event*.

## 8.5 further extensions to the metagrammar

The metagrammars presented in Chapters 6 and 7 model three alternations and one fairly productive individual argument structure construction from a constructionist perspective. While the focus in these models is on a selection of English verb alternations, other constructions of the language can easily be added to the grammar to allow for a wider coverage of sentence types. This section will sketch the steps involved in the course of extending the grammar. Subsection 8.5.1 will incorporate another English verb alternation into the metagrammar. Subsection 8.5.2 will describe how morphological issues and phenomena like passivization can be included in the metagrammar.

### 8.5.1 *The conative alternation*

The conative alternation (Levin 1993: 41, Levin & Rappaport Hovav 2013: 54) allows certain transitive verbs to appear in the conative construction. This construction involves a prepositional phrase headed by *at* and denotes an attempted action without entailing that the action itself is actually completed. The conative construction can be added to the current metagrammar using the same mechanisms that have been employed so far.

The two relevant constructions are exemplified in (53) (taken from Levin 1993: 41).

(53) a. Margaret cut the bread.
b. Margaret cut at the bread.

As discussed by Levin (1993: 5–11) and Levin & Rappaport Hovav (2013: 57), the availability of the conative construction for a given verb depends on the compatibility of the verb's semantic type with the meaning facets *motion* and *contact* (see also Section 2.2 of this thesis). This circumstance can be leveraged when it comes to constraining with which verbs the conative construction can be instantiated.

To allow the metagrammar to model sentences involving the constructions involved in the conative alternation, the relevant lexicon entries must first be added. The lemmas *Margaret*, *cut*, *at* and *bread* are added to the `lemma.mg` file, and their inflected forms are added to the `morph.mg` file. The semantic frame for *cut* looks similar to the frame that has been defined for the verb *slice*.

Words like *Margaret* and *bread* will be specified to anchor the existing *Propernoun* and *Commonnoun* classes, respectively. The verb *cut* will be specified to anchor the *RegularTransitiveVerbFamily* class, analogous to the verb *slice* which is already covered in the metagrammar.

The preposition *at* will be specified to anchor the yet-to-be-defined *ConativeConstruction* class, whose syntactic structure is inherited from the existing *PrepositionalAdjunct* class.

A frame describing the meaning contributed by the conative construction is shown in Figure 73 on page 182. The frame is of the type *starting* because the construction merely expresses an attempted action, which is not necessarily completed. The value of the THEME attribute of the *starting* event corresponds to the completed event that would be described by the non-conative use of the verb.

In the metagrammar presented in Chapter 6, certain preposition-anchored auxiliary constructions involve semantic frames that embed the semantic contribution of the sen-

$$\boxed{1}\begin{bmatrix} \textit{event, starting} & \\ \text{AGENT} & \boxed{2} \\ \text{ACTOR} & \boxed{2} \\ \text{THEME} & \boxed{3}\begin{bmatrix} \textit{event, motion, contact} & \\ \text{AGENT} & \boxed{2} \\ \text{ACTOR} & \boxed{2} \\ \text{THEME} & \boxed{4} \\ \text{UNDERGOER} & \boxed{4} \end{bmatrix} \end{bmatrix}$$

Figure 73: Semantic frame for the *ConativeConstruction* class.

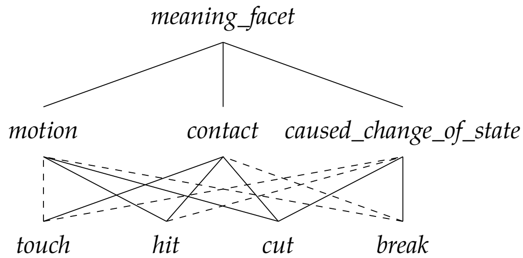


Figure 74: Types, subtype relations and incompatibility constraints reflecting the meaning facets associated with the conative construction. Dashed lines indicate explicit incompatibility between types.

tence they adjoin to into the value of one of their attributes. For instance, the *PrepositionalCausedMotionConstruction* class allows a prepositional phrase like *to the door* to be added to a sentence like *Queequeg pushed Daggoo*, and the *transitive_action* frame for *Queequeg pushed Daggoo* gets embedded into the CAUSE subframe of the *causation* frame provided by the auxiliary construction.

Analogous to this process, the conative construction sketched here can embed a verb's lexical meaning under its THEME attribute.

The final step is to extend the existing type hierarchy with the types required to cover the additional construction. As mentioned above, the availability of the conative construction depends on certain meaning facets. Levin (1993: 5–11) discusses four example verbs, *cut*, *hit*, *touch* and *break*, that have different properties with respect to participation in three alternations including the conative alternation. She associates each verb's alternation behavior with some meaning facets that, according to her, are predictive of participation in the three alternations she is concerned with. A partial frame type hierarchy representing her perspective on the example verbs and meaning facets is shown in Figure 74.

In the argument by Levin, the example verb *cut* is regarded as having a lexical meaning facet *change_of_state*. In this thesis, the closely related verb *slice* has been treated as lexically causative instead. Its EFFECT subframe is of the type *change_of_state* and expresses an UNDERGOER's change to a certain result state. This is in contrast to pure *change_of_state* verbs like *break*.

Levin's example verbs *touch* and *hit* would fall into the *transitive_ action* type in the frame type hierarchy presented in this chapter, while *break* falls into the *change_of_state* type and *cut* falls into the *causation* type. The conative construction can thus be instantiated by *transitive_action* verbs and by *causation* verbs; in the metagrammar, it will not be necessary to explicitly prevent the verb *break* from instantiating it. The description of the construction class can simply require the verb in sentences to which it adjoins to have a type that is compatible with *transitive_action* or *causation*, thereby automatically excluding all pure *change_of_state* verbs.

This leaves the meaning facets *motion* and *contact*. The construction-controlling meaning facets modeled so far in this thesis were all restricted to a common supertype. For instance, the meaning facet *agentive_manner_of_motion* effectively selected for a subset of all *intransitive_action* verbs. In the context of the conative construction, the meaning facets need to be modeled in a way that ensures that event verbs with different, mutually incompatible supertypes can be related to them without conflicts. One possibility to implement this would be to add *motion* and *contact* directly as subtypes of *event*. Alternatively, a sibling type of *event* could be created that covers all meaning facets of the sort discussed here. The complete implementation of the conative construction is outside the scope of this chapter, but several options are available should this metagrammar be extended at a later time. The modular, constraint-based nature of the XMG framework makes it easy to update the type hierarchy if desired.

The *ConativeConstruction* class discussed in this section will of course also receive a syntactic description. More factorization opportunities beyond the ones already covered may be available. The provided sketch of the required additions to the metagrammar shows that the architecture developed here is well-suited to allow for extensions to cover more alternations and constructions.

### 8.5.2 *Syntactic variation and morphology*

The metagrammars presented in this thesis part cover a number of English constructions that are relevant to certain verb alternations. The type hierarchy is the main tool for distinguishing between superficially-similar constructions, based on the types associated with the verbs appearing in the input sentences as well as the types of entities in the different argument positions. The syntactic realizations of the constructions under discussion follow from the XMG class hierarchy as proposed by Crabbé et al. (2013) (see Section 5.3.2 of this thesis). Covering additional syntactic realizations is a matter of extending the XMG class hierarchy to describe more possible configurations. Thanks to the modular nature of XMG and the absence of a requirement for a structural isomorphism between trees and frames (see Kallmeyer & Osswald 2013: 270), implementing additional syntactic realizations would not change the semantic descriptions of the constructions.

While the metagrammars successfully model the differences between the various constructions, certain simplifications have been made here for the sake of space and readability. For instance, the argument order for the constructions has been explicitly fixed, with the subject always appearing as the first argument and all verbs having active forms. This is not necessary in the XMG framework; classes can be underspecified with respect to the relative position of nodes. Of course, not constraining the position of nodes at all would lead to overgeneration in the grammar, and a class defined like that would yield numerous linguistically impossible realizations when the metagrammar is compiled.

Softening the constraints on argument order in the constructions slightly would pave the way for including variations of the constructions under discussion. For instance, passive realizations of certain constructions can easily be added. Figure 15 on page 66 shows how Crabbé et al. (2013) suggests treating realization alternatives like passive and active forms. At the level of tree fragments, different classes exist to describe *ActiveVerbForm* and *PassiveVerbForm*. Classes on the level of diathesis alternatives then describe the realization of intransitive or transitive diatheses by referring to either the class responsible for active verb forms or the class responsible for passive verb forms. In the case of transitive diathesis with an active verb, classes describing the syntactic functions (arguments) *Subject* and *Object* are taken into account. For the passive diathesis, only the *Subject* is necessary. The model described by Crabbé et al. 2013 describes these classes on a purely syntactic level. However, as shown in the metagrammar hierarchies for the work performed in this thesis (see for instance page 119), constructions appear on a separate level. For their syntactic realization, they can refer to descriptions of diathesis alternatives. Their semantic descriptions are part of the construction class constraints.

Within this structure, passives do not pose a problem. Analyzing a sentence like *The horse was jumped over the fence* simply requires the addition of a passive realization of the *TransitiveDiathesis* class. Since the *InducedActionConstruction* imports its syntactic description directly from the *TransitiveDiathesis* class, the passive realization would immediately also be available to *InducedActionConstruction*. For more on describing passive sentences in TAG, see Kroch & Joshi (1985: 60) and Abeillé et al. (1990: 25). Note that the metagrammar approach does not require any assumptions with respect to a transformational process connecting active and passive forms; neither one is underlying or derived, they are both simply described, and facets that they share are described only once and then structurally inherited wherever they are needed.

Whether the sentence is a valid instance of the induced action construction then comes down to the question whether the type of the verb *jump* and the types of the entites described by *horse* and *fence* are permitted in the relevant slots of the construction. Because each of these types is compatible with the construction's type requirements, the induced action reading will be available. In contrast, a sentence like *The horse was pushed over the fence* would not license the induced action reading, because *push* is not compatible with the *agentive_manner_of_motion* type.

Note that an argument described in the construction's description does not necessarily need to be expressed at the surface. The lexicon entry for the verb *eat* contains a THEME that may be unexpressed in case the verb appears in a sentence without an object. Analogously, the definition of the *InducedActionConstruction* (see page 125) contains an ACTOR who may remain unexpressed when the construction is realized as a passive sentence without a syntactic object, such that the THEME is unified with the syntactic subject position.

The linguistic literature referenced throughout this thesis tends to predominantly focus on certain canonical forms of the constructions, but as Müller & Wechsler (2014a: 32) note, variants like nominalizations of resultatives may pose additional issues. They are concerned with cases like the ones shown in (54) and argue that the only way to analyze them would be to assume that lexical rules are required to license the resultative constructions, because morphological processes need to take place after processes that introduce arguments.

(54) a. Er tanzt die Schuhe blutig / in Stücke.
he dances the shoes bloody into pieces

b. die in Stücke / blutig getanzten Schuhe
the into pieces bloody danced shoes

c. * die getanzten Schuhe
the danced shoes

A metagrammar with LTAG and frames from a more morphology-oriented perspective is presented by Zinova (2021). She presents an analysis of a number of Russian verbal prefixes, with a focus on semantic composition that is triggered by operations at the morphological and syntactic levels.

As shown by Zinova (2021), in a metagrammar with LTAG and frames there is no need to assume any ordering of processes on the different levels. Instead, constraints in the metagrammar encode which affixes are licensed in a given environment and which ones are blocked. Other constraints encode the number and type of verbal arguments. Well-formed expressions in the language described by the grammar must fulfill all constraints on all levels. As Zinova puts it,

> The level of the metagrammar is well-suited for capturing derivational morphology processes: it allows for a general description of derivational patterns that can be accompanied by a change of the argument structure. (Zinova 2021: 229)

Concerning the practical implementation in XMG, nominalizations would also integrate without problems (albeit not as seamlessly as passives) into the proposed metagrammars from the previous chapters of this thesis. The constructions described here always refer to a sentence's main verb, whose type and arguments are taken into account in order to allow or disallow each candidate construction analysis. Nevertheless, the tools provided by XMG would be sufficient to add alternative syntactic configurations, such as nominalizations. For pointers on the syntactic modeling of nominalizations in TAG, see Kroch & Joshi (1985: 51). Again, the semantics associated with the constructions would come into play at a less abstract level in the construction-specific XMG classes, which would receive access to additional syntactic realizations.

The focus of the work presented here is on modeling the mechanisms leading to the availability of different interpretations for sentences exhibiting a similar basic syntactic structure. Including other forms of the constructions, such as the variants illustrated in (55), is beyond the scope of this thesis, not least because the judgments from linguistic literature on which the metagrammar is based also predominantly focus on the realizations of the constructions in their "canonical" forms.

(55) a. The tissue was sneezed off the table. (*passive CMC*)

b. Es wurde ein Taschentuch vom Tisch geniest. (*impersonal passive CMC*)

c. The soldiers were marched around the yard. (*passive IAA*)

d. ? the sneezing of the tissue off the table (*nominalized CMC*)

e. ? the marching of the soldiers around the yard (*nominalized IAA*)

Modeling further cases like the ones shown in (55) would necessitate a thorough analysis of the sentences and a well-motivated assignment of grammaticality judgments and expected interpretations for each construction realization. For instance, should (55e) be interpreted as an induced action, or should the nominalization simply express the action of *marching* performed by *the soldiers*? Making these judgments and decisions is beyond the goals of this thesis. However, as discussed in this section, the tools provided by LTAG also make it possible to model, for instance, passives (Kroch & Joshi 1985: 60, Abeillé et al. 1990: 25), including passives with impersonals (Kroch & Joshi 1985: 63), and nominalizations (Kroch & Joshi 1985: 50), and it would therefore be possible to incorporate these phenomena into the grammar model as well. Kallmeyer et al. (2016) show how passive voice and the macrorole assignments occuring in that context can be implemented with XMG.

## 8.6 CONCLUSION

The combined metagrammar can be compiled in the same way as the original metagrammars, and the parser can derive the desired structures for all example sentences discussed in Chapters 6 and 7. The combined metagrammar files and their compiled versions can be found in the thesis repository.

Although the metagrammars presented in this thesis part follow a constructionist perspective, certain distinctions made by Goldberg (1995) are not applicable here. For instance, she distinguishes between argument roles provided by constructions and participant roles provided by verbs. In the implementation developed here, both constructions and lexical frames for verbs can contribute roles from the same role inventory. For instance, the basic transitive construction for activity verbs provides an ACTOR and an UNDERGOER role, which is congruent with the roles provided by the lexical frames for transitive verbs like *push*. In contrast to this, the lexical frame for *eating* events always contributes an UNDERGOER, which is not necessarily expressed at the surface, as *eat* can also appear intransitively. Finally, *becoming_broken* events always have an UNDERGOER, but an ACTOR role is only added (profiled) by the causative construction.

Goldberg's semantic coherence principle (Goldberg 1995: 50, Goldberg 2002: 342) requires roles contributed by a verb's lexical semantics and by a construction to be semantically compatible in order to be able to be fused. In the metagrammars developed here, the compatibility of roles is modeled by the type constraints indicating subtype relationships and explicit incompatibility relations in the type hierarchy. The correspondence principle requires expressed profiled participant roles contributed by verbs to fuse with profiled argument roles contributed by constructions. This principle is implemented here through the interface-based linking between attributes in lexical semantic frames and nodes in the syntactic structure of a construction.

The metagrammars developed here show that causativity is a relevant property of a number of constructions. The verbs covered in the example sentences were subtypes of *activity*, *change_of_state* or *causation*. Constructions can add a causative meaning to lexically non-causative verbs, as seen in the *InducedActionConstruction* or the *PrepositionalCausedMotionConstruction* classes, or they can completely unify with the frames contributed by lexically causative verbs, as seen with the sentences involving the verbs *insert* and *slice*. The construction classes featuring in the metagrammars presented here thus constitute a good foundation for future work on modeling related phenomena at the syntax-semantics interface in XMG or a comparable framework.

The use of typed, recursive semantic frames in the sense of Kallmeyer & Osswald (2013) allowed an optimal representation of the interactions between meanings contributed by verbal lexicon entries, argument lexicon entries and constructions. The derivation of a tree and frame structure for an input sentence based on a tree adjoining grammar is not transformational, but based on constraints; there is no ordering of operations, they all apply simultaneously following the rules specified in the grammar. Construction instantiation can be viewed as a similar constraint-based process. A construction can add an attribute to the semantic frame of a verb it interacts with, which is exemplified in some realizations of the constructions responsible for directional prepositional phrases in the CMC/IAA metagrammar. Alternatively, a construction can embed the verb's semantic contribution into an attribute in the semantic frame provided by the construction itself, which is exemplified in the causative constructions that are anchored by non-causative verbs, for instance in the inchoative construction. In all cases, meanings that are contributed by verbs, arguments and constructions all work together to bring about the overall meaning expressed by the sentence in which they appear.

Goldberg (2002: 349) argues against studying alternations as pairs of paraphrases or near-paraphrases, and instead for putting the focus on the individual argument structure constructions that characterize each alternation. She claims that this perspective makes it possible to identify generalizations that would not be apparent if alternations are researched separately from each other.

In this thesis, each alternation is modeled in a way that focuses on making exactly the constructions available to each verb that are licensed either by the verb's lexical semantics, or by the alternations in which the verb participates. There is no point in the metagrammar or in the lexicon where the two alternants of an alternation are directly associated with a verb as a pair of possible realizations. Instead, construction availability is determined via the interplay of each verb's lexical semantics, the lexical semantics of the observed arguments, the semantic requirements of the construction, and semantic type and attribute constraints in the frame type hierarchy.

This approach makes it possible to optimally exploit the factorization mechanisms provided by XMG. When alternations have an overlap in the argument structure constructions they license, as is the case for the causative-inchoative alternation and the instrument subject alternation, then all verbs participating in these alternations can access a shared set of metagrammar classes, which reduces redundancy and makes grammar maintenance and extensions easy. Unanchored metagrammar classes can also allow constructions to share part of their syntactic or semantic descriptions. For instance, in the combined metagrammar, both *GeneralCausedChangeOfState* and *GeneralCausedMotion* inherit parts of their semantic description from their common superclass *Causation*. In addition to these design decisions facilitating the computational treatment of verb alternations, there is also no linguistic motivation to treat the causative construction of one alternation as completely separate from the causative construction associated with another alternation.

The implemented metagrammars presented in this thesis part thus confirm Goldberg's position regarding the value of surface generalizations over that of alternations as paraphrase pairs.

This thesis part also implements the proposal by Levin (1993) that alternation participation is guided by certain meaning facets that are inherent to a verb's lexical semantics. However, not all alternation-specific constructions are exclusively characterized by such meaning facets here. There are also cases where constructions impose semantic requirements on the verb's arguments instead of or in addition to requirements for the verb itself. For instance, in the instrument subject construction, the INSTRUMENT that is located in the syntactic subject position must be compatible with a type that has here been referred to as *autonomous*, although a more fine-grained distinction of autonomous and non-autonomous entity types would be possible and may reflect the availability of this construction more accurately.

The design of the metagrammar makes it possible to define constraints at various levels: in the lexicon entries for verbs and their arguments, in the type hierarchy, and in the construction classes that can be anchored by specific lexical items. With these tools, the metagrammar can also reflect the constraints posited by Montemagni (1994: 351), who argues that certain verbs impose type requirements on their arguments in one alternant but not in the other one.

There are several purposes to developing a metagrammar modeling alternations and constructions in a language. Firstly, of course, any attempt to parse sentences in a language can benefit from taking semantic information into account instead of being limited to syntactic and morphological categories and structures. Secondly, expressing hypotheses about the relevant factors for construction availability in such a model can reveal whether any conflicts or contradictions exist in the theory that need to be addressed. New verbs or entities can be added to the metagrammar to test how well the developed theory generalizes to new contexts.

The metagrammars presented in this thesis part are completely handcrafted, with reference to theoretical literature about the phenomena under investigation, but without the involvement of any data-based learning routines to generate parts of the metagrammar without human supervision. In order to create a model like this on a larger scale, covering more than a handful of constructions, it would be desirable to be able to learn the relevant semantic distinctions automatically to a certain extent. For instance, the affordances of INSTRUMENTS for *eating* events have only been sketched here with the type *can_manipulate_food* – identifying which combination of properties it is, exactly, that is required from an object in order for it to be able to manipulate food, would make more fine-grained distinctions possible.

# Part IV

# ALTERNATION IDENTIFICATION

This part of the thesis discusses ways to identify whether a verb participates in a specific alternation or not, given that the verb can be observed in all relevant syntactic patterns that are associated with the constructions of the alternation. Here, the challenge is to determine whether the observed instances of the verb in a specific syntactic pattern do in fact instantiate the expected, alternation-specific construction.

# 9 ALTERNATION IDENTIFICATION: BACKGROUND AND TASK DEFINITION

## 9.1 INTRODUCTION

Part III of this thesis illustrated and discussed the way alternating verbs behave differently from non-alternating verbs, with a focus on modeling the required distinctions between constructions that share syntactic patterns but are associated with different semantic interpretations. The present part of the thesis is concerned with the task of determining for a given verb whether it participates in a specific alternation. Classifiers based on various feature sets will be implemented for the causative-inchoative alternation, the instrument subject alternation, and the together reciprocal alternation (intransitive).

The constructions associated with a verb alternation provide a set of syntactic patterns in which all verbs that participate in this alternation can appear. However, since alternation-specific constructions may share their syntactic patterns with other constructions that are unrelated to the alternation, the task of identifying verbs that participate a particular alternation is non-trivial: observable distributions of verbs in syntactic patterns are not sufficient to decide whether a verb alternates or not. Instead, sentences that instantiate similar syntactic patterns may be instances of different constructions, and thus be associated with different semantics. This can be the case even if the sentence's verb participates in a specific alternation.

As a result, applications or tasks that rely on syntactic cues to identify, for instance, semantic roles in a sentence need to be able to distinguish between alternating and non-alternating verbs, because the role assignment in each constructions follows certain idiosyncratic rules (Goldberg 1995: 1). Giving systems for tasks like frame-semantic parsing or semantic role labeling access to information about alternation participation is likely to improve their performance. The metagrammars presented in Part III relied on hard-coded constraints indicating alternation participation of individual verbs via a hand-crafted semantic type hierarchy. This part of the thesis will explore to what extent the alternation behavior of verbs can be determined based on unsupervised or semi-supervised features, where such type hierarchies are not available (or not entirely reliable).

Lexical resources like VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007) partially denote alternation participation of individual verbs or verb classes, but are not comprehensive; more verbs might participate in an alternation beyond the ones explicitly associated with the alternation in VerbNet.

Additionally, transferring these resources to other languages takes time, effort and linguistic expertise, as shown by Majewska et al. (2018). They study to what extent VerbNet classes and their members can be translated from English to other languages by human translators. With regard to diathesis alternations, they observe that alternation behavior can differ between languages and language families to an extent requiring a certain degree of language-specific tuning for each target language. For instance, English subcategorization frames that express an instrument in a prepositional phrase, as in *Paula hit the ball with a stick*, are realized in Polish without a preposition and with an instrumental case marker attaching to the instrument NP. Causativity, which plays a role in alternations like the causative-inchoative alternation, is in some languages marked morphologically with an affix attaching to the verb, while this is not the case in English. Some constructions involved in English alternations have no counterpart in another language; for instance, there is no resultative construction in Italian or in Polish. Creators of VerbNet-like resources in other languages can benefit from having access to a way to recognize, with little human annotation effort, which verbs in the language at hand participate in an alternation.

### 9.1.1 *Goals of this thesis part*

This thesis part aims to develop a set of classifiers for identifying verbs that participate in a selection of verb alternations in English. For each verb observed in a set of syntactic patterns, an alternation-specific classifier will assign a binary label denoting the verb's participation or non-participation in the alternations whose constructions are associated with those syntactic patterns. Such classifiers can be a valuable asset for extending VerbNet-like resources with more verbs per alternation, as well as transferring such resources to other languages than English. Ideally, they would be able to generalize from a few examples for a verb alternation in the given language to unseen verbs in that language, using unlabeled data or data with labels that are easy to acquire, such as dependency tree information (provided that a reliable parser for the target language exists).

As discussed in Chapter 2, verb alternations are characterized by the constructions they license for participating verbs. Each construction has its own syntactic structure and semantic interpretation. While syntactic patterns can be observed directly from dependency-parsed corpora, differentiating between different constructions that share syntactic patterns requires additional effort.

On the purely syntactic side, corpora can be used as a source of attestations for the verbs that are being assessed: if all relevant syntactic patterns for a specific alternation are felicitous for a specific verb, it is likely that that verb appears with each of them with a certain frequency throughout the corpus. However, the absence of instances of syntactic patterns for a specific verb cannot be understood as a reliable indicator of infelicity. This is a well-known issue in corpus linguistics (Chomsky 1957: 16, Kübler & Zinsmeister 2014: 17) as well as child language acquisition (Bowerman 1988: 74), although some (e.g., Stefanowitsch 2008: 527) argue that corpus-linguistic methods can in fact be used to distinguish between accidental gaps and significant ones. Corpora are generally considered a useful tool to study the presumed generalizations and felicity conditions on specific linguistic phenomena; as Bresnan et al. (2007) state it,

> [...] linguistic intuitions of ungrammaticality are a poor guide to the space of grammatical possibility [..., and] usage data reveals generalizations which we are sometimes blind to. (Bresnan et al. 2007: 8–9)

Since it is possible for different constructions to share the same syntactic structure, observing a verb in all alternation-specific syntactic patterns is not sufficient by itself to label that verb as participating in the alternation. If a pattern associated with an alternation is generally frequent in the language under investigation, then any observed instance of a verb in that syntactic pattern may in fact instantiate a construction that is not related to the alternation, and merely coincidentally shares the same syntactic form.

For instance, the constructions involved in the causative-inchoative alternation (CIA) are characterized syntactically by simple transitive and intransitive patterns, as illustrated in (56):

(56) a. The storm broke Ahab's compass. (*causative*)
b. Ahab's compass broke. (*inchoative*)

Transitive and intransitive sentences are so frequent in corpora of English that their occurrence with one specific verb does not necessarily indicate that this verb participates in this alternation. Another plausible explanation for the occurrence of both transitive and intransitive patterns for a specific verb could be that the verb participates in an unexpressed object alternation, or that it has multiple senses that take different sets of arguments.

On the other hand, alternations like the spray-load alternation (Levin 1993: 50) involve constructions that are syntactically characterized by more specific constraints, as illustrated in (57):

(57) a. Tashtego loads the bucket with spermaceti oil. (*"with" variant*)
b. Tashtego loads spermaceti oil into the bucket. (*locative variant*)

The *with* variant requires a prepositional phrase headed by *with*, while the locative variant requires a directional prepositional phrase. Sentences fulfilling these requirements are generally less frequent in corpora than the patterns discussed above for the causative-inchoative alternation. In fact, the dependency annotations that are provided with ENCOW (Schäfer & Bildhauer 2012, Schäfer 2015), the corpus that will be used as a data source for the experiments in this thesis part, show that the corpus contains over 1,000,000 simple intransitive sentences and over 9,000,000 simple intransitive sentences, while the subset of transitive sentences with a prepositional phrase headed by *with*, as in (57a), contains 325,040 sentences, and the subset of transitive sentences with a prepositional phrase headed by *into*, as in (57b), only consists of 64,069 sentences.[1] If the frequency of the syntactic patterns is correlated with the number of constructions that can be realized with the patterns, then corpus attestations of verbs in the patterns characterizing the spray-load alternations would be more indicative of alternation participation than attestations of the patterns characterizing the causative and inchoative constructions.

For these reasons, a purely syntactic approach to identifying alternating verbs is not sufficient. Instead, an additional step is required that determines whether the syntactic patterns in which a verb is observed do in fact instantiate the constructions relevant to a given alternation. The task becomes more difficult by the fact that the absence of evidence for one syntactic pattern does not equal evidence of infelicity – in other words, not observing one of the patterns for a verb in corpora is no sufficient indicator that

1 The numbers mentioned here are based on the annotations generated by MaltParser which are part of the downloadable ENCOW corpus. Since the data was parsed around 2015, misparses are expected to occur. However, even with misparses, the orders of magnitude that separate the frequencies of the different syntactic patterns are nevertheless notable.

the alternation is impossible for that verb. If, on the other hand, all relevant syntactic patterns are observed, this does not indicate that the verb at hand necessarily participates in the alternation, unless their semantic interpretations are consistent with the constructions involved in the alternation.

Concerning the acquisition of the relevant semantic properties to distinguish superficially-identical constructions, a complete, detailed semantic analysis of each sentence is not necessary. Instead, the semantic properties can be approximated by an unsupervised or semi-supervised setup. This will allow a classifier to estimate whether the attestations of a verb in question are consistent with the alternation-specific constructions (in which case the verb participates in the alternation) or not (in which case it does not participate).

Two issues pose challenges to this strategy of using corpus data to find attestations of verbs instantiating alternation-specific constructions. The first is the impact of specific argument types on the availability of constructions, as discussed by Montemagni (1994), and later by Romain (2018: 21), who argues that alternating verbs can have different "alternation strengths", so that participation in alternations is not actually a binary property, but a spectrum. The second is the impact of pragmatics and other factors on the choice between constructions involved in an alternation, as discussed by Bresnan et al. (2007).

The argument by Montemagni (1994) has been summarized previously in Section 3.3.3 of this thesis. In short, the issue is that verbs can impose type requirements on their arguments specifically in the context of certain constructions. For instance, the verb *break* can apply to a wide range of undergoer entities, including body parts. *Break* is generally viewed as a verb that participates in the causative-inchoative alternation. However, body part undergoers do not seem to be allowed as subjects in the inchoative construction, as illustrated in (58) (taken from Montemagni 1994: 351):

(58) a. John broke the window.
b. The window broke.
c. John broke his arm.
d. * His arm broke.

As a result, the instances of verbs like *break* in a corpus may show a gap for this type of undergoer, seemingly an indicator against the verb's participation in the causative-inchoative alternation. In a world with infinite, high-quality corpus data available, one might be able to conclude that *break the window* participates in the alternation, but *break an arm* does not, and similarly with other verbs exhibiting this type of behavior. In the actual experiments conducted in this thesis part, the classifiers rely on general trends of entity types for each syntactic argument slot for each candidate verb, and may tend to misclassify verbs that behave like *break*.

The argument by Romain (2018: 21) is related to that of Montemagni (1994). She also argues for taking into account not only verbs and constructions, but also the undergoer argument. While neither Montemagni nor Romain are concerned with automatically labeling verbs as participating or not participating in a particular alternation, their point is still relevant here: Levin (1993) and VerbNet label verbs as either participating in an alternation or not participating, but do not relate the availability of each construction to the content of the verbal arguments. Those cases may negatively impact the performance of the classifiers developed in this thesis.

Bresnan et al. (2007) are concerned with the choice speakers make when more than one construction is available to express a verb's arguments, as is the case in certain verb alternations. This is exemplified with the double object construction and the prepositional object construction, both of which are involved in the dative alternation. The authors find that factors like argument weight, givenness, definiteness, animacy, and pronominality of the verb's arguments are fairly predictive of the choice of one construction over the other. They report 94% accuracy on a task classifying 2360 instances of the two constructions of the dative alternation, covering 38 verbs. Notably, none of their features focus on the verb in each instance; only the properties of the arguments are taken into account to predict which construction is preferred.

These discourse factors may also play a role for the alternations under investigation here. As a result, certain arguments may predominantly appear in one slot of one construction involved in the alternation, and rarely or never be observed in the corresponding slot of the alternate construction. Again, the classifier may conclude that the verbs for which this is the case do not participate in the alternation. However, this phenomenon is certainly impacted by the semantic similarity or difference between the constructions belonging to an alternation: if the constructions are semantically very similar, then the discourse factors are more likely to play a role than for constructions that are semantically very dissimilar, because in those cases, the choice between the available constructions would presumably be guided more strongly by semantic factors.

Note that the task of identifying verbs that can instantiate alternation-specific constructions does not amount to developing a construction classifier. In the context of this thesis, verbs will be classified on the basis of types, not tokens. All attestations of each verb are taken into account in order to make this lemma-level decision. Individual instances of the verb are not classified as to whether or not they instantiate an alternation-specific construction.

An example of a construction classifier is presented by Hwang & Palmer (2015), who are concerned with recognizing instances of the caused-motion construction, specifically in contexts where the verb itself does not lexically encode motion. The syntactic structure for this construction involves a direct object and a prepositional phrase; however, this syntactic structure is actually associated with not just one construction, but a series of (related) constructions (see Chapter 6 of this thesis).

Candidate sentences for this task are sourced from several constituency-parsed corpora, based on the required syntactic pattern for the construction. The sentences are manually annotated (Hwang, Zaenen & Palmer 2014), distinguishing four different types of the caused-motion construction: displacement, change of scale, change of possession, and change of state. The classifier has access to features characterizing the verb, the preposition, the complement of the preposition, the cause argument, and the undergoer argument. The baseline only uses the verb lemma as an indicator for or against a caused-motion label.

For the task of determining whether a sentence instantiates the caused-motion construction or not, the authors achieve a best $F_1$ score of 83.5% (annotator agreement: 0.606) on the Broadcast News portion of OntoNotes (Weischedel et al. 2011).

In contrast to that approach, the experiments presented in this thesis operate on the level of verb lemmas, using all observed instances of each candidate verb in certain syntactic patterns as indicators for or against alternation participation. No annotated construction dataset is available for the candidate verbs for the alternations that will be investigated. Implementing a construction classifier for each alternation-specific construction would be an interesting addition to the experiments presented here, but is outside the scope of this thesis. Note also that the construction classifier by Hwang, Zaenen & Palmer (2014) performs well, but does not achieve perfect accuracy – identifying specific constructions is not a solved task. The more recent work by Dunietz, Levin & Carbonell (2017) focuses less on the syntactic patterns that characterize verb alternations, but approaches the identification of causal constructions in a more general way, and is thus not applicable in the context of this thesis.

### 9.1.2 *Hypotheses for this thesis part*

The present part of this thesis is concerned with operationalizing the alternation identification task and implementing and discussing several approaches to it. The experiments and discussions are designed to assess the hypotheses listed below. The remainder of this chapter is dedicated to defining and preparing the alternation identification task. Chapter 10 presents a series of feature sets for the classification and discusses their performance for several English verb alternations. Section 10.6 will discuss the results with regard to the hypotheses detailed above.

H1. *Verbs that participate in alternations whose constructions have generally frequent syntactic patterns are harder to identify than verbs that participate in alternations with less frequent syntactic patterns.* As argued above, syntactic patterns can be associated with multiple unrelated constructions. Thus, alternations whose patterns appear numerous times in a corpus may cause ambiguity problems, making it more difficult to determine which of several possible constructions is instantiated in each case. The number of participating verbs in each alternation may also play a role.

H2. *Using only the number of attestations of certain syntactic patterns for a verb in a corpus is a weak strategy to identify verbs in any alternation; it may be more successful for alternations with less frequent patterns than for alternations with frequent syntactic patterns.* For alternations whose constructions are less impacted by ambiguity, using the number of attestations of the different syntactic patterns associated with the alternation may be sufficient to determine whether a given verb participates in it or not. A frequency-based baseline will be implemented in order to determine whether this is the case.

H3. *Alternating verb identification is improved by features that approximate the selectional preferences of each candidate verb for its argument slots in the relevant syntactic patterns.* Syntactic patterns can be associated with multiple constructions. Features that approximate selectional preferences may lead to better classification results compared to the frequency-based baseline, because they provide a means of distinguishing between constructions with identical syntactic structures.

H4. *Features that approximate the acceptability of an alternation's constructions for a candidate verb can lead to better classification results, especially for alternations whose patterns are attested infrequently.* Each alternation involves two or more constructions that license different syntactic patterns. A classifier may benefit from features that encode to what extent each observed instance of one of these patterns can be transformed into the alternative pattern(s) of the relevant alternation and yield an acceptable resulting sentence. Since this method augments observed instances with hypothetical alternate instances of the alternation's patterns, it may lead to stronger results than the features approximating selectional preferences based on observed instances, particularly in cases where not enough instances of a verb in a pattern are observed to generalize over slot fillers.

## 9.2 RELATED WORK

This thesis is by no means the first work to attempt to classify verbs with respect to alternation participation. Computational linguistics research has been done on this type of task throughout the past two decades, involving various corpora, lexical resources, classification features, human annotation strategies and language models. In attempting to compare the results, it is important to acknowledge the fact that the wide range of resources that have been employed makes it difficult to relate the different approaches to each other or rank them.

This section gives an overview of existing research on the classification of verbs participating in verb alternations. No qualitative or quantitative comparison will be attempted, for the reason stated above. Occasionally, a publication will be cited whose main goal is not to classify verbs with respect to alternation participation, but instead to cluster verbs semantically, or to present a new method for learning semantic properties of verbs that is merely being illustrated with an alternation classification task. These references are included because they still contribute to the landscape of existing methods for studying verbs and taking their alternation behavior into account, in a broader sense.

The majority of published work in this area is concerned with the causative-inchoative alternation, often accompanied by one or more additional alternations, such as the simple reciprocal intransitive alternation in McCarthy & Korhonen (1998), the conative alternation in McCarthy (2000), or the dative alternation, the spray-load alternation, there-insertion, and the unexpressed object alternation in Kann et al. (2019). Lapata & Brew (1999) attempt to predict Levin verb classes and evaluate their results for verbs in the dative or benefactive alternation. The only existing publication involving the identification of verbs participating in the instrument subject alternation, which features in this thesis, is based on preliminary work by myself (Seyffarth & Kallmeyer 2020). The together reciprocal alternation (intransitive) does not seem to be covered in any existing published work in this area.

The publication of Levin (1993), and later, VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007), made it possible to work on a shared basis of canonical alternating verbs in English. At the same time, one cannot make corpus-based claims about verbs appearing in these resources that are unattested in the corpus one uses. A supposedly alternating verb can also be attested in a corpus, but not in one or both of the relevant constructions. On top of that, the distribution of alternation-specific constructions in a corpus seems to also be influenced by text type and modality: Bresnan et al. (2007: 25) note that the distribution of the constructions involved in the English dative alternation that they learn from the Switchboard corpus (Godfrey, Holliman & McDaniel 1992), which contains recorded telephone conversations, are only weakly predictive of the distribution of the same constructions in the Wall Street Journal Treebank corpus (Marcus, Santorini & Marcinkiewicz 1993), which contains written text covering news and finance.

Additionally, for classification tasks, the selection of classes to contrast with each other and of instances representing these classes is also an individual decision. Verbs in the causative-inchoative alternation may be contrasted with verbs in the unexpressed object alternation, as in Seyffarth (2019a), or simultaneously with verbs in that alternation and with unergative verbs, as in Merlo & Stevenson (2001), or with verbs that never appear in at least one of the constructions involved in the causative-inchoative alternation, as in Baroni & Lenci (2009), who populate their negative class with verbs that are specifically listed by Levin (1993) as not participating in the causative-inchoative alternation. The different classes may be balanced, as in McCarthy (2000), who uses 46 verbs per class for the causative-inchoative alternation and 6 verbs per class for the conative alternation, or unbalanced, as in Baroni & Lenci (2009), who use 232 alternating verbs and 170 non-alternating verbs.

Corpora that have been used as the basis for alternation participation classification include the Switchboard corpus with three million words (Godfrey, Holliman & McDaniel 1992) in Bresnan et al. (2007), the WSJ corpus with roughly three million words (Marcus, Santorini & Marcinkiewicz 1993) in Bresnan et al. (2007), Merlo & Stevenson (2001), the BNC corpus with 100 million words (Leech 1993) in Lapata & Brew (1999), and the Brown corpus with one million words (Francis & Kučera 1964, 1971, 1979) in Merlo & Stevenson (2001). External tools for the acquisition of subcategorization frames are occasionally used, such as the Briscoe & Carroll (1997) system in McCarthy (2000), Gsearch (Corley et al. 2001) in Lapata & Brew (1999) or the Preiss, Briscoe & Korhonen (2007) system in Sun, McCarthy & Korhonen (2013).

Most work in this area regards syntactic subcategorization frames as instances of the corresponding constructions. Whether this is sufficient to identify alternating verbs depends on the context of the specific task. Sun, McCarthy & Korhonen (2013) attempt to learn verb clusters based on the verbs' alternation behavior. Their approximation of a diathesis alternation is derived from the pairwise joint probability of different (syntactic) frames, obtained by integrating over all possible diathesis alternations. If a frame pair's joint probability is higher than chance, assuming an independent distribution of frames, they view the verb as likely to participate in the alternation that is characterized by those two frames. As the authors themselves note (Sun, McCarthy & Korhonen 2013: 740), this strategy is missing an approximation of selectional preferences, which they suggest could be added via distributional data over argument heads.

In contrast to this, Merlo & Stevenson (2001) conduct a three-class classification task on verbs that can appear in transitive or intransitive environments for different reasons, with each class being able to instantiate a different pair of constructions. Obviously, for this task, the distribution of syntactic subcategorization frames for the candidate verbs is not a particularly informative indicator for class membership, and semantic features are defined to attempt to distinguish between the three construction pairs characterizing the different classes.

A curious combination of these different approaches can be found in Baroni & Lenci (2009). They base their classifier on distributional information about observed verb arguments, but contrast their set of verbs in the causative-inchoative alternation with a set of non-alternating verbs that are assumed by Levin (1993) to not be able to instantiate both alternation-specific constructions anyway. Under these circumstances, the authors report that their model is "completely successful in detecting the distinction" (Baroni & Lenci 2009: 7).

Finally, each work in the area of alternating verb classification presents its own set of methods for distinguishing the classes. The spectrum of proposed features for the distinction includes subcategorization frame distributions, as in Sun, McCarthy & Korhonen (2013), supervised, linguistically motivated features on argument heads, as in Merlo & Stevenson (2001), Sun & Korhonen (2009), referring to the WordNet noun hierarchy to group argument types, as in Schulte im Walde (2000), distributional information about argument heads, as in Baroni & Lenci (2009), and verb and sentence embeddings, as in Kann et al. (2019). Most of the work cited in this section relates directly to one or more of the feature types that will be presented in Chapter 10 of this thesis. The specific methods described in the referenced publications will be reported on in the respective sections of that chapter.

## 9.3 task setup

This section gives an overview of the methods and tools employed for the experiments in this thesis part. Details on the various steps of the process can be found in Sections 9.4, 9.5 and 9.6 and in Chapter 10.

### 9.3.1 *Task and method overview*

The classification of verbs as participating or not participating in each alternation will be implemented with support vector machine (SVM) classifiers. This algorithm is well-suited for binary classification tasks on a balanced dataset. Different sets of linguistically motivated features will be defined in the following chapter, and the performance of the alternation-specific classifiers with each feature set will be interpreted and evaluated separately. An introduction to SVM classifiers will be given in Section 9.3.2.

Before the classification experiments can be set up, a set of alternations will be chosen for which the classifiers will be implemented. Since the hypotheses given in Section 9.1 make predictions about the effects of syntactic pattern frequency, the number of participating verbs, and differences in slot-specific selectional preferences, the goal is to compare classification results on alternations that vary in these properties. The selection of alternations for the experiments will be described in Section 9.4. The selected alternations are the causative-inchoative alternation, the instrument subject alternation, and the together reciprocal alternation intransitive. The syntactic patterns characterizing the constructions involved in these alternations will be described in terms of dependency relations in order to facilitate the creation of alternation-specific corpora.

In this thesis, the classification is set up as a binary task: for a given candidate verb for an alternation, the alternation-specific classifier will determine based on corpus data whether the verb participates in the alternation or not. Romain (2018: 21) notes that verbs may have different degrees of alternation strength, that is, a verb may participate in an alternation but have a weak or strong preference for one of the relevant constructions; this distinction will not be made in this thesis.

Corpus attestations for verbs in the relevant syntactic patterns will be sourced from a dependency-parsed English corpus, ENCOW. A subcorpus for each alternation is created, containing all sentences extracted from the corpus that contain the dependency relations characterizing the constructions associated with the alternation. Details on the corpus and the creation of alternation-specific subcorpora will be provided in Section 9.5.

The members of the positive class ("positive verb set") for each alternation will be sourced from VerbNet and reduced as necessary to ensure that each verb is seen in the corpus data at least once in the pattern associated with each relevant construction. Additionally, a "negative verb set" for each alternation will be created that contains verbs that are also observed in the relevant syntactic patterns, but do not participate in the alternation according to VerbNet.

Theoretically, the negative verb set for an alternation could comprise every single verb in the language under investigation except the ones that are members of the positive verb set. However, this would lead to an imbalance in the class sizes that would make it difficult to evaluate and compare classification results. Instead, each negative verb set will contain exactly the same number of verbs as its corresponding positive verb set.

Furthermore, verbs are only added to the negative verb set for an alternation if they are observed in the corpus at least once in each relevant syntactic pattern for the alternation. Without this condition, classifying verbs would amount to simply checking if they are attested in the alternation-specific syntactic patterns or not.

The strategy for collecting positive and negative verb sets will be described in detail in Section 9.6.

To train the alternation-specific classifiers, each sentence in each subcorpus is treated as an example of the relevant alternation construction if the sentence's root verb is a member of the positive verb set for the alternation. If the root verb is in the negative set, the sentence is regarded as an example of a construction that shares the same syntactic pattern but is otherwise unrelated to the alternation. This allows the classifiers to learn differences in the behavior of verbs in the alternation and verbs outside the alternation. Suppose that an alternation is associated with a construction A, characterized by syntactic pattern X, and another construction B, characterized by syntactic pattern Y. Then, the classifier will regard all instances of pattern X involving a known alternating verb as instances of construction A, and analogously for pattern Y and construction B. For instances of patterns X and Y with a known non-alternating verb, the classifier assumes that the patterns are not instances of constructions A and B, but of other constructions that merely share the same syntax.

Treating each occurrence of a verb in a certain syntactic pattern as an example of a particular construction, just because the verb is known to license that construction, is a simplification. Construction identification is a nontrivial task. Instead of attempting to determine whether individual sentences instantiate a particular construction or not, the classifiers implemented here will label verbs at the lemma level, based on all attestations of the verbs in the relevant syntactic patterns. The underlying assumption is that even if alternating verbs occasionally appear in sentences that do not instantiate the relevant construction but a syntactically similar one, it is still more likely for sentences with these verbs to instantiate the construction than it is for sentences with non-alternating verbs.

Chapter 10 presents each set of classification features and reports the accuracy scores of classifiers using each feature set for the different alternations. The frequency-based baseline classifier feature (Section 10.2) operates purely on the frequencies of the different pattern types for each alternation and puts them in relation to each other. The lemma-based classifier features (Section 10.3) attempt to approximate selectional preferences on the different argument slots in the alternation-specific patterns based on observed slot filler lemmas for each candidate verb. The vector-based classifier features (Section 10.4) attempt to generalize over observed slot filler lemmas using distributional information. Finally, the perplexity-based classifier features (Section 10.5) compare each attested instance of a verb in one of the alternation-specific patterns with a generated alternate version of that sentence in which the other pattern type is instantiated. The perplexity scores for these pairs of observed and generated pattern instances are an attempt to approximate the acceptability of the constructions associated with the respective syntactic patterns.

Since the negative set for each alternation contains verbs that are observed in all relevant syntactic patterns, presumed negative candidates may in fact be verbs that actually participate in the alternation, but are missing from the positive set for some reason. This is particularly likely for verbs that are labeled as "false positives", that is, verbs for which the classifier predicts a positive label while the gold label is negative. A procedure involving manual annotation of these false positives in order to discover new alternating verbs is presented and discussed in Section 10.7.

### 9.3.2 *Support Vector Machine classifiers*

The classifiers for the alternation identification task in this thesis part will be implemented as support vector machine (SVM) classifiers. Previous work in this area (Seyffarth & Kallmeyer 2020) showed that SVMs are a good choice for the identification of verbs participating in alternations, beating competing algorithms such as k-nearest neighbor and Naive Bayes. SVM classifiers are a popular choice for classification tasks in natural language processing and have been used in a range of contexts, including relation extraction (e.g. Hong 2005), dependency parsing (e.g. Nivre, Hall & Nilsson 2006), spelling correction (e.g. Duan et al. 2012), semantic similarity quantification (e.g. Vo, Caselli & Popescu 2014), identification of cognates (e.g. Jäger, List & Sofroniev 2017), and many others.

SVM classifiers take as their input a set of data points that are represented in a multidimensional vector space according to their values for each feature, such that each feature corresponds to one dimension in the vector space. During training, the goal of the classifier is to place a separating hyperplane in the vector space that creates a maximum margin between instances belonging to different classes (Vapnik 1995). The term "support vector" refers to the data points from the labeled training set that are closest to the optimal separating hyperplane. Features that are more discriminative make it easier for the classifier to determine the optimal location and angle of the separating hyperplane.

During classification, the feature values for an unseen data point determine its placement to one side or the other of the separating hyperplane, which leads to the classifier assigning to the data point the class label that is associated with that side of the hyperplane. Multi-class SVMs are also possible, but will not be used here, since the verbs under consideration here will be labeled with binary labels denoting whether or not they participate in each alternation.

The classification scripts developed in the context of this thesis rely on the SVC implementation with a radial basis function kernel that is part of the scikit-learn toolkit (Pedregosa et al. 2011) for Python (using Python 3.9.2, with version 0.24.2 of `sklearn`). The SVC class is based on libsvm (Chang & Lin 2011).

SVM classifiers can be tuned by configuring several hyperparameters. C is a regularization parameter that impacts the width of the margin around the separating hyperplane: a wider margin is associated with a simpler decision function, but may also lead to a larger number of misclassified training examples. $\gamma$ is applicable to SVMs using the radial basis function kernel and determines the radius of influence of each training example. Both parameters can be used to increase classification accuracy on the training set, but pose a certain risk of overfitting if the distribution of instances in the test set is not learned well by the trained classifier. The parameters are documented in detail on the scikit-learn website at `https://scikit-learn.org/stable/auto_examples/svm/plot_rbf_parameters.html`.

The linguistically motivated features for the classifiers in this thesis are expected to be predictive of alternation participation to different extents, depending on the alternation. In other words, the instances in each alternation-specific training set may cluster differently in a way that warrants different values for the hyperparameters. At the same time, the goal of this thesis part is to determine for each alternation which features work well and which ones do not. For the sake of comparability of the classification results across different sets of training and test instances, the recommended default hyperparameter values for SVC will be used. C will be set to a value of 1, and $\gamma$ will be set to the value of the inverse of the standard deviation of the current training set.

In a later step, the classifiers will be improved further based specifically on the features that have been found to be most indicative of participation in each of the alternations (see Section 10.6). The hyperparameters will be set per-alternation in order to determine the maximal achievable classification scores for the features for each alternation.

Due to the relatively small set of instances in these classification tasks, each of the classifiers will be trained and evaluated with 10 rounds of 10-fold cross-validation, based on random splits of the data into a train portion and a test portion. Each bucket will contain the same number of positive and negative instances (stratified k-fold split).

## 9.4 ALTERNATIONS UNDER INVESTIGATION IN THIS THESIS PART

The alternations that will be examined in this part of the thesis need to fulfill a set of conditions in order to be usefully employed to test the hypotheses formulated in Section 9.1:

1. They should vary in the frequency of their syntactic patterns; at least one should have more frequent syntactic patterns, at least one should have more infrequent patterns.

2. They should vary in their semantic properties; for instance, the selectional preferences that are imposed on different argument slots should have an overlap in at least one of the alternations, and be more distinct in at least one.

3. They should vary in size; at least one of the alternations should encompass a large number of verbs, and at least one should apply to a smaller number of verbs.

Since this thesis part aims to design as well as evaluate several approaches to alternation identification, the focus will be on alternations for which reliable gold examples are available. VerbNet 3.4 (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007, available for download in `json` format at `https://uvi.colorado.edu/nlp_applications`) marks 17 different alternations explicitly. In addition, VerbNet contains information about causative and inchoative uses of verbs in a class, so that the causative-inchoative alternation is also available based on VerbNet data.

### 9.4.1 *Choosing alternations: the causative-inchoative alternation, the instrument subject alternation and the together reciprocal alternation (intransitive)*

Table 3 on page 204 lists all alternations that are explicitly marked in VerbNet. Alternations with fewer than 30 participating verbs will not be considered for the classification experiments. Based on the table and the properties of the alternations, three alternations are selected that will be used in the classification experiments here. They are specifically chosen because their properties differ in ways that are relevant to the hypotheses formulated for this thesis part.

THE CAUSATIVE-INCHOATIVE ALTERNATION (CIA). As discussed already in Part III of this thesis, the causative-inchoative alternation (Halliday 1968: 184, Fillmore 1970a: 253, Vendler 1972: 210, Hale & Keyser 1986: 607, Croft 1986: 231, Levin 1993: 27) allows verbs to appear in a transitive, causative construction or in an intransitive, inchoative construction. Participating verbs typically express some change of state of the theme. The inchoative construction is characterized by an intransitive syntactic frame in which the theme is expressed as the syntactic subject; no external cause of the change of state is expressed in these sentences. The causative construction is characterized by a transitive syntactic frame, where the syntactic subject is the position of the agent or cause that causes the change of state undergone by the theme, which is expressed in the syntactic direct object position.

Existing work on automatic alternation identification often focuses on the causative-inchoative alternation (see Section 9.2 of this thesis). While VerbNet does not assign a label for this alternation to verb classes explicitly, some verb classes in VerbNet are marked as licensing a causative frame as well as an inchoative one, which corresponds to participation in the causative-inchoative alternation. Nine VerbNet classes participate in this alternation, containing a total of 453 verb lemmas. This makes it one of the largest alternations in VerbNet in terms of participating verbs.

The constructions involved in the causative-inchoative alternation are associated with syntactic patterns that are generally frequent in English. Expressing the theme in the subject position in intransitive sentences is unusual in English and often points to an inchoation; most verbs that allow this type of construction also allow its causative counterpart. A challenge for the automatic identification of verbs that participate in this alternation consists in distinguishing between verbs that exhibit intransitive behavior without expressing an inchoation, and verbs for which the intransitive pattern is strongly associated with an inchoative meaning.

Table 3: Alternations that are explicitly marked in VerbNet 3.4. The number of verbs for each alternation is summed up over all participating VerbNet classes. Examples are from the relevant chapters in Levin (1993). Rows with gray text indicate infrequent alternations (fewer than 30 participating verbs).

| Alternation | Verbs | Construction 1 | Construction 2 |
|---|---|---|---|
| Instrument Subject Alternation | 558 | David broke the window with a hammer | The hammer broke the window |
| Pro-Arb Object Alternation | 251 | The sign warned us against skating on the pond | The sign warned against skating on the pond |
| Together Reciprocal Alternation Transitive | 177 | I creamed the sugar into the butter | I creamed the sugar and the butter together |
| Locatum Subject Alternation | 113 | I filled the pail with water | Water filled the pail |
| Together Reciprocal Alternation Intransitive | 52 | The eggs mixed with the cream | The eggs and the cream mixed together |
| Material/Product Alternation Transitive | 41 | Martha carved a toy out of the piece of wood | Martha carved the piece of wood into a toy |
| Simple Reciprocal Alternation Transitive; Prepositional Variant | 34 | I mixed the sugar into the butter | I mixed the sugar and the butter |
| Fulfilling Alternation | 33 | The judge presented a prize to the winner | The judge presented the winner with a prize |
| Benefactive Alternation; double object | 30 | Martha carved a toy for the baby | Martha carved the baby a toy |
| Apart Reciprocal Alternation Transitive | 22 | I broke the twig off (of) the branch | I broke the twig and the branch apart |
| Apart Reciprocal Alternation Intransitive | 22 | The twig broke off (of) the branch | The twig and the branch broke apart |
| Location Subject Alternation | 13 | We sleep five people in each room | Each room sleeps five people |
| Total Transformation Alternation Transitive | 11 | The witch turned him into a frog | The witch turned him from a prince into a frog |
| Total Transformation Alternation Intransitive | 11 | He turned into a frog | He turned from a prince into a frog |
| Attribute Object Possessor-Attribute Factoring Alternation | 10 | I sensed his eagerness | I sensed the eagerness in him |
| Material/Product Alternation intransitive | 8 | That acorn will grow into an oak tree | An oak tree will grow from that acorn |
| With Preposition Drop Alternation | 4 | Jill met with Sarah | Jill met Sarah |

The different semantic roles that appear in sentences instantiating the causative or inchoative construction are expected to be subject to different selectional preferences. Since participating verbs typically denote a specific change of state that the theme undergoes, the theme must semantically be able to be in the relevant state, while the cause must fulfill the conditions for causativity that are in place generally, independently of the specific verb in the sentence at hand. For the verb *open*, for instance, the set of possible themes is the set of entities that can be in an open or non-open state, and the set of possible causes is the set of entities that are able to affect a change between those two states in a theme, for instance by a physical action.

In this thesis, selectional preferences on participants of events have so far been expressed in the form of constraints on types or valid attributes of entities in a semantic frame (see Part III). For the purposes of the current thesis part, these constraints will not be explicitly specified. Instead, they will be hypothesized as a unifying property of entities that can appear in certain semantic roles (and thus, syntactic positions of constructions).

With respect to the causative-inchoative alternation, this amounts to a requirement on the undergoer of the event to have an attribute that can be changed by the inchoative or causative event expressed by the verb. The assumption is that change-of-state verbs typically only apply to undergoers that are unified by their ability to possess certain attributes related to the particular change of state denoted by the specific verb.

THE INSTRUMENT SUBJECT ALTERNATION (ISA). The instrument subject alternation (Fillmore 1970b: 126, Levin & Rappaport 1988: 1071) has also already been covered in Part III of this thesis. It allows verbs to express the instrument role either in the syntactic subject position, or as an adjunct in a phrase headed by the preposition *with* or the verb *using*. This alternation is fairly large and contains 543 verb lemmas according to VerbNet.

The instrument subject construction imposes no further syntactic conditions on a sentence except that there needs to be a subject. This means that filtering the original corpus data for this syntactic context will result in an extremely large subcorpus, and that attestation in the syntactic pattern associated with this construction is expected to be a fairly weak indicator for alternation participation, since it is extremely frequent for English sentences to have a syntactic subject.

The agent subject construction requires a *with* or *using* phrase modifying the verb. Observing a verb in a sentence with a *with* phrase may be an indicator for that verb's participation in the instrument subject alternation, but since *with* is highly polysemous, that phrase will not necessarily always express an instrument. It might alternatively express something like accompaniment or spatial positioning.

Another challenge for the identification of verbs in this alternation is that explicitly mentioning an instrument is often optional. When canonical instruments are used for an action described by a particular verb, this is often not expressed on the surface in a sentence. For instance, you would typically *stir a drink* with a *spoon*; this does not mean that a *spoon* will always be mentioned explicitly in sentences with the root verb *stir*. In fact, the instruments that are explicitly mentioned may even be more likely to be atypical instruments for the action described by the verb; for instance, if someone stirs their drink with a *fork*. For more on the problem of using corpus data to learn about aspects of the lexical semantics of terms that are common world knowledge and thus often not explicitly stated on the surface, see Herbelot (2013).

The fact that mentioning an instrument is often optional will have an impact on the subcorpus collection for this alternation. Even when a verb describes an action that is typically conducted with the help of some instrument, that verb may not be attested with explicit instruments at all in the corpus, or it may be attested with explicit instruments only infrequently, possibly more often with atypical instruments than with typical ones.

The instrument subject construction does not express an agent or cause explicitly on the surface. In the other construction, which includes a *with* or *using* phrase containing the instrument, the subject typically expresses an agent who controls the instrument in order to perform the action described by the verb.

As with the causative and inchoative constructions, the relevant semantic roles come with distinct selectional preferences: the instrument is a concrete or abstract entity that can be used to perform the action described by the verb, while the agent is an animate entity that is capable of controlling an instrument to perform the action described by the verb. Note that instruments do not have to belong to a specific semantic category, like *tools*: one can use a *key* to open a door, but also a *crowbar* or even a *foot*, or a *magical incantation* in the context of a fairy tale story. These instruments are semantically diverse, but share the common property that they facilitate or enable the action described by the verb, *open*. Although the selectional preferences on the instrument slot are highly diverse, they typically contrast well with the selectional preferences for the agent role, which is subject to expectations of animacy and agency that do not apply to the instrument role.

THE TOGETHER RECIPROCAL ALTERNATION (INTRANSITIVE) (TRAI). The together reciprocal alternation (intransitive) (Fillmore 1970a: 261, Vestergaard 1973: 87, Levin 1993: 64) involves verbs that have two or more co-themes. As illustrated by the example sentences for this alternation in Table 3 on page 204, these themes can either both appear in the subject position, connected by a conjunction and accompanied by the adverb *together*, or one theme can be located in the subject position with the other being located in a prepositional phrase. Levin (1993: 64) mentions specific prepositions (*with*, *into*, *to*) for the different verbs that are listed as participating in this alternation.

The together reciprocal alternation (intransitive) is a much smaller alternation in terms of participating verbs than the two alternations described above: according to VerbNet, only 50 verbs participate in this alternation.

Moreover, the syntactic pattern characterizing the *together* construction is far less frequent than the patterns for the constructions associated with the other alternations. The ENCOW corpus only contains 2484 sentences that are intransitive and have the word *together* as a direct dependent of the root verb. Since this pattern is so infrequent, many alternating verbs are likely to not be attested with it even though it would be feasible and felicitous. On the other hand, the pattern may turn out to be specific enough that observing a verb in it is a strong indicator for participation in this alternation.

In the description of this alternation by Levin (1993: 64), the *together* construction always appears with two co-themes connected by a conjunction in the subject position. Since the number of attestations of the *together* pattern in ENCOW is so low, the pattern can be extended for the purposes of the experiments conducted here to also include plural subjects (as in *The boats crashed together*).

Concerning the construction characterized by a prepositional pattern, the same issues apply that were mentioned above for the agent subject construction with a *with* phrase. Prepositions tend to be polysemous, so that observing an intransitive verb with a prepositional phrase is expected to be only a weak indicator for that verb's participation in the together reciprocal alternation (intransitive), even when the same verb is also observed in a *together* sentence.

The causative-inchoative alternation and the instrument subject alternation are what McCarthy (2001: 3) calls role-switching alternations, since they each involve one semantic role – theme for the causative-inchoative alternation, instrument for the instrument subject alternation – that is expressed in different syntactic positions in the different constructions involved in the given alternation. This is not the case for the together reciprocal alternation (intransitive). The only relevant semantic role for this alternation is the (co-)theme, and the same set of semantic arguments is expressed in both constructions that are involved in this alternation, albeit in different syntactic positions. This means that the same selectional preferences are imposed on all relevant syntactic positions and semantic roles for each verb that participates in this alternation: each of the different positions needs to be filled by something that can semantically function as a theme of the event described by the given verb.

For the causative-inchoative alternation and the instrument subject alternation, the different selectional preferences on the semantic arguments should be useful for determining whether observed instances of each syntactic pattern are examples of the relevant construction or not. For the causative-inchoative alternation, intransitive sentences should express a theme in the subject position, so a sentence whose subject is unlikely to fill the theme role would be an indicator against alternation participation. With respect to the instrument subject alternation, the same holds for instruments in the subject position of sentences without a *with* or *using* phrase. For the syntactic patterns involved in the together reciprocal alternation (intransitive), non-alternating verbs may assign different semantic roles to the various syntactic argument slots instead of expressing a theme in each slot.

### 9.4.2 *Pattern types with different numbers of alternation-relevant syntactic arguments*

While the three alternations under investigation here differ in terms of their size, the association between syntactic positions and semantic roles in their constructions, and the selectional preferences imposed on the semantic arguments, the remainder of this thesis part will regularly refer to properties that all three alternations have in common. To make this easier, a number of alternation-independent terms will now be introduced that can then be used to describe classification features in an alternation-agnostic way.

Although the constructions involved in the causative-inchoative alternation can be realized with two syntactic patterns, the constructions for the instrument subject alternation involve three syntactic patterns and the constructions for the together reciprocal alternation (intransitive) involve four syntactic patterns, for the purposes of these experiments, each of these alternations can be described in terms of two pattern types. The term *pattern type* is introduced here to describe the syntactic behavior exhibited by alternating verbs; however, non-alternating verbs can and will exhibit these pattern types as well, because the patterns are not necessarily linked to the alternation-specific constructions.

For the causative-inchoative alternation, each construction corresponds to one pattern type. For the instrument subject alternation, the patterns associated with the agent subject construction involving a *with* or *using* phrase can be summarized into one pattern type that contrasts with the pattern of the instrument subject construction. For the together reciprocal alternation (intransitive), the patterns of all three variants of the prepositional construction form one pattern type that contrasts with the pattern of the *together* construction.

Mentions of pattern types in this thesis part always refer exclusively to syntactic patterns, not to their corresponding alternation-specific constructions. This is because instances of each pattern type are also possible and attested for non-alternating verbs.

The classification features defined in Chapter 10 of this thesis will be based on what will be referred to as the *one alternation-relevant syntactic argument (1ARSA) pattern type* and the *two alternation-relevant syntactic arguments (2ARSA) pattern type* for each alternation. The 1ARSA pattern is always the one containing fewer syntactic arguments:

- For the causative-inchoative alternation, the 1ARSA pattern is the intransitive pattern of the inchoative construction that expresses the theme in the subject position. The 2ARSA pattern is the transitive pattern of the causative construction that expresses the agent in the subject position and the theme in the direct object position.

- For the instrument subject alternation, the 1ARSA pattern is the transitive or intransitive pattern of the instrument subject construction. The 2ARSA pattern type encompasses the *with* and *using* patterns of the agent subject construction, both of which express the instrument in the respective phrase added for it. Note that the 2ARSA pattern may have an additional syntactic argument, a direct object, which is not "alternation-relevant" in the same sense as the argument that expresses an instrument in sentences with alternating verbs.

- For the together reciprocal alternation (intransitive), the 1ARSA pattern is the intransitive pattern of the *together* construction, which expresses themes in the subject position only. The 2ARSA pattern type comprises the patterns of the different realizations of the prepositional construction, which expresses themes both in the subject position and in the prepositional phrase.

The 1ARSA and 2ARSA pattern types for each alternation are illustrated in (59) through (61).

(59) Causative-inchoative alternation:

a. *The glass* empties. — *1ARSA pattern*

b. *Bildad* empties *the glass*. — *2ARSA pattern*

(60) Instrument subject alternation:

a. *Starbuck's lance* scratched their backs. — *1ARSA pattern*

b. *Starbuck* scratched their backs with *his lance*. — *2ARSA pattern*

c. *Starbuck* scratched their backs using *his lance*. — *2ARSA pattern*

(61) Together reciprocal alternation (intransitive):

a. *The oil and the butter* blend together. *1ARSA pattern*

b. *The oil* clings to *the butter*. *2ARSA pattern*

c. *The oil* blends into *the butter*. *2ARSA pattern*

d. *The oil* merges with *the butter*. *2ARSA pattern*

The terms *1ARSA* and *2ARSA* do not express a claim as to which of the various patterns involved in an alternation are the underlying forms or derived forms, or which of them is more simple or more complex. Goldberg (Goldberg 1995: 7, Goldberg 2013: 15) notes that Construction Grammar does not posit underlying syntactic or semantic forms and is not transformational. Instead, the terms are merely introduced here as a way to distinguish the different pattern types in a way that generalizes across the different alternations and their specific realizations of 1ARSA and 2ARSA patterns.

If a verb is observed both in the 1ARSA pattern type and the 2ARSA pattern type of an alternation, this is not sufficient evidence to label the verb as participating in that alternation. Verbs should only be classified as alternating verbs if their attestations in each pattern type are likely to be instances of the alternation-specific constructions, and not of unrelated, superficially-similar constructions.

Additionally, alternation-independent terms are needed to refer to the different alternation-relevant syntactic slots in the 1ARSA and 2ARSA patterns for each alternation. If a verb participates in a given alternation, certain slots in that alternation's 1ARSA and 2ARSA pattern types are expected to express the same semantic argument; for verbs that do not participate in this alternation, those slots are expected to express different semantic arguments.

Each of the alternations under investigation here has three relevant slots that need to be distinguished. For the causative-inchoative alternation, the three slots are the subject slot of the 1ARSA pattern, the subject slot of the 2ARSA pattern, and the direct object slot of the 2ARSA pattern. For the instrument subject alternation, the three slots are the subject slot of the 1ARSA pattern, the subject slot of the 2ARSA pattern, and the slot headed by *with* or *using* in the 2ARSA pattern. For the together reciprocal alternation (intransitive), the different slots are the subject slot of the 1ARSA pattern (if there are multiple subjects, connected with a conjunction, they are expected to share a semantic role, so it makes sense to summarize all slot fillers in the whole subject phrase as one category); the subject slot of the 2ARSA pattern; and the prepositional object slot in the 2ARSA pattern.

The subject slot of the 1ARSA pattern type for each alternation will here be referred to as the *external argument of the 1ARSA pattern* (*EA1*). The subject slot of the 2ARSA pattern type for each alternation will be called the *external argument of the 2ARSA pattern* (*EA2*). The third relevant slot for each alternation will be referred to as the (*alternation-relevant*) *internal argument of the 2ARSA pattern* (*IA2*). For the causative-inchoative alternation, the IA2 slot is the syntactic object position in the 2ARSA pattern. For the other alternations, this slot is located inside a prepositional phrase or an adverbial phrase in the 2ARSA pattern. The 2ARSA pattern type for the instrument subject alternation can, and often does, involve a direct object; however, in this thesis, the term *IA2* will only ever refer to the argument slot inside the *with* or *using* phrase, as the focus here is not on the set of possible undergoers of verbs in this alternation, but on possible instruments.

Verbs that participate in one of the alternations under investigation here will assign different semantic roles to the EA1 slot than verbs that do not participate in the alternation. The IA2 slot will express the same semantic role as the EA1 slot if the verb alternates. For verbs that do not participate in the alternation, the EA1 position will typically express the same semantic role as the EA2 position. This effect is illustrated by the sentences in (62) and (63), where different semantic roles are expressed in the EA1 slot depending on whether the verb participates in the instrument subject alternation or not.

(62) ISA-alternating verb:

a. *His lance* tickled their backs. (*1ARSA*)

b. *He* tickled their backs with his lance. (*2ARSA*)

c. *He* tickled their backs using his lance. (*2ARSA*)

(63) Non-alternating verb:

a. *He* eats the steak. (*1ARSA*)

b. *He* eats the steak with a fork. (*2ARSA*)

c. *He* eats the steak using a fork. (*2ARSA*)

The patterns involved in the instrument subject alternation pose a particular challenge, as it is typically not mandatory to explicitly mention instruments, especially when the instrument that is being used in the situation described by the sentence is typical for the event described by the verb. A sentence like *He tickled their backs* is plausible and acceptable. The verb *tickle* participates in the alternation, but the sentence instantiates the 1ARSA pattern without expressing an instrument in the subject position. In the previous part of this thesis, the distinction between such sentences and sentences like (62a) was made based on a frame type hierarchy with type constraints for individual arguments, such that entities belonging to certain types were unavailable for the instrument or agent role. For the classification of alternating and non-alternating verbs, no such type resource will be available. The assumption that the EA1 argument always expresses an instrument when the verb participates in the alternation is a simplification. The results of the classification will give an indication of the extent to which this phenomenon impacts the classifiers.

Verbs that participate in the causative-inchoative alternation will typically express a theme in their EA1 slot and in their IA2 slot, and an agent or a cause in their EA2 slot. Verbs that participate in the instrument subject alternation can express an instrument in their EA1 slot and in their IA2 slot, and an agent in their EA2 slot. Verbs that participate in the together reciprocal alternation (intransitive) will express themes in their EA1 slot, their IA2 slot, and their EA2 slot.

Most of the classifier features presented in the following chapter rely on the different associations between relevant syntactic positions and semantic roles to distinguish alternating verbs from non-alternating ones. Different sets of features will be defined whose aim is to encode which of the relevant syntactic slots for a given verb are most likely to be linked to the same semantic arguments. In other words, the features attempt to capture whether the attested instances of each verb in the relevant patterns are predominantly instances of the alternation-specific constructions, or of alternative, superficially-similar constructions.

## 9.5 corpus data for the classifier training sets

This section describes the corpus on which the classification experiments will be based. Only a subset of the sentences in the original corpus will actually be seen by the classifiers during training, since the classification features are exclusively concerned with attestations of alternation candidates in the various pattern types for each alternation. Therefore, the corpus will be filtered down to create a subcorpus for each alternation and pattern type, from which attestations of each alternation candidate will be collected to determine the candidate's feature values for the classification.

### 9.5.1 *The ENCOW corpus of English web text*

As mentioned previously, the classifiers will be trained on corpus data containing instances of verbs that participate in each alternation and instances of verbs that do not participate in the alternations, but nevertheless appear in the relevant syntactic patterns. The source for these alternation-specific subcorpora will be the ENCOW corpus (Schäfer & Bildhauer 2012, Schäfer 2015). That corpus contains 9,578,828,861 tokens of English web text, crawled in 2012 and 2014 and parsed between 2015 and 2018, available for download in a sentence-shuffled version. The dependency trees for the corpus were generated with MaltParser (Nivre 2003). More information about the corpus is available on its website at `https://web.archive.org/web/20210924132944/https://corporafromtheweb.org/encow16/`.[2]

Since the corpus consists of web text, it contains a certain amount of unedited, user-generated text. This means that there are some expected variations in style and spelling that will impact the experimental results. Additionally, since the corpus was parsed between 2015 and 2018, the quality of the dependency trees generated by MaltParser is below the current state of the art for dependency parsing on English text. While neither the MaltParser website nor the ENCOW website provide detailed information about the accuracy of the specific model that was used to process the original ENCOW data, the most recent evaluation of MaltParser before the publication of ENCOW (Nivre, Hall & Nilsson 2006) gives a labeled attachment score for English of 86.3, evaluated on the Penn Treebank. For the noisier texts that are included in ENCOW, the parser is likely to have been less accurate than that.

The dependency structure of sentences in the corpus plays an essential role in the classification process, since instances of syntactic patterns will be collected based on the dependency relations that characterize each alternation-specific construction. It is therefore crucial to have access to more reliable dependency information. This is why the original dependency labels provided with ENCOW will be double-checked by running a newer parser, udpipe (Straka, Hajič & Straková 2016, Straka 2018)[3], on the sentences collected in the initial filtering step. Only sentences that fulfill the syntactic conditions for a pattern according to both parsers are then used for the classification.

2 At the time of conducting the experiments reported here, the original website was still available online. Since then, the website has changed to only allow access to the corpora for logged-in users of institutional accounts.

3 Here: version 1.2.0.3, as sourced from `https://pypi.org/project/ufal.udpipe/`. The model used for the parser in these experiments is `english-ewt-ud-2.5-191206.udpipe`, available for download at `https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-3131`.

Although udpipe and the English model for it are newer than the parser and the model that were used to create the original ENCOW annotations, the labeled attachment score for English is reported on the udpipe website to also be around 86. The strategy of using only sentences for which both parsers yield the desired dependency structures is meant to exclude as many misparsed sentences as possible from the subcorpora. This strategy reduces the size of the subcorpora, but at the same time improves the reliability of the remaining data. Earlier attempts using a heuristic with manually-defined rules for the inclusion or exclusion of sentences from ENCOW according to their MaltParser labels were not pursued further because the rules were highly specific to the types of errors introduced by MaltParser, which would not generalize well to other, newer parsers.

While the strategy of using only sentences that fit the desired syntactic patterns according to both MaltParser and udpipe generates more reliable corpora for the experiments, the two tools do not always assign the same lemmas to the tokens in each sentence. This is problematic, because a successful lemmatization step is necessary in order to be able to make predictions about verbs and argument slot fillers on the lemma level. Straka, Hajič & Straková (2016) implement the morphological analysis in udpipe with the help of a guesser that applies a lemma rule which they describe as

> [...] the shortest formula for generating a lemma from a given form, using any combination of "remove a specific prefix", "remove a specific suffix", "append a prefix" and "append a suffix" operations. (Straka, Hajič & Straková 2016: 4293)

They evaluate this lemmatization algorithm on the English portion of the Universal Dependencies 1.2 treebanks with an accuracy of 99.4; the method used to achieve this score is described as

> [...] the success rate, i.e., the ratio of words, for which the analyser produces among others the correct analysis, for the morphological analyser. (Straka, Hajič & Straková 2016: 4294)

The lemmatization in the downloadable ENCOW data was produced using TreeTagger (Schmid 1994, 1999), with the Penn Treebank tag set (Santorini 1990) and the morphological analyzer from Karp et al. (1992). The ENCOW website does not specify which version of the underlying models was used when the corpus was created, or what the accuracy of the lemmatizer was at the time. In the course of setting up the experiments for this thesis part, a comparison of verb lemma lists from both sources – udpipe and the ENCOW lemmas – showed that the rule-based lemmatization from udpipe tends to overgeneralize; for instance, the verb form *revealed* is mapped to the nonexistent lemma *revealead*, presumably because the suffix *-led* resembles a known verb from the training data. Other examples of incorrect lemmatization from udpipe include *submitt* as the lemma for the word form *submitted*, *tabl* as the lemma for the word form *tabled*, and *suppli* as the lemma for the word form *supplied*.

Using inconsistent lemmas for the experiments conducted here would mean losing some generalization power and introducing some unwanted noise, since forms of the same verb would not necessarily be mapped to the same verb lemma, and the same holds for the lemmas of the observed argument slot fillers for each verb. The features for the classification task should be robust to inflection and other morphological properties, so a successful mapping from observed tokens to their underlying lemmas is important to avoid losing relevant information. For the purposes of this thesis, the parse trees for each sentence in the subcorpora will be sourced from the combination of MaltParser and udpipe output, while the lemmas will be taken from the output of MaltParser. This decision has practical benefits, since the more accurate lemmatization by MaltParser is already part of the downloadable original ENCOW corpus. In future work, improvements in the areas of dependency parsing and lemmatization may yield better classification results based on the same input data; for now, the lemmas and dependency trees will be sourced from this combination of both tools, MaltParser and udpipe.

### 9.5.2 *Collecting alternation-specific subcorpora of ENCOW*

Each of the three alternations under investigation here involves a set of constructions that can be described in syntactic terms, which is necessary for the step of collecting alternation-specific and pattern-specific subcorpora on which to train the classifiers. The dependency trees given in (64) through (66) describe the relevant dependency relations in each alternation's 1ARSA and 2ARSA pattern. The alternation-relevant syntactic arguments are framed and labeled as EA1, EA2 or IA2. Arcs and labels in gray denote relations that are not required in order for a sentence to instantiate the current pattern. For instance, sentences instantiating the patterns associated with the instrument subject alternation can be intransitive or transitive.

The data available for download on the ENCOW website is annotated with the 2008 version of Stanford Typed Dependencies, as documented here: `https://downloads.cs.stanford.edu/nlp/software/dependencies_manual.pdf`. The descriptions for the alternation-specific patterns below follow the Universal Dependencies 2.5 standard from de Marneffe, MacCartney & Manning (2006), de Marneffe & Manning (2008), de Marneffe et al. (2014), which is also the format returned by udpipe.

(64) Dependency patterns for the causative-inchoative alternation:

a. 1ARSA pattern:

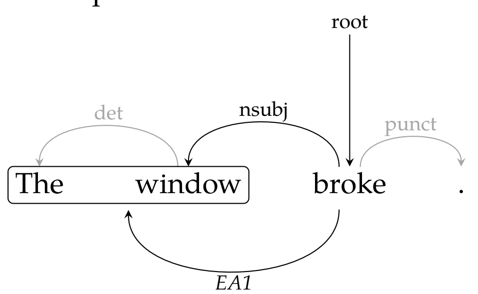


b. 2ARSA pattern:

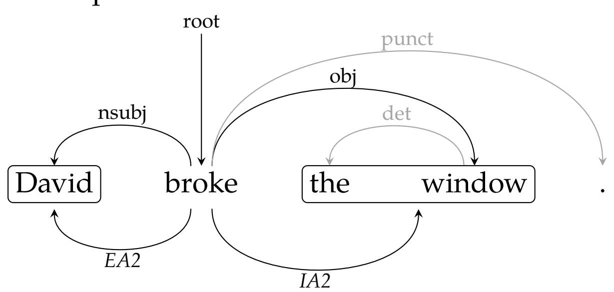

(65) Dependency patterns for the instrument subject alternation:

a. 1ARSA pattern:

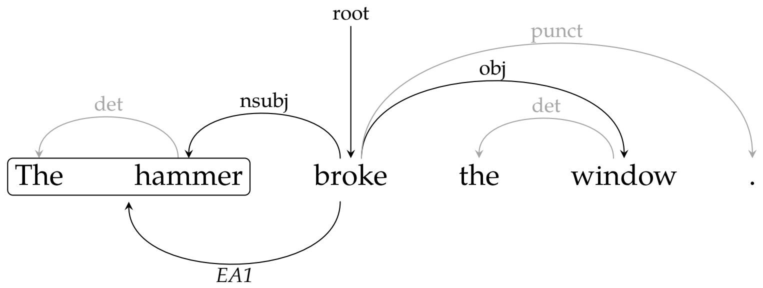


b. 2ARSA pattern:

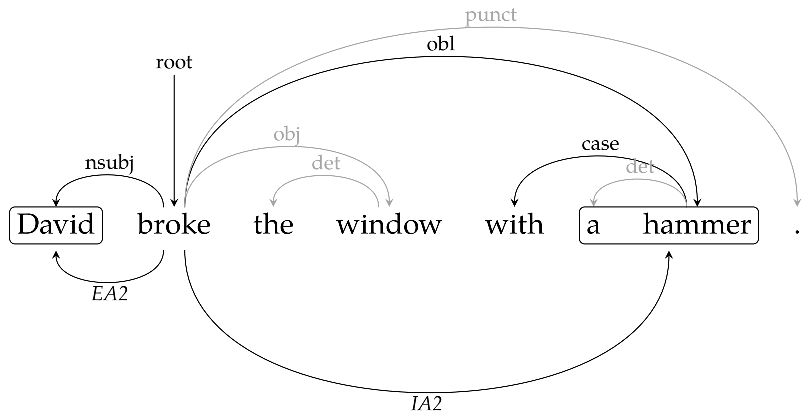


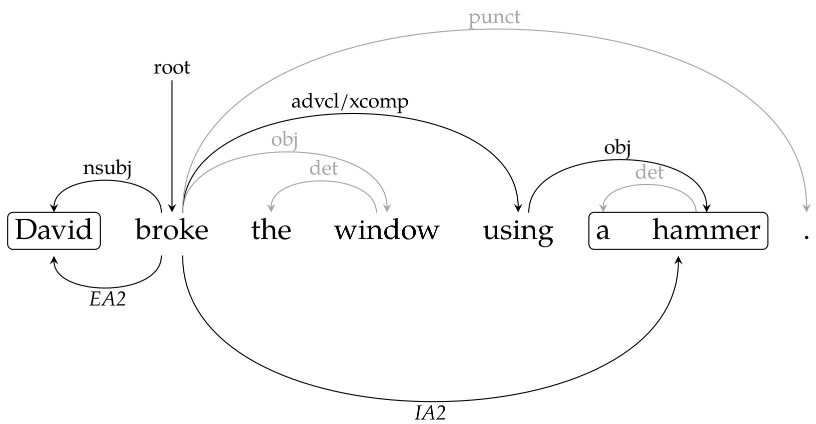

(66) Dependency patterns for the together reciprocal alternation (intransitive):

a. 1ARSA pattern:

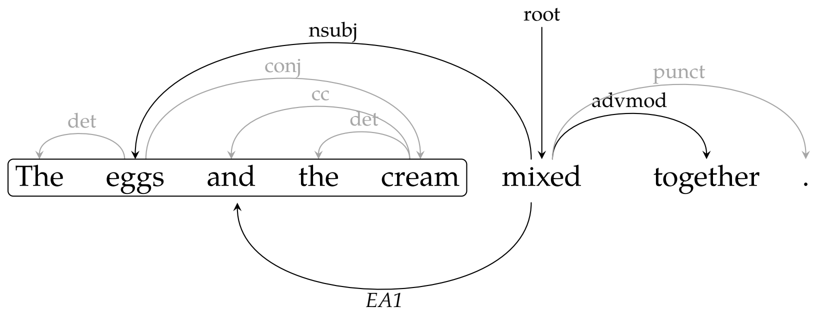


b. 2ARSA pattern:

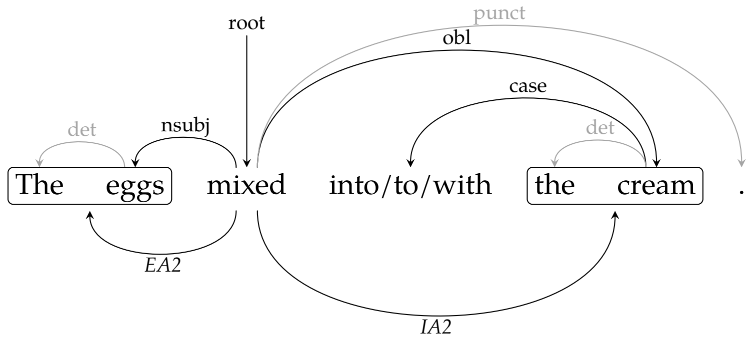


These patterns are the basis for the collection of alternation-specific subcorpora as a source for the classification features. A sentence is only included in one of the subcorpora if its top-level dependencies (dependencies that have the sentence's `root` element as their head) correspond to the relations given in the relevant pattern description. Tokens with a `punct` relation to the `root` element are disregarded. The order of the dependencies in a sentence is not taken into account.

In the next step, all subcorpora are re-parsed with udpipe, due to the known quality issues with the dependency annotations that are provided with the original corpus. Re-parsing all sentences in ENCOW would be computationally expensive and time-consuming, so the parsing efforts for this thesis are limited to the alternation-specific subcorpora. Once all sentences are parsed, another filtering step is conducted in which each sentence from each subcorpus is checked for the conditions for that subcorpus as well as all other subcorpora. Each sentence that fulfills the conditions of one or more alternation patterns is included in the subcorpora for the relevant alternation patterns.

At this stage, some additional checks are performed on each subcorpus to ensure that each sentence is usable for the experiments to be implemented in the following. This mainly concerns the perplexity-based features, in which each sentence from the corpora is compared to one or more alternate versions of that sentence. The alternate versions will be generated with a rule-based script. Sentences whose syntactic properties make this reordering step impossible or lead to systematically odd resulting sentences are discarded. For more on this, see Section 10.5.

Reducing each subcorpus to only sentences that can be reordered without issues is yet another step that limits the amount of available data for the experiments. While it would be possible to reduce the subcorpora only for the perplexity-based classifier, and use the larger corpora for the other approaches, this would make a comparison of

Table 4: Sizes of final subcorpora.

| Alternation | Pattern | # Sentences | # Verb types |
|---|---|---|---|
| CIA | 1ARSA | 235,496 | 4,006 |
| | 2ARSA | 1,073,240 | 6,385 |
| ISA | 1ARSA | 2,682,746 | 8,063 |
| | 2ARSA (with) | 182,631 | 3,787 |
| | 2ARSA (using) | 3,157 | 594 |
| TRAI | 1ARSA | 1,149 | 221 |
| | 2ARSA (into) | 16,226 | 860 |
| | 2ARSA (to) | 88,910 | 1,438 |
| | 2ARSA (with) | 67,399 | 1,689 |

the performance of the different approaches more difficult and less informative. Since one of the goals of this chapter is to determine which approaches are most well-suited for which alternations, it is more appropriate to use one fixed set of subcorpora that are usable for all approaches, instead of reducing the corpora only for a subset of the classification features.

These considerations and filtering steps lead to the final reduction of subcorpora. After removing sentences whose udpipe-generated parse trees yield results that are not congruent with the desired syntactic patterns, and then removing sentences that cannot be reordered, the final corpora for the experiments have the sizes given in Table 4.

## 9.6 VERBS UNDER INVESTIGATION IN THIS THESIS PART

Each alternation-specific classifier will be trained based on a positive verb set (with alternating verbs) and a negative verb set (with verbs that do not participate in that alternation). The source for the positive verb sets will be VerbNet 3.4 (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007). Some alternations are explicitly marked in VerbNet, while others are not marked directly but can be derived from other information that is given in VerbNet, e.g. available subcategorization frames or semantic roles for a verb or verb class.

For the instrument subject alternation and the together reciprocal alternation (intransitive), VerbNet classes that participate in one of these alternations are marked with the name of the alternation, as part of the `description` field in the data provided for the class. For the causative-inchoative alternation, alternation participation is not marked on the level of the verb class, but on the level of individual frames that belong to the class. A VerbNet class is regarded as participating in the causative-inchoative alternation if at least one of its frames is marked as `(basic) transitive; causative (variant)` and at least one of its frames is marked as `intransitive; inchoative`. Extracting alternating verbs for all three alternations from VerbNet based on these indicators results in the initial alternation-specific verb set sizes presented in Table 5 on page 217. Note that no word sense disambiguation step is involved: if a verb lemma is a member of multiple VerbNet classes that participate in an alternation, then that

Table 5: Number of verbs participating in each of the three alternations under investigation, according to VerbNet.

| Alternation | Number of participating verbs according to VerbNet |
|---|---|
| CIA | 453 |
| ISA | 543 |
| TRAI | 50 |

lemma is only counted once. Verbs that are members of alternating classes as well as non-alternating classes are only counted as alternating verbs.

### 9.6.1 *Creating positive and negative verb sets*

As noted by Lapata (1999: 397), there may be a mismatch between the verbs reported by Levin (1993) as participants in an alternation, and the verbs attested in corpus data with the relevant patterns. Using data from the BNC corpus, Lapata searches for the syntactic frames involved in the dative alternation and the benefactive alternation, and identifies, on the one hand, verbs that are listed in Levin (1993) as alternating but are not attested in the relevant frames in the corpus, and on the other hand, verbs that do participate in one of the alternations according to corpus data, even though they are not listed as such in Levin (1993).

While VerbNet covers more verbs than Levin (1993), and has been repeatedly updated over several years to incorporate additional verbs or to add alternations to verb classes, a certain mismatch between the corpus data and VerbNet is still expected. As mentioned previously, absence of evidence of felicity is not evidence of infelicity. It is therefore undesirable to automatically label a verb as non-alternating if it is not attested in all relevant syntactic patterns. The classifiers will be trained on subcorpora that exclusively contain sentences in the relevant syntactic patterns, and will not be able to make any informed claims about verbs that do not appear in these patterns. For this reason, the original positive verb sets from VerbNet for each alternation will be reduced to only those verbs that occur in the corpus at least once with each of the pattern types associated with the alternation's constructions.

A challenge to be considered here is that VerbNet only marks verb classes that participate in an alternation, but the absence of an alternation flag in a verb class does not necessarily mean that the verb class and the alternation are incompatible. It could instead mean that a connection between the verb class and the alternation was simply never considered, or that the alternation is marked in other ways than an explicit flag (as is the case for verbs participating in the causative-inchoative alternation), or that the participation of this verb class in this alternation is contested. For the experiments conducted in this thesis, it is necessary to have access to a negative set as well as a positive set of alternating verbs.

There are two purposes to the negative sets that will be used here. The first one relates to casting the problem as a binary classification task. The second purpose is to collect a set of possible additions to the positive set for each alternation, beyond the positive verbs that are found in VerbNet. Section 9.1 hinted at the fact that more verbs might participate in an alternation than are marked as such in VerbNet; if this is the case, any set of presumed negative instances may contain a number of good additional candidates for the alternation. Section 10.7 in the next chapter will go into more detail about this idea.

Since this part of the thesis is concerned with classification procedures that aim to distinguish verbs in a given alternation from verbs that do not participate in that alternation, the strategy for the selection of negative training and test instances has a direct impact on the success and validity of the results reported for these classifiers. For instance, if the negative set exclusively contains verbs that never appear in the relevant syntactic patterns, then observing an individual verb in the patterns would be a completely reliable indicator for alternation participation. If, on the other hand, the negative verbs can appear in the same syntactic patterns as the positive ones, the classifiers need to rely on additional features encoding the differences between alternating and non-alternating verbs.

Thus, the strategy for collecting negative verbs for this type of classification task is a crucial aspect of the experimental setup. In Baroni & Lenci (2009), the negative verbs for a classification task focusing on the causative-inchoative alternation are sourced from the chapter in Levin (1993) about the alternation. There, Levin gives examples of verbs that do not participate in the alternation and for which at least one of the syntactic patterns associated with the alternation is invalid. For instance, the class of GIVE verbs is mentioned as an example of a class that does not participate in the alternation, as illustrated in (67):

(67) a. They gave the bicycle to me.
  b. * The bicycle gave to me.

Baroni & Lenci (2009) base their classification task on these negative verbs and report flawless classification accuracy.

In a different approach to a similar classification task, Merlo & Stevenson (2001) aim to distinguish verbs belonging to three different optionally-intransitive classes. Every candidate verb in their setup can appear in a transitive or intransitive context, and linguistically motivated features are used to identify the most likely class for each candidate. Their evaluation shows that the members of different classes cannot be distinguished completely reliably in this setup; in other words, their candidate sets are more challenging to label correctly. In this thesis, the selection of alternation candidates will follow the example of Merlo & Stevenson (2001) rather than that of Baroni & Lenci (2009).

Another important property of the positive and negative verb sets for the experiments is the overall frequency of each verb in a corpus. If, for instance, non-alternating verbs appear more frequently in a subcorpus than their alternating counterparts, this imbalance between the sets will impact the way the classifiers assign labels to verbs. Therefore, care should go into ensuring that verbs in the negative set have a frequency in the corpus that is comparable to that of the verbs in the positive set. Practically, this will be achieved here by creating the negative verb set for each alternation as follows: for each verb in the positive set, first determine its overall frequency across the alternation-specific subcorpora. Then, from the set of all verbs observed in these subcorpora, select one as a negative counterpart to the current positive verb, such that the

Table 6: Number of verbs in the positive and negative set for each alternation, after reducing the original positive set to only verbs that are attested in each relevant pattern at least once.

| Alternation | Positive/negative verb set size |
|---|---|
| CIA | 269 |
| ISA | 291 |
| TRAI | 13 |

overall frequency of the negative verb in the subcorpora is the same or close to the frequency of the positive verb.

VerbNet and Levin (1993) treat alternation availability as a property on the lexical level, not on the level of individual instances of verb usage. Therefore, all classification setups that are described in this chapter approach the task as follows: given a verb lemma, e.g. *decide*, determine whether that verb participates in a specific alternation, e.g. the causative-inchoative alternation, based on a set of linguistically motivated features, whose values are based on observed instances of the verb in the ENCOW subcorpus containing all pattern types for that alternation. The features are slightly adapted for each alternation, but essentially represent the same type of information. They will be motivated, defined and evaluated in Chapter 10.

Based on these considerations and the alternation-specific subcorpora collected previously, the final positive and negative verb sets for each alternation can now be collected. Table 6 presents the verb set sizes after reducing the positive verb sets to only those that are attested at least once in each relevant pattern type for their alternation, and selecting the corresponding number of negative verbs from the remaining verb lemmas observed in the corpora.

It is not surprising that the size of the verb sets shrinks so noticeably when only attested verbs are taken into account, even though each subcorpus contains a large number of attested verb types (see Table 4 on page 216.). Since the criteria for the creation of the subcorpora are purely syntactic in nature, each subcorpus is expected to contain a large number of sentences that only superficially seem to instantiate one of the relevant constructions. Furthermore, some of the conditions on the subcorpora are not very restrictive; for instance, the subcorpus for the 1ARSA pattern of the instrument subject alternation comprises all sentences from ENCOW that have an `nsubj` element as a top-level dependency under the root verb. This filter is so unspecific that the large number of verbs observed in this syntactic pattern is no reliable indicator for the prevalence of actual alternating verbs in the data.

Out of the 50 verbs originally provided by VerbNet as positive instances for the together reciprocal alternation (intransitive), only 13 are attested in both the 1ARSA and the 2ARSA pattern. If attestation in at least one of the patterns was sufficient for the inclusion of a verb in the candidate pool, the process would leave 44 verbs in the set instead of reducing it to 13. However, since attestation in all relevant syntactic patterns is the basic condition that all positive and negative candidates for an alternation must exhibit for the classification task implemented here, the experiments will be focusing on those 13 positive and 13 negative candidates that are selected based on attestation data.

Table 7: Number of verbs in the frequency-filtered positive and negative set for each alternation, such that only verbs are included that appear at least 20 times in the relevant subcorpora.

| Alternation | Positive/negative verb set size |
|---|---|
| CIA | 178 |
| ISA | 226 |
| TRAI | 10 |

In order to classify an alternation candidate based on its occurrences in the subcorpora, those instances have to be representative of the possible contexts for that verb. For verbs that are attested very infrequently, this can pose problems, for instance if the single observed occurrence of a verb in one of the patterns of an alternation is actually a misparsed sentence. A second group of verb sets for the alternations is therefore defined under the condition that each candidate must appear at least a total of 20 times in the relevant corpora. The sizes of the verb sets under this condition are given in Table 7. The classifiers will be evaluated on both conditions, and the results will be discussed and interpreted in the respective sections in the next chapter.

Based on these verb sets and the subcorpora comprising their attestations in the alternation-specific patterns, the classifiers can now be trained. The different features and the performance of each classifier based on each feature set will be discussed in the following chapter.

# 10 ALTERNATING VERB CLASSIFICATION

## 10.1 INTRODUCTION

This chapter is primarily concerned with the implementation of three types of features for the classification of candidate verbs for the three alternations under investigation. The feature sets are evaluated in comparison to a frequency-based baseline (Section 10.2) that only takes the ratio of 1ARSA pattern attestations and 2ARSA pattern attestations of each verb into account. Section 10.6.2 interprets the classification results with respect to the hypotheses formulated in Section 9.1 in the previous chapter. An attempt to use the classifier results to extend the set of positive verbs for an alternation is described in Section 10.7.

The first set of features (Section 10.3) will be based on observed lemmas in the different argument slots of the syntactic patterns associated with each alternation's constructions. At least one semantic argument can be expressed in different slots of the 1ARSA and 2ARSA pattern type for each alternation. If a verb participates in the alternation, the observed lemmas in one of these slots should be similar to the observed lemmas in the corresponding slot in the other pattern.

The second set of features (Section 10.4) will be based on vector representations of the observed lemmas in the different argument slots of the syntactic patterns associated with each alternation. Different slots may express the same semantic argument and thus impose the same selectional preferences on their slot fillers, even if little or no lemma overlap is observed in these slots. The vector-based features aim to approximate the semantic category of slot fillers to estimate how likely it is that two slots in different patterns contain the same semantic argument.

The third set of features (Section 10.5) will be based on a rule-based extension of each pattern-specific subcorpus. For each sentence instantiating one pattern of an alternation, an alternate sentence is generated that instantiates the contrasting pattern. Arguments are placed in the positions they would take if the verb alternates. The assumption is that the resulting sentences are more acceptable if the verb participates in the alternation, and less acceptable if it does not alternate. Scores for each pair of sentences are retrieved from a large neural language model in order to estimate the acceptability delta for the alternating and non-alternating verbs.

Finally, a hybrid approach is presented that uses the classifier output from the different feature types and introduces a manual annotation step. The goal is to identify verbs that seem to be incorrectly labeled as participating in an alternation, but which may in fact be good candidates for the positive set for that alternation. This approach is iterative, as the results from the annotation step can directly be used to train and run the classifier again in order to trigger the next annotation round.

## 10.2 FREQUENCY-BASED BASELINE

### 10.2.1 *Introduction*

The success of the features developed in this chapter for the classification of verbs in the three alternations will be compared to a frequency-based baseline that will be described in this section. This baseline uses as input to the classifier only information on the number of attestations of each verb in each pattern type of an alternation. The features that will be implemented in the following sections are expected to outperform this baseline if they successfully capture meaningful semantic properties of the verbs and their arguments.

Existing work on alternation identification and related verb classification tasks historically often used observed syntactic subcategorization frames (SCFs) as main indicators of alternation participation or membership in some type of verb class. In those approaches, the pure presence or absence or the frequency of the different relevant syntactic patterns in corpus data served as features for the verb classification tasks. For more on these types of work, see Section 10.3.2.

The success of a purely SCF-based approach to alternating verb identification depends, among other things, on the criteria guiding the creation of the set of candidate verbs. For most alternations, Levin (1993) contrasts verbs that participate in the alternation with others that allow one of the patterns, but not the other, or verbs that are semantically similar to verbs in the alternation but cannot be used in the relevant patterns. This is a useful contrast to make when explaining how the behavior of alternating verbs systematically differs from that of verbs outside the alternation. In the context of an attempt to distinguish alternating verbs from non-alternating verbs, using these negative examples is likely to lead to deceptively good results, since the negative verbs were specifically chosen by Levin by virtue of them not being expected to exhibit the required behavior.

In this thesis, due to the selection strategy for the positive and negative verb sets for this task, the pure presence or absence of attestations of a verb in a set of patterns cannot be used as a reliable indicator for alternation participation. All verbs in both sets are certain to occur with each relevant syntactic pattern at least once in the corpus, which would compel the classifiers to always assign the positive class to every single candidate verb.

### 10.2.2 *Related work on verb classification based on SCF frequencies*

Syntactic subcategorization frames are an essential feature in many less recent studies on verb clustering or classification tasks, such as Schulte im Walde (2000, 2006), Sun, McCarthy & Korhonen (2013).

Merlo & Stevenson (2001) classify verbs that share their basic syntactic frames but belong to three different classes, such that the shared syntactic patterns are associated with different thematic role assignment patterns. In their experimental setup, they are concerned with three verb classes that are optionally intransitive for different reasons, which means they all license an intransitive subcategorization frame. This illustrates that categorizing verbs purely based on the syntactic patterns in which they appear is necessarily an inexact strategy.

At the same time, the distributions of subcategorization frames for verbs in a corpus can still be informative as a general description of each verb's behavior. Contributions like the German SCF lexicon created by Schulte im Walde (2002) thus provide useful resources for determining alternation *candidates*.

Lapata & Brew (1999) approach the task of alternating verb identification by determining the joint probability of a verb and a syntactic subcategorization frame based on corpus frequencies. This approach requires knowledge of the number of occurrences of each verb *outside* the relevant SCFs as well as the number of occurrences *in* one of the relevant SCFs. At the same time, Lapata & Brew do not approximate selectional preferences, and their classification is purely based on the distribution of verbs in the various relevant syntactic patterns.

In the scope of this thesis, the approach by Lapata & Brew is not applicable. Firstly, the focus of the present work is on determining to what extent selectional restrictions of alternating and non-alternating verbs can be approximated to aid classification; in other words, the slot fillers for the candidate verbs' arguments will be the main focus of the classifier features. Secondly, because of the focus on selectional restriction approximation, the subcorpora for the experiments presented here consist exclusively of sentences instantiating the relevant syntactic patterns for each alternation. Thus, no information about the distribution of candidate verbs outside these patterns is available.

### 10.2.3 *Method*

The naive frequency-based feature that will be implemented in this section is meant to determine the baseline classification accuracy for the alternation identification task. Features that approximate selectional preferences in some way are expected to outperform the feature based on the simple observation of syntactic pattern frequencies.

Note that in this thesis, the selection of non-alternating candidate verbs ("negative verbs") for each alternation already follows basic frequency conditions with respect to the relevant syntactic patterns: for each known positive verb, a negative verb is selected which has roughly the same number of observed instances across the two or more syntactic patterns that are relevant to the alternation. This means that a simple check as to whether a verb appears in all relevant patterns or not would be expected to perform at exactly 50% accuracy on a verb set that consists of a balanced number of alternating and non-alternating verbs.

feature 00-01: *Frequency ratio between 1ARSA and 2ARSA pattern.* This feature encodes the relative frequency of the 1ARSA and 2ARSA pattern type for each alternation. This will make possible trends in pattern distributions visible: do alternating verbs instantiate their different patterns at more balanced frequencies than non-alternating verbs? This might be the case if non-alternating verbs appear in one of the patterns only coincidentally, whereas the pattern selection of alternating verbs might be expected to be more systematic. Possible explanations for coincidental instances of non-alternating verbs in an alternation-relevant pattern include misparsed sentences, idiosyncratic usages, usage errors, or some other phenomenon that is not associated with the alternation but may lead to similar surface behavior in individual cases.

To make the feature values for infrequent verbs and more frequent verbs comparable, this feature is defined as the ratio of 1ARSA pattern instances to all attestations of the

Table 8: Scores (accuracy, precision, recall, $F_1$) for the frequency-based baseline feature.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 00-baseline | 57.56 | 57.35 | 60.74 | 58.63 |
| ISA | 00-baseline | 52.70 | 51.90 | 54.97 | 53.51 |
| TRAI | 00-baseline | 30.33 | 8.67 | 23.00 | 10.17 |

current candidate verb. This leads to feature values between 0 and 1, although 0 and 1 will never occur due to the data containing only candidates that appear at least once in each pattern.

For alternations that have a 1ARSA or 2ARSA pattern with more than one possible surface form (the instrument subject alternation and the together reciprocal alternation (intransitive)), the frequencies for all patterns belonging to the same pattern type will be summed up to determine the value for this feature. For instance, in order to determine the frequency of the 2ARSA pattern for the instrument subject alternation, the frequency of each verb in the *with* subcorpus is added to the frequency of the same verb in the *using* subcorpus. For the instrument subject alternation, this makes more sense than calculating the average frequency of these two separate variants of the 2ARSA pattern, since both of them contrast with the 1ARSA pattern in the same essential way.

For the together reciprocal alternation (intransitive), Levin (1993: 64) specifies particular prepositions for the verbs listed as participating in this alternation; however, in order to facilitate a comparison between the classification performance on the three different alternations discussed here, the classifier will not attempt to predict the specific preposition with which a candidate verb may participate in the together reciprocal alternation (intransitive). Since verbs can participate in this alternation with more than one possible preposition (for example, according to Levin, *blend* allows the prepositions *with* and *into*), and since the corpora for this alternation are smaller than the ones for the other alternations, summarizing the different variants of the prepositional pattern type is expected to lead to better results than attempting to make more detailed predictions based on the sparse data that is available.

### 10.2.4 *Results*

The results of the frequency-based baseline feature on the full verb sets for each of the three alternations are shown in Table 8. The feature set `00-baseline` contains only feature 00-01, which expresses the relative frequency of a verb in an alternation's 1ARSA pattern in relation to the total frequency of that verb in the alternation-specific subcorpus.

The frequency-ratio baseline performs at a wide range of scores for the different alternations. For the causative-inchoative alternation, it is relatively strong (for such a naive feature) with an $F_1$ score of 58.63 and an accuracy of 57.56, while for the instrument subject alternation, it is weaker with an $F_1$ score of 53.51 and an accuracy of 52.70, and for the together reciprocal alternation (intransitive), it performs poorly with an $F_1$ score of 10.17 and an accuracy of 30.33.

Table 9: Scores (accuracy, precision, recall, $F_1$) for the frequency-based baseline feature, with a minimum alternation candidate frequency of 20.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 00-baseline | 63.79 | 64.70 | 63.10 | 63.21 |
| ISA | 00-baseline | 56.84 | 55.90 | 63.65 | 59.13 |
| TRAI | 00-baseline | 39.50 | 5.00 | 11.00 | 6.67 |

For the together reciprocal alternation (intransitive), the subcorpora are extremely unbalanced, with just over 1,100 sentences in the *together* pattern and tens of thousands of sentences in each of the prepositional patterns. As a result, the 1ARSA pattern is much less frequent than the 2ARSA pattern for the majority of verbs in the test set, whether they participate in the alternation or not. The fact that counting the occurrences of the different patterns leads to poor classification performance for this alternation points out the limits of this naive frequency-based baseline. While the sizes of the subcorpora for the 1ARSA and 2ARSA pattern are also unbalanced for the other two alternations, the difference is not as impactful as it is for the together reciprocal alternation (intransitive).

Table 9 shows how the baseline performs on the reduced verb sets with a minimum number of 20 attestations per verb.

For the causative-inchoative alternation and the instrument subject alternation, discarding infrequent alternation candidates improves the accuracy and $F_1$ scores of the baseline classifier by several points. For the together reciprocal alternation (intransitive), the minimum frequency condition changes the verb set sizes only slightly, which may be one reason why the scores are not improved by this condition.

As mentioned above, the 1ARSA pattern of the together reciprocal alternation (intransitive) is the rarest one, with only 1,149 sentences covering 221 verb types. The alternating verbs *cling*, *combine*, *connect* and *merge* are each attested only a single time in the *together* pattern. With such a small number of attestations, there is no sufficient evidence for individual alternation candidates to determine whether they appear in the 1ARSA pattern coincidentally, or because it is systematically licensed by the *together* construction of the alternation. The following sections will examine whether the linguistically motivated classification features can alleviate this sparsity problem and achieve better scores for the together reciprocal alternation (intransitive) and the other alternations.

Overall, the baseline feature leads to moderate classification scores. Discarding infrequent verbs is somewhat helpful, but only for the causative-inchoative alternation and the instrument subject alternation. The linguistically motivated features should outperform this frequency-based baseline if they accurately encode the behavior of candidate verbs.

## 10.3 LEMMA-BASED FEATURES

### 10.3.1 *Introduction*

The constructions associated with an alternation express certain semantic roles in different syntactic slots. For verbs participating in the causative-inchoative alternation, the theme is expressed in the EA1 slot and in the IA2 slot; for verbs participating in the instrument subject alternation, the instrument is expressed in the EA1 slot and in the IA2 slot; and for verbs participating in the together reciprocal alternation (intransitive), the themes are expressed in the EA1 slot, in the EA2 slot, and in the IA2 slot.

The lemma-based features for the classifiers are concerned with the lemmas that are found in each of the relevant syntactic slots for each alternation candidate. If the EA1 slot can contain, for instance, an entity filling the instrument role, and the IA2 slot can also contain this type of entity, then there may be a certain overlap in the sets of lemmas observed in these positions. On the other hand, if the verb in question does not alternate, these slots express different semantic roles, so the lemma overlap is expected to be smaller or empty.

Some related work in this area has made use of slot filler lemmas as indicators for various verb classification tasks in the past (see Section 10.3.2). Features based on observed slot filler lemmas have the advantage of being computationally inexpensive, since all relevant information can be directly extracted from a corpus. At the same time, they have the disadvantage of being sensitive to sparsity, since a generalization over different slot fillers can only reliably be made if enough instances of a verb with enough similar slot fillers are available. Infrequent verbs or sentences with unique slot filler lemmas can pose problems here.

Out of all types of features for the alternation-specific classifiers developed in this chapter, the lemma-based features are the most naive. All feature sets aim to approximately represent selectional preference patterns that characterize the alternation-specific constructions. The lemma-based features are expected to be less robust against sparsity than the vector-based and perplexity-based features, since those feature sets have some generalization power that is lacking here.

Verbs that happen to impose the same or similar selectional preferences on all three slots may be misclassified based on these features, but for verbs with different selectional preferences, the features may encode enough information about those preferences to lead to a successful classification.

### 10.3.2 *Related work on lemma-based alternation identification*

Out of the works referenced in Section 9.2 of this thesis, the approach by McCarthy (2000) is most similar to the one implemented here, although her candidate set is smaller. She predicts participation in the causative-inchoative alternation and the conative alternation, using observed slot filler distributions in corpus data to learn alternation behavior of verbs. In her approach, candidate verbs for an alternation are selected based on the syntactic patterns that characterize the alternation's constructions. In the next step, selectional preference models are created to predict whether the slot fillers in the different syntactic patterns indicate alternating behavior or not.

The author uses the SCF acquisition system of Briscoe & Carroll (1997), which categorizes the observed patterns in a parsed corpus into 161 distinct SCF classes. McCarthy removes SCFs that are less frequent than would be expected by chance. The result is an SCF lexicon in which each verb is annotated with the frequency of each SCF, as well as all observed instances of argument heads for each of the relevant slots.

Next, the selectional preferences of each verb for each syntactic slot are approximated based on the observed slot filler lemmas. This is implemented with the method by Li & Abe (1998), which represents preferences for a slot in the form of disjoint classes that partition the leaves of the WordNet noun hierarchy at the time. For each slot, a conditional probability for each of these classes is determined based on the observed slot filler distributions. Frequencies of ambiguous slot fillers, i.e., observed argument heads that are associated with more than one WordNet class, are divided between the available WordNet classes equally.

Based on this tree-cut model, McCarthy calculates the degree of similarity between the preferences of slots in the different patterns involved in an alternation that are associated with each other, i.e., slots in the different syntactic positions of the patterns that are associated with the same semantic role. The similarity for alternating verbs is hypothesized to be stronger than for non-alternating verbs.

McCarthy evaluates this approach using data from 19.3 million words of parsed text from the BNC corpus (Leech 1993). Only verbs that are observed with 10 or more different argument heads in each relevant slot are taken into account, and only argument heads that are found in the WordNet hierarchy at the time are used. From all candidates that are collected in this way, positive and negative examples for each alternation are selected by hand, with the goal of "[obtaining] an even split between candidates which did participate in the alternation and those which did not" (McCarthy 2000: 258). The process involves four human judges, and only candidates with a κ agreement of 75% or more are used for the experiments. This leads to a set of 46 positive examples and 53 negative examples for the causative alternation (from which 7 negative examples are then randomly removed to yield a balanced set of positive and negative examples), as well as 6 positive and 6 negative examples for the conative alternation.

The results are reported in terms of the correlation between the predictions from the model and the human-sourced predictions. With different cuts of the tree model, the causative alternation is identified by the similarity measure with a highly significant correlation to the human judgments, with mean accuracy scores between 71 and 73. For the conative alternation with only 12 candidate verbs, the reported mean accuracy score is 67%. McCarthy (2000: 259) notes that false positives tend to occur when the preferences of different target slots in the alternation are near neighbors in WordNet; for instance, she reports that verbs like *eat* and *drink* have a high probability mass for the "entity" class in both target slots, which makes the slots harder to distinguish using this model.

An alternative implementation presented by McCarthy relies on lemmas directly. This method does not group slot fillers into WordNet classes, but instead directly measures the similarity of presumed related slots of different syntactic patterns in the form of a multiset overlap with respect to the argument heads observed in the relevant slots. While the tree-cut model is reported to result in more false positives than false negatives, the lemma-based approach tends to be more "conservative" and leads to more false negatives. The lemma-based model only achieves a mean 60% accuracy score for the causative alternation, and a 58% mean accuracy score for the conative alternation. This method is therefore weaker than the class-based approach that leverages WordNet information.

The assumption underlying the approach by McCarthy is also the motivation for the experiments performed in the context of this thesis: if the selectional preferences for each syntactic slot can be modeled at a sufficient level of detail, then the similarity or difference between the preferences for syntactic slots can be used to determine whether a verb is likely to participate in an alternation or not. As shown by McCarthy, the pure lemma-based approach is sensitive to sparsity and is outperformed by an approach that attempts to generalize over observed argument heads.

While relying on the WordNet hierarchy improves the classification results, it still means that only a limited set of argument heads in the different slots can be used as indicators for or against the alternation. Recall that the experiments are restricted to verbs that are observed with 10 or more different argument heads in each relevant slot, and only argument heads that actually appear in the WordNet hierarchy at the time are used. In the following sections of the present chapter, vector-based and perplexity-based features will be developed in an attempt to generalize over observed slot fillers without the necessity of relying on hand-made resources like WordNet.

McCarthy (2000: 256) notes that the tree-cut model and the lemma-based model are only applicable to classifying whether or not verbs participate in role-switching alternations. This is because the similarity between the different slots of the patterns involved in an alternation is only of interest if that similarity is an indicator for or against the alternation. Out of the alternations under investigation in this thesis, the causative-inchoative alternation and the instrument subject alternation fall into this category.

Merlo & Stevenson (2001) take an approach to alternation identification that is also based on observed slot fillers for each candidate verb in the different syntactic patterns allowed by the respective alternation. They are concerned with a three-way classification of optionally intransitive English verbs that belong to different classes: unergatives, unaccusatives, and object-drop verbs. Their strategy involves the definition and implementation of a set of linguistically motivated features characterizing the nature of the arguments in each slot of each pattern for each verb.

Merlo & Stevenson use a tagged corpus comprising 65 million words and a parsed corpus comprising 29 million words as sources for the values of their five features for all candidate verbs. Instead of using a lemmatizer, they opt to base the feature values only on observations of the past participle form of each candidate verb, ending in -*ed*. For each of the three classes, 20 verbs from Levin (1993) are chosen (19 for the unaccusative class) such that each candidate appears in the corpora at least 10 times (with one exception, *jogged*, which appears only 8 times).

The five features defined by Merlo & Stevenson aim to distinguish whether a verb's slot fillers in its transitive subject, transitive object and intransitive subject position are more agent-like or more theme-like. Their three classes differ in their assignment of these thematic roles to the syntactic slots, so a good estimate of the thematic role assignment pattern for a verb directly contributes to choosing the most likely class for that verb.

The five features are named *transitivity*, *causativity*, *animacy*, *passive voice*, and *VBN tag*. The *causativity* feature is similar to the lemma overlap feature used by McCarthy (2000). Merlo & Stevenson (2001: 379–380) qualify this feature with a caveat that unergatives are expected to appear in transitive contexts only infrequently, so the lemma overlap is expected to mainly be useful for distinguishing unaccusatives on the one hand, which will have a high lemma overlap between the transitive object slot and the intransitive subject slot, and unergatives and object-drop verbs on the other hand, which will have a lower overlap of slot fillers for these slots.

The *transitivity* feature indicates how frequent the transitive context is for each verb; based on a linguistically motivated scale of markedness contrasting the three verb classes, Merlo & Stevenson (2001: 379) formulate the expectation that unergatives will be the least frequent in the transitive, unaccusatives have intermediate frequency in the transitive, and object-drop verbs are the most frequent in the transitive. The *animacy* feature encodes the extent to which slot fillers are animate and thus likely to fill the agent role. With this feature, Merlo & Stevenson (2001: 384) approximate animacy based on an animacy hierarchy, citing Dixon (1994), Silverstein (1976); the hierarchy indicates that pronouns are the most animate. The feature value is calculated as the ratio of occurrences of pronoun subjects to all subjects for each verb. The *passive voice* and *VBN tag* features are used as additional indicators for transitivity, since both of them imply the availability of a transitive use of the verb.

The *causativity* feature alone achieves an accuracy score of 55.7% in a classification task involving 20 verbs from each of the three classes, making it the strongest individual predictor of alternation participation for the verbs in their dataset. Using all features raises the overall classification accuracy to 69.8%.

In summary, both McCarthy (2000) and Merlo & Stevenson (2001) show that features based exclusively on measuring the overlap between observed slot filler tokens are partially predictive of alternation participation, but adding WordNet categories to characterize the slot fillers in one specific slot semantically or adding features indicating particular semantic properties, such as animacy, improves the classification scores. This suggests that the lemma-based features presented in this section of the thesis will also fall short of the vector-based and perplexity-based features, which represent the semantics of each verb's slot fillers in the different syntactic slots of the relevant patterns in a more general way than the features relying on arguments being attested in multiple slots that correspond to each other.

### 10.3.3 *Method*

The features implemented in this section are based exclusively on the observed lemmas in the various slots of the relevant syntactic patterns for each candidate verb for an alternation. Candidate verbs with a small number of observed instances in the relevant subcorpus are expected to be subject to sparsity issues (which may be alleviated by using more complex feature types) to a larger extent than more frequent candidates.

The fundamental assumption underlying the features in this section is that each semantic role is associated with a category of entities that can typically fill that role, and these entities are in turn associated with a set of lemmas that can be used to describe them. The set of all observed lemmas in a particular argument slot of a verb is thus expected to reflect the semantic role associated with that argument slot.

Consider the verbs *open* and *eat*. Each of these can appear in the 1ARSA and 2ARSA pattern of the causative-inchoative alternation, but only *open* participates in that alter-

Table 10: Examples of sentences for the CIA-alternating verb *open* and the non-alternating verb *eat*.

| | 1ARSA pattern | 2ARSA pattern |
|---|---|---|
| *open* | The door opens.<br>The gate opened.<br>It opened.<br>The window opens.<br>The jar opened. | Ishmael opened the door.<br>The wind opened the door.<br>I opened the chest.<br>I open every drawer in the kitchen.<br>My mother opens the window.<br>The witch opened the gate. |
| *eat* | Bildad ate.<br>My brother is eating.<br>Everybody eats.<br>People eat. | Sharks eat tuna.<br>My brother ate an entire pizza.<br>My cat eats buttons.<br>People eat pizza. |

nation. Table 10 shows a few example sentences with these verbs in the two syntactic patterns for illustrative purposes.

Given the attestations of the two verbs shown in Table 10, the lemma-based features will encode the overlap in the lemmas observed in the EA1 slot, IA2 slot, and EA2 slot of each verb. If the data reflects the assumptions detailed above, the overlap should be strongest between the EA1 slot and the IA2 slot for alternating verbs, like *open*. For non-alternating verbs like *eat*, the overlap should be strongest between the EA1 slot and the EA2 slot. The same should hold for the instrument subject alternation.

For verbs participating in the together reciprocal alternation (intransitive), the expectations for slot overlap are as follows. In the 1ARSA pattern, the subject, or EA1 slot, is expected to encode themes. The same holds for the EA2 and the IA2 slots. The lemma overlap features should indicate a larger overlap between the IA2 slot on the one hand and the EA1 and EA2 slots on the other hand if the verb alternates, and a lower value if the verb does not alternate.

The sentences in (68) through (71) illustrate the expected difference. Since the verb *mix* participates in the together reciprocal alternation (intransitive) and the highlighted arguments are both valid theme slot fillers, each of them can be expressed in any of the available theme slots in the alternation's two constructions characterized by the 1ARSA and 2ARSA syntactic patterns. In contrast to this, the verb *pour* does not participate in the together reciprocal alternation (intransitive). The element in the IA2 position expresses a location here and cannot be moved to the subject position. Swapping the positions of the two highlighted elements leads to an ungrammatical sentence (as in (71b)), while the alternating verb allows either order of arguments (as in (69b)).

(68) TRAI-alternating verb, 1ARSA pattern:
  a. The *eggs* and the *cream* mix together.
  b. The *cream* and the *eggs* mix together.

(69) TRAI-alternating verb, 2ARSA pattern:
  a. The *eggs* mix into the *cream*.
  b. The *cream* mixes into the *eggs*.

(70) Non-alternating verb, 1ARSA pattern:
  a. * The *cream* and the *bowl* pour together.
  b. * The *bowl* and the *cream* pour together.

(71) Non-alternating verb, 2ARSA pattern:
  a. The *cream* pours into the *bowl*.
  b. * The *bowl* pours into the *cream*.

The fact that the 2ARSA pattern can be realized with three different prepositions poses another challenge for the identification of verbs that participate in the together reciprocal alternation (intransitive). Levin (1993) specifies for each alternating verb the preposition(s) with which that verb can appear in the patterns associated with the alternation. In the experiments presented here, the classifier only attempts to identify the verbs that participate in the alternation and can appear in at least one of the prepositional patterns, without attempting to predict which of the prepositions is or are possible for a given verb.

#### 10.3.3.1 *Feature definition*

The slot filler overlap features will be implemented in several different ways here. Each alternation involves an EA1 slot, an IA2 slot, and an EA2 slot. In order to determine which of these slots can share slot fillers, two features per slot pair will be created. The features are asymmetrical, meaning that for each pair of slots A and B, one feature will reflect the number of slot fillers in slot A that are also attested at least once in slot B, and another feature will reflect the number of slot fillers for slot B that are also attested at least once in slot A. Furthermore, the feature values for each pair of slots will be determined once on a type level, and once on a token level. This leads to six type-level slot overlap features, and six token-level slot overlap features, which will be described in the following.

features 01-01 through 01-06: *Type-based slot filler lemma overlap.* In the type-level approach, the set of possible slot filler lemmas for one of the relevant slots is viewed as a representation of the selectional preferences of the given verb for that slot. This view does not take the frequency of individual slot fillers into account, and thus does not distinguish between one-off slot fillers and those that are observed frequently. A sentence in which someone *stirs their coffee* with their *finger* would add the slot filler *finger* to the set for the verb *stir*, while sentences in which the instrument used to stir a drink is a *spoon* would contribute the slot filler *spoon* to the set. These two possible slot fillers are then treated equally in terms of their contribution to the type-based slot overlap features.

The main goal of the type-based slot filler overlap features is to determine whether individual attestations of possible slot fillers are already informative with regard to the selectional preference of the current verb for that slot. If this is the case, there is no need to take slot filler frequency into account as well.

The value for these features is determined as the percentage of slot fillers in the relevant slot in one pattern type that are also attested *at least once* in the corresponding slot in the alternative pattern type. This leads to a range of feature values between 0 and 1.

In the example sentences given in Table 10 on page 230, the verb *open* appears five times in the 1ARSA pattern, with five different slot fillers in the EA1 position. Out of these five slot filler lemmas, three (*door*, *gate*, *window*) are also observed in the IA2 position for the same verb. Thus, the feature determining type-level slot overlap between these two slots will have a value of 0.6 – independently of the frequency of each slot filler in the EA1 slot or the IA2 slot, and independently of the number of observed slot fillers in the IA2 slot. In the other direction, the slot overlap feature between the IA2 slot and the EA1 slot will also receive a value of 0.6, since three of the five slot filler types that appear in the six sentences instantiating the 2ARSA pattern are also attested in the EA1 slot.

Contrast this with the example sentences given in the table for the non-alternating verb *eat*. It appears four times in the 1ARSA pattern, with four different slot fillers in the EA1 position, and none of these slot fillers are attested in the IA2 slot. Thus, the feature determining the overlap between these slots will have a value of 0 for this verb.

The feature encoding the slot filler overlap between the EA1 slot and the EA2 slot will yield the following results. For the *open* sentences, the feature will have a value of 0, since none of the subject slot fillers in the 1ARSA pattern also appear as subjects in the 2ARSA pattern. For the *eat* sentences, the feature will have a value of 0.5, since two of the four observed EA1 slot filler lemmas (*brother*, *people*) are also observed in the EA2 slot.

FEATURES 01-07 THROUGH 01-12: *Token-based slot filler lemma overlap.* Analogous to the type-based features described previously, these features express the overlap of the sets of relevant slot filler lemmas on the token level. Essentially, this means that the frequency of slot filler lemmas will be taken into account, such that frequent lemmas are weighted more strongly than infrequent ones.

In Table 10, the verb *open*, which participates in the causative-inchoative alternation, appears five times in the 1ARSA pattern. Three of the tokens observed in the EA1 position, *door*, *gate* and *window*, are also attested in the IA2 position. The token-level lemma overlap feature comparing these two slots will therefore have a value of 0.6. The feature encoding the overlap in the other direction ("how likely is it for a token that appears in the IA2 position to be observed in the EA1 position?") will have a value of 0.66, since four of the six sentences have a lemma in the relevant position that is also attested in the corresponding slot in the other pattern.

Note that the term *token* here does not refer to inflected word forms. Instead, each instance of an argument lemma counts as a token, in contrast to the previous set of features, for which each lemma counted only once. Table 11 on page 233 lists all lemma-based features for the classification.

Table 11: Overview of all lemma-based features for the classification of candidate verbs for each of the three alternations. Feature sets 01a and 01b will be evaluated separately.

| Lemma-based feature sets |
|---|
| 01a-lemma-features-type-based: |
| 01-01_EA1_slot_filler_lemmas_attested_in_IA2_slot |
| 01-02_IA2_slot_filler_lemmas_attested_in_EA1_slot |
| 01-03_EA1_slot_filler_lemmas_attested_in_EA2_slot |
| 01-04_EA2_slot_filler_lemmas_attested_in_EA1_slot |
| 01-05_EA2_slot_filler_lemmas_attested_in_IA2_slot |
| 01-06_IA2_slot_filler_lemmas_attested_in_EA2_slot |
| 01b-lemma-features-token-based: |
| 01-07_EA1_slot_filler_lemma_tokens_attested_in_IA2_slot |
| 01-08_IA2_slot_filler_lemma_tokens_attested_in_EA1_slot |
| 01-09_EA1_slot_filler_lemma_tokens_attested_in_EA2_slot |
| 01-10_EA2_slot_filler_lemma_tokens_attested_in_EA1_slot |
| 01-11_EA2_slot_filler_lemma_tokens_attested_in_IA2_slot |
| 01-12_IA2_slot_filler_lemma_tokens_attested_in_EA2_slot |

Table 12: Scores (accuracy, precision, recall, $F_1$) for the lemma-based feature sets (type-based and token-based).

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 01a-lemma-features-type-based | 61.13 | 67.20 | 42.84 | 52.46 |
| ISA | 01a-lemma-features-type-based | 52.31 | 51.62 | 63.05 | 55.49 |
| TRAI | 01a-lemma-features-type-based | 51.67 | 42.00 | 50.00 | 40.70 |
| CIA | 01b-lemma-features-token-based | 60.47 | 68.09 | 40.61 | 50.83 |
| ISA | 01b-lemma-features-token-based | 54.90 | 53.95 | 71.70 | 61.22 |
| TRAI | 01b-lemma-features-token-based | 36.17 | 31.83 | 51.00 | 40.23 |

### 10.3.4 *Results*

The results of the lemma-based feature sets on the full verb sets for each of the three alternations are shown in Table 12.

Concerning the difference between the scores for the type-level feature set and the corresponding token-level features, there is no general trend across alternations for one of these to systematically lead to better classification results. This is most likely due to sparsity in the data: if most slot fillers are attested exactly once in the relevant position, there will be only a minimal difference in predictions depending on whether the frequency is taken into account or not.

For the causative-inchoative alternation, the type-level features perform slightly better, but for the instrument subject alternation, the token-level features raise the accuracy and $F_1$ score of the classifier by a few points. For the together reciprocal alternation (intransitive), both lead to the same $F_1$ score, although the accuracy is higher for the type-level feature set.

The best accuracy for these features is achieved on the verb set for the causative-inchoative alternation. The subcorpora for the causative-inchoative alternation are among the largest ones created for these experiments. The more attestations of a verb are taken into account when determining the feature values, the higher the likelihood is for a slot filler to also appear in the corresponding slot in another pattern. For the instrument subject alternation, it may be the case that lemmas that could plausibly appear in two corresponding slots are less likely to in fact be attested in both slots, because the subcorpus for the 2ARSA pattern type is smaller than the smallest of the two subcorpora for the causative-inchoative alternation. For the together reciprocal alternation (intransitive), the lemma overlap features are even less likely to capture whether two slots express the same semantic argument, since the subcorpora are even smaller.

Table 13 on page 235 shows how the lemma-based features perform on the reduced verb sets with a minimum number of 20 attestations per verb.

For most alternations and lemma-based feature sets, discarding infrequent candidates leads to slightly higher accuracy and $F_1$ scores. For the causative-inchoative alternation and the together reciprocal alternation (intransitive), the type-level lemma feature sets yield higher accuracy values than their token-level counterparts.

Table 13: Scores (accuracy, precision, recall, $F_1$) for the lemma-based feature sets (type-based and token-based), with a minimum alternation candidate frequency of 20.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 01a-lemma-features-type-based | 62.06 | 64.63 | 53.66 | 58.09 |
| ISA | 01a-lemma-features-type-based | 55.73 | 56.37 | 54.07 | 53.34 |
| TRAI | 01a-lemma-features-type-based | 63.50 | 49.50 | 63.00 | 55.33 |
| CIA | 01b-lemma-features-token-based | 59.29 | 62.80 | 49.51 | 54.58 |
| ISA | 01b-lemma-features-token-based | 58.24 | 58.36 | 58.45 | 57.70 |
| TRAI | 01b-lemma-features-token-based | 46.00 | 36.50 | 54.00 | 43.67 |

Overall, the diversity of verb meanings and of slot filler semantics seems to not be sufficiently covered by these features that are based purely on the sets of observed slot filler lemmas.

## 10.4 vector-based features

### 10.4.1 *Introduction*

The type-based and token-based lemma overlap features presented and discussed in the previous section are intended to approximate the selectional preferences for different semantic arguments of a verb. However, they are unable to group semantically related lemmas together, so they have no ability to generalize across individual slot fillers. The goal of this section is to explore whether classification features that leverage distributional information about slot fillers are more successful at the alternation identification task than the pure lemma-based features.

Semantic information beyond the level of individual slot filler lemmas can be retrieved from handcrafted resources, like VerbNet (Kipper, Dang & Palmer 2000, Kipper et al. 2006, 2007), FrameNet (Ruppenhofer et al. 2006), WordNet (Miller 1995), or similar lexical databases. A disadvantage of such resources is that they may not be comprehensive and may lack generalization power. Unsupervised approaches to representing meaning are less exact, but also easier to transfer to other languages or corpora. Instead of describing word meaning based on a manually-designed taxonomy or semantic composition rules, approaches that rely on distributional information describe the semantics of words in terms of their distribution in a corpus in relation to other words (Sahlgren 2008).

The distributional hypothesis, the idea of "knowing a word by the company it keeps", goes back to ideas discussed by Harris (1954) and Firth (1957). This family of approaches to modeling natural language gained popularity with the success of word2vec (Mikolov et al. 2013), GloVe (Pennington, Socher & Manning 2014), and other distributional frameworks that were able to achieve reasonable performance on benchmark tasks like word analogy identification or sentiment analysis without needing access to labeled training data (for an overview, see e.g. Baroni, Dinu & Kruszewski 2014, Schnabel et al. 2015).

In these distributional frameworks, the instances of a word throughout a training corpus are observed with respect to their immediate context. The frequencies of all context words are the basis for the vector that is used to represent the target word; a dimensionality reduction step makes the approach more robust against sparsity issues. The resulting word vectors, or embeddings, can be used to determine whether two words have similar distributions, which tends to correlate with similar or related meanings. Depending on the context size and the details of how context words are being counted, word vectors that are close to each other in the vector space may be more likely to be semantically related (e.g. *coffee* and *milk*) or syntactically related (e.g. *coffee* and *drink*), as noted by e.g. Bullinaria & Levy (2007: 518), Kiela & Clark (2014), and Andreas & Klein (2014).

For the purposes of identifying verbs that participate in a specific alternation, a distributional approach to approximating the selectional preference overlap between the different relevant slots may be more successful than the approach purely based on lemma overlap, which relies on a certain number of specific slot fillers being attested in the different relevant syntactic patterns at least once. As an indicator for slot filler overlap, distributional representations should be more robust against sparsity than the pure lemma-based features.

The sentence pair in (72) illustrates the motivation behind the vector-based features. The highlighted tokens are the themes of the sentences; the theme slots are filled by different lemmas. The pure lemma-based approach would therefore regard this sentence pair as an indicator against the verb's participation in the causative-inchoative alternation.

(72) a. *The number of whales* is diminishing. (*1ARSA pattern*)

b. Whaling is diminishing *the whale population*. (*2ARSA pattern*)

If the verb *diminish* participates in the alternation, the EA1 slot and the IA2 slot should largely accept semantically related fillers. In turn, if the slot fillers are more semantically dissimilar, that could be an indicator that the verb does not participate in this alternation.

Whether the slot filler lemmas *population* and *number* belong to a shared, specific semantic type is not the main question here. What is more important is whether these slot fillers are *more* similar to each other than they are, respectively, to the slot fillers in the third relevant slot of the alternation. In this case, finding that *number* is more similar to *population* than it is to *whaling* – the slot filler in the EA2 slot of this alternation – makes it most likely that the EA1 slot and the IA2 slot in the sentence pair do in fact encode the same semantic argument for the verb *diminish*. In other words, the sentence pair follows the expected role assignment for the alternation. Since individual sentences and words may not be sufficient to determine which slots are more similar to each other, values for the vector-based features will be calculated by averaging the similarities over all observed slot fillers for each relevant slot.

However, there is no requirement for all semantic arguments of alternating verbs to belong to different semantic types. Thus, if a verb imposes the same selectional preferences on more than one of the relevant syntactic slots, the vector-based features would yield inconclusive evidence, because no informative difference could be observed between the vector representations for different slots in the relevant syntactic patterns.

Of course, this is a problem that all classification features attempting to approximate selectional preferences will be subject to.

The classifiers based on vector-based features would also have difficulties identifying alternating verbs where the selectional preferences are less specific. For instance, if a verb requires a semantic argument filler to be of the type *entity*, then the different slot fillers for that argument slot may be so heterogeneous that they cannot be used to distinguish alternating verbs from non-alternating verbs. On the other hand, if the type for an argument slot of a verb is relatively restricted, for instance if the verb requires an argument of the type *consumable substance*, there is a higher chance for slot fillers for this argument to be more homogeneous and distinct from the slot fillers for other arguments of that verb.

Different verb alternations may be subject to these issues to different extents. Additionally, the semantics of each verb that participates in an alternation may come with more or less strict selectional preferences for the different argument slots.

For verbs participating in the causative-inchoative alternation, the selectional preferences imposed on a theme require the entity in that slot to be able to undergo the described change of state or event. For instance, the verb *melt*, which participates in the causative-inchoative alternation, can only be meaningfully applied to a substance that can be solid or liquid under certain circumstances. Melting *butter* is a likely scenario, while melting *wood* is typically impossible.

For verbs participating in the instrument subject alternation, the instrument role needs to be filled by an entity that can be used as an instrument for the act described by the verb. This can either be an instrument that is typical for the verb (such as a *spoon*, which is a likely instrument for *stirring*), or it can be an entity that is an unlikely instrument but is being used in the scenario described by the sentence to perform the act described by the verb (such as a *crowbar*, which is an unlikely but possible instrument for *opening a door*). Compared to verbs participating in the causative-inchoative alternation, the selectional preferences for this slot of this alternation are less strict, since coercion plays a larger role than in the context of the causative-inchoative alternation – in the real world, it is easier to use an atypical instrument to do something than to impose a change of state on something that does not typically exhibit that state, as in the case of *wood* that *melts*.

Finally, for the together reciprocal alternation (intransitive), alternating verbs express themes in all relevant slots of all constructions that are associated with the alternation. The vector-based features will therefore approximate selectional preferences by estimating to what extent the slot fillers for the different slots of the 1ARSA and 2ARSA pattern are semantically similar. If the IA2 slot fillers of a verb are noticeably distinct from the EA1 or EA2 slot fillers for the same verb, a likely explanation is that the verb does not participate in the together reciprocal alternation (intransitive), and expresses a non-theme argument in the prepositional phrase and/or a non-theme argument in the subject slot.

Since distributional word representations can be trained on unannotated data, the vector-based approach to the identification of alternating verbs can be transferred to other languages without requiring manual annotation, although some variation in reliability is expected for languages that are morphologically richer than, for instance, English, if no lemmatization is performed prior to learning the vectors (Cotterell & Schütze 2015).

In recent years, contextualized word embeddings have become popular in natural language processing research, with notable early contributions being ELMo embeddings (Peters et al. 2018) and BERT (Devlin et al. 2019). Static word vectors, like those produced by word2vec or GloVe, produce one representation per term, determined by taking into account all occurrences of that term in the training data; for contextualized embeddings, each occurrence of a term in a specific context is represented with a different, context-specific representation (Ethayarajh 2019). These contextualized embeddings outperform static embeddings on benchmark tasks like GLUE (Wang et al. 2018), and have therefore become popular over the years.

For the purposes of the experiments conducted here, static embeddings are preferable over contextualized ones. Throughout the alternation-specific subcorpora, the various slot filler lemmas are observed in different contexts; in order to find semantic relationships between the slot fillers in these different positions, using one static embedding per lemma is more appropriate than determining different contextualized representations for each occurrence of the slot filler throughout the corpus.

### 10.4.2 *Related work on vector-based alternation identification*

While the idea of generalizing over observed slot fillers to approximate selectional preferences for certain slots is not new, early approaches (e.g. Schulte im Walde 2000, McCarthy 2000) typically did so by referring to a static lexical resource like WordNet. As distributional methods became more popular and more practical thanks to the availability of the necessary frameworks and computational power, they began to also be used in the context of alternation identification tasks. Sun & Korhonen (2009), whose aim is not to identify alternating verbs but to cluster verbs into semantic classes, use spectral clustering (Ng, Jordan & Weiss 2001) with a fixed number of clusters to categorize slot fillers into semantically-related sets. The vectors that are used to represent the different slot fillers are based on high-dimensional similarity matrices.

Baroni & Lenci (2009) present an approach to modeling semantic relations between words in a system they call a distributional semantic memory. This resource is based on corpus data from the 2-billion word web corpus ukWaC (Ferraresi et al. 2008) and takes the shape of a network with weighted links between concepts. The links are syntactically-informed: the authors focus on noun-to-noun and noun-to-verb connections and collect distributional information based on certain manually-defined types of dependency paths between the concepts in question, e.g. direct paths between nouns and verbs, as in *kill+obj+victim*. For the 20,000 most frequent nouns and 5,000 most frequent verbs as target concepts, their resulting database of concept links (weighted by mutual information) consists of 69 million tuples.

Baroni & Lenci apply this distributional memory to various tasks, including word similarity rating, noun categorization, analogy recognition, and the identification of verbs in the causative-inchoative alternation. Their dataset for the alternation identification task consists of 232 causative-inchoative alternating verbs from Levin (1993) as positive instances, and 170 non-alternating verbs from the same source as negative instances. For each verb in the dataset, the weighted links from the distributional memory that involve the verb and a transitive subject, the verb and an intransitive subject, or the verb and a direct object are taken into account. These links, they argue, should reveal that alternating verbs have a higher similarity between intransitive subjects and transitive objects ("the things that are broken also break"), while this should not be the case for non-alternating verbs. The cosines for all observed subjects and objects are calculated, and paired t-tests are conducted to determine whether the difference in similarity between transitive subjects and direct objects on the one hand and intransitive subjects and direct objects on the other hand is statistically significant.

According to the authors, this approach is "completely successful in detecting the distinction" (Baroni & Lenci 2009: 7). While this is encouraging with regard to the usefulness of distributional information for alternating verb identification, there are some differences between the setup in their work and the setup pursued in the context of this thesis. Most notably, the non-alternating verbs listed in Levin (1993: 27–30) for this alternation – the set on which the experiments by Baroni & Lenci are based – do not necessarily appear in both syntactic patterns associated with the constructions of the alternation. For instance, one set of the negative verbs that are listed are Give verbs, illustrated with the following example:

(73) a. They gave the bicycle to me.
b. * The bicycle gave to me.

To what extent intransitive sentences involving, for instance, the verb *give*, appear in the ukWaC corpus is not reported in the paper. Baroni & Lenci do not specify how they handle verbs for which one or more of the sets of links are empty. It is possible that those verbs are excluded from the set from the beginning – Levin (1993) lists over 300 alternating verbs for the causative-inchoative alternation, and the set used by Baroni & Lenci only contains 232 positive verbs, so some prior selection is involved in Baroni & Lenci's setup.

Since many of the non-alternating verbs listed in Levin (1993) rarely or never behave similarly to alternating verbs, the approach by Baroni & Lenci cannot be compared directly to the approach taken in this thesis, where all candidate verbs for an alternation appear in all relevant syntactic patterns at least once in the corpus.

Kann et al. (2019) implement and describe two approaches to identifying acceptable syntactic frames related to alternations for individual verbs.[1] The first approach is based on verb embeddings, while the second one uses sentence embeddings instead. The five alternations they investigate include the causative-inchoative alternation, which is also discussed in this thesis, and the spray-load alternation and the dative alternation, each of which allows participating verbs to express the same set of arguments in different syntactic slots, similar to the together reciprocal alternation (intransitive), which is discussed in this thesis.

In order to train a classifier that can be used to predict the acceptability of a given verb with a given frame, the authors create two datasets that are labeled for acceptability. To create the datasets, known acceptable sentences are used to semi-automatically derive unacceptable counterparts. 515 alternating verbs are sourced from the sections in Levin (1993) on the five alternations. Each of the five alternations is associated with two frames; each of the verbs in the dataset is labeled as to whether it participates in each of the five alternations, that is, which of the ten frames under investigation in the paper are expected to be acceptable for the verb.

1 In this thesis, the term *frame* typically refers to a semantic representation, while the syntactic environments that Kann et al. (2019) are concerned with would here be termed *constructions*. In the paragraphs summarizing this work, the term *frame* will be used as in the referenced paper.

Sentences for the dataset are generated semi-automatically, using hand-selected combinations of verbs and plausible arguments, resulting in acceptable or unacceptable sentences. This method leads to a dataset with all sentences following the simplest possible scheme for each frame, while being reliable in the sense that all argument slot fillers are typical or at least unsurprising for the given verb.

Kann et al. (2019: 290) note that there are no English verbs that can appear in the inchoative frame, but not in the causative one. As a result, no negative examples for this combination can be included. They also note that the verbs sourced from Levin (1993) are only categorized there as being grammatical or ungrammatical for a specific alternation; acceptability judgments on the eight frames in their experiments that are unrelated to the given alternation are augmented semi-automatically.

Kann et al. approach the task of classifying sentences as acceptable or unacceptable with two methods, one based on word embeddings, the other on sentence embeddings. The approach using sentence embeddings will be discussed in a later section of this thesis. For the approach using word embeddings, they use 300-dimensional GloVe embeddings trained on 6 billion tokens (Pennington, Socher & Manning 2014) as well as a set of embeddings trained on 100 million tokens from the BNC corpus, using a single-directional LSTM (Hochreiter & Schmidhuber 1997). The classification task is then implemented in the form of a multi-layer perceptron with a single hidden layer, which takes as its input the embedding for the given sentence's root verb. Each classifier predicts whether each of the two frames of one alternation is acceptable for the given verb or not.

While Kann et al. do not refer to the task as *alternation identification*, they do set up their classifiers to predict for each verb and syntactic frame whether the combination is acceptable or not, using one classifier per alternation. In other words, a prediction of two acceptable frames from an alternation for a given verb can be regarded as a prediction of alternation participation for that verb, while a prediction of only one or no acceptable frame from the alternation corresponds to a prediction of non-participation.

Results are reported in terms of accuracy and also in terms of Matthews correlation coefficient (Matthews 1975), which is more robust for unbalanced classification tasks than $F_1$ or accuracy. For four of the ten frames, the GloVe embeddings tend to predict the (positive) majority class. According to the authors (Kann et al. 2019: 293), this is an indicator that the datasets do not contain enough examples to generalize across verbs, and they mention that a balanced dataset might reduce the positive class bias and make the results more interpretable.

The causative frame is among the ones that lead to only majority-class predictions in this setting, while the inchoative frame is among the highest-performing ones (MCC $> 0.45$). As mentioned by the authors, there are no verbs in the dataset that allow the inchoative frame without also allowing the causative frame. This means that the prediction of an acceptable combination of a verb with the inchoative frame can be regarded as a prediction of that verb's participation in the causative-inchoative alternation; however, since the dataset is biased toward the positive class, i.e., it contains more positive instances in this category than negative ones, and the evaluation is presumably conducted on the same number of inchoative examples as causative examples, the results do not imply a better-than-chance overall classification accuracy for this alternation.

The fact that the dataset used by Kann et al. is unbalanced for all five alternations discussed there makes a comparison to results from different approaches challenging. Out of the five alternations featuring in Kann et al.'s work, only the causative-inchoative alternation is discussed both there and in this thesis.

Kann et al. use embeddings for a sentence's root verb as the main indicator guiding the classifier's decision. This indicates an underlying assumption that the distributional representation for the verbs that is learned by GloVe encodes enough information about the behavior of the verb that a prediction of acceptable frames is possible. As shown in the results by Kann et al. for this task, that assumption is not justified, since the classifiers do not perform above the majority-class baseline. In the experiments conducted in the context of this thesis, verb embeddings as sole indicators for or against alternation participation will therefore not be pursued. Instead, embeddings for the arguments observed for each verb will be used as the main source for the vector-based features used by the classifier. In other words, the experiments will be set up in a way that more explicitly models the assumptions about slot overlap, by involving features derived from the different slot fillers instead of relying on embeddings for the verbs alone.

### 10.4.3 *Method*

For each of the alternations under investigation, the values of the vector-based features for each alternation candidate will be derived from all observed instances of that candidate in the relevant subcorpora. In practice, this means that each vector feature will represent the similarity of the set of all lemmas in one slot to the set of all lemmas in another slot. If this approach is a reliable approximation of selectional preference similarity across the slots, it should allow the classifier to decide whether the current verb belongs in the alternating class or the non-alternating class.

The vector-based features will be presented in the following. Table 14 on page 242 shows examples of sentences involving the verb *open*, which participates in the causative-inchoative alternation, and the non-alternating verb *eat*. These examples will illustrate how the vector-based features estimate selectional preference similarity for different slots of a verb.

The slot fillers for the EA1, EA2, and IA2 slots for these verbs can be represented in a vector space. The more semantically similar two items are, the closer they are located to each other. Figure 75 on page 242 shows a two-dimensional representation of the slot fillers that appear in the sentences in Table 14. Items that fill the same semantic argument slot are grouped via colors: yellow items are entities that can *eat*, green items are entities that can *be eaten*, red items are entities that can *open* and blue items are entities that can *be opened*. For the sake of simplicity, each slot filler only appears once in the figure, even if it occurs multiple times in the example sentences in the table.

Dotted lines frame all items that can appear in the same slot for a verb. The figure contains one group of items that can appear in the EA1 slot for *open*, one group of items that can appear in the IA2 slot for the same verb, and one group of items that can appear in the EA2 slot; likewise for the items that can appear in the three slots for the verb *eat*. The $\times$ inside each set denotes the position of the centroid for the items in that set.

Table 14: Examples of sentences for the CIA-alternating verb *open* and the non-alternating verb *eat*.

| | 1ARSA pattern | 2ARSA pattern |
|---|---|---|
| *open* | The door opens.<br>The gate opened.<br>It opened.<br>The window opens.<br>The jar opened. | Ishmael opened the door.<br>The wind opened the door.<br>I opened the chest.<br>I open every drawer in the kitchen.<br>My mother opens the window.<br>The witch opened the gate. |
| *eat* | Bildad ate.<br>My brother is eating.<br>Everybody eats.<br>People eat. | Sharks eat tuna.<br>My brother ate an entire pizza.<br>My cat eats buttons.<br>People eat pizza. |

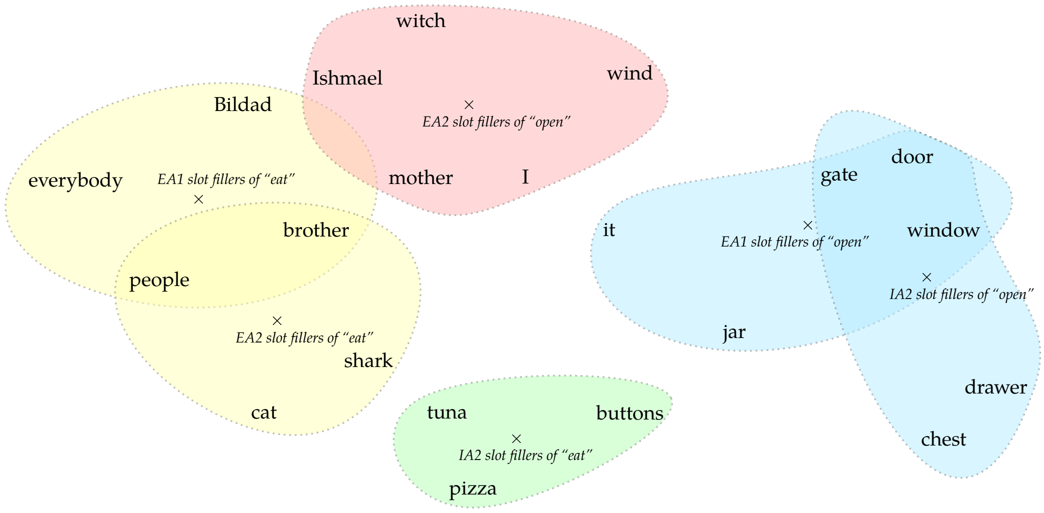


Figure 75: Two-dimensional representation of the relative positions of vectors for the slot fillers that appear in the sentences in Table 14.

In order to implement the features for the classification, the first required step is to collect vectors for each observed slot filler. As mentioned earlier, contextualized embeddings are not an option for this approach to alternating verb identification, so only static vectors will be considered here. The vectors for these classification features can either be derived from the available data, or from some freely available pre-trained resource. Learning vectors for a specific dataset has the advantage that each vector truly reflects the distributional behavior of that word in that dataset; the main disadvantage is that sparsity in smaller training corpora will impact the size and generalizability of the vector set.

Several types of techniques for the acquisition of static word embeddings of the type needed here are available at the time of writing, the most popular ones being word2vec (Mikolov et al. 2013), GloVe (Pennington, Socher & Manning 2014), and fastText (Bojanowski et al. 2017). Newer techniques for the collection of word vectors tend to be more successful on benchmark tasks like word similarity test sets. There are also benchmark tasks for which representations of subword units are more successful than vectors that represent full lemmas; however, for the task pursued in this thesis, since the feature values will be based on embeddings for lemmas instead of word forms, the advantage of subword embeddings is expected to be negligible.

FastText vectors are built from vectors of substrings of characters contained in the given word. In order to test whether subword-level embeddings lead to a better performance in the alternation identification task than embeddings for full words, the vector features will be implemented once based on GloVe vectors and once based on fastText vectors. Word2vec will not be used. Pretrained fastText vectors will be sourced from `https://fasttext.cc/docs/en/english-vectors.html` (crawl-300d-2M.vec.zip: 2 million word vectors trained on Common Crawl, containing 600B tokens), and pretrained GloVe vectors will be sourced from `https://nlp.stanford.edu/projects/glove/` (Common Crawl, 840B tokens, 2.2M vocab, cased, 300d vectors).

#### 10.4.3.1 *Feature definition*

For each slot of a candidate verb in each pattern for an alternation, a set of vectors is collected that corresponds to the set of observed slot fillers. Frequency is taken into account: slot fillers that appear many times are assumed to be more typical for the given slot than slot fillers that only appear once or twice. After collecting the embeddings for all slot fillers, a centroid is calculated for each slot. The centroid is the vector that represents the average of a set of vectors; it will be used to represent the set of argument slot fillers for the current slot. For slots that are subject to very strict and specific selectional preferences, the centroid is more likely to be close to the vectors of the observed slot fillers, whereas more heterogeneous slot fillers for a slot will tend to be more distant from their centroid and from each other.

The three slot types that are relevant in all three alternations – the EA1, EA2 and IA2 slots – will be compared in terms of the cosine distance of the centroids of their observed slot fillers. Each vector-based classification feature will correspond to the pairwise cosine distance between two slots; since this measure is symmetrical, three features are necessary to compare three slots to each other.

The value range of the vector features is defined by the range of the cosine function. The minimal value is 0 and indicates completely identical vectors. The maximal value is 1 and indicates maximally distinct vectors.

While the vector-based features are expected to somewhat alleviate the sparsity issues associated with the lemma-based features, it can still be the case that an observed slot filler cannot be linked to a vector and thus cannot contribute directly to the vector-based features. The vector sets used for the features each contain around 2 million items, so that proper names, misspellings or infrequent words cannot be expected to be covered by the vector set. In the context of static embeddings, using one single UNK vector for unknown words is not advisable with respect to the goals of these experiments, since it would conflate the different slot fillers and not yield any useful information about the semantics of the different lemmas. Instead, two other strategies to handle missing vectors will be pursued and evaluated.

First, missing lemmas will simply not be taken into account in the creation of a centroid for a given verb's slot fillers in a particular position. For instance, if a verb appears with 9 different slot filler tokens, out of which 4 cannot be represented with an embedding from the pre-trained vector set, then only the remaining 5 vectors will be used to represent the set of slot fillers in this position for this verb. It is also possible for verbs to be attested exclusively with slot filler lemmas that cannot be represented with embeddings from the pre-trained vector set. In these cases, the value of features involving this slot will automatically be set to 1, which corresponds to maximally different vectors.

It was previously mentioned that predicting an alternation participation label for a verb that is not attested, or not attested with all necessary patterns, is not justified. In a similar vein, claiming maximally distinct slot fillers based solely on the fact that representations for one of the sets of slot fillers are not available is not ideal. A second strategy for these cases inserts placeholder vectors that are filtered by the slot in which the missing lemma was observed. For instance, if there are 4 lemmas without vector representations in the set of EA1 slot fillers for a given verb, then 4 instances of a slot-filler-specific placeholder vector will be inserted into the set. The value of this placeholder vector is determined as a centroid of all slot fillers in the EA1 slot across all verbs in the current candidate pool. While this conflates the observed slot fillers for alternating verbs with those of non-alternating verbs, it may be more appropriate than using one UNK placeholder vector independently of the currently relevant slot. The impact of these different approaches will be discussed in Section 10.4.4.

FEATURE 02-01: *Cosine distance between EA1 and IA2 slot fillers.* This feature approximates the semantic similarity of items observed in the EA1 position and items observed in the IA2 position. As shown in Figure 75 on page 242, the items in these two positions for the verb *open* have some lemma overlap (*door*, *gate*, *window* appear in both slots), and the centroids of these two sets are fairly close to each other. For the verb *eat*, the sets of items in these two slots have no overlap, and their centroids are further apart.

The lemma overlap and short distance between the centroids for the verb *open* are a result of the fact that this verb participates in the alternation, and thus expresses the same semantic role (the theme) in the two slots. The verb *eat* does not participate in the alternation, and expresses different roles in the slots.

Note that lemma overlap between the slots is not a requirement for the vector-based features. Even without overlap, the shorter distance between the centroids of the relevant slots would be indicative of alternation participation.

FEATURE 02-02: *Cosine distance between EA1 and EA2 slot fillers.* This feature approximates the semantic similarity of items observed in the EA1 position and items observed in the EA2 position. Figure 75 on page 242 shows that these sets have some overlap and a short distance between their centroids for the verb *eat*, while they are completely distinct and further apart for the verb *open*.

As mentioned earlier, observing that a verb imposes similar selectional preferences on its EA1 and IA2 slots is not sufficient to conclude that the verb participates in an alternation. The verb may simply have multiple arguments on which it imposes similar or identical selectional preferences. For instance, a verb may express its agent and patient in different syntactic slots, but require both of them to be *animate*.

Adding feature 02-02 is meant to support the classifier in distinguishing these cases. If a verb simply imposes similar selectional restrictions on different semantic roles, both feature 02-01 and 02-02 will indicate a short distance between centroids for the respective slots. Otherwise, if feature 02-02 indicates a short distance between the centroids

Table 15: Overview of all vector-based features for the classification of candidate verbs for each of the three alternations. Feature sets 02a, 02b, 02c and 02d will be evaluated separately.

| Vector-based feature sets |
|---|
| 02a-vector-features-glove-without-placeholder: |
| 02-01a_glove_cosine_between_EA1_and_IA2_slot_fillers_without_placeholders |
| 02-02a_glove_cosine_between_EA1_and_EA1_slot_fillers_without_placeholders |
| 02-03a_glove_cosine_between_IA2_and_EA1_slot_fillers_without_placeholders |
| 02b-vector-features-glove-with-placeholder: |
| 02-01b_glove_cosine_between_EA1_and_IA2_slot_fillers_with_placeholders |
| 02-02b_glove_cosine_between_EA1_and_EA1_slot_fillers_with_placeholders |
| 02-03b_glove_cosine_between_IA2_and_EA1_slot_fillers_with_placeholders |
| 02c-vector-features-fasttext-without-placeholder: |
| 02-01c_fasttext_cosine_between_EA1_and_IA2_slot_fillers_without_placeholders |
| 02-02c_fasttext_cosine_between_EA1_and_EA1_slot_fillers_without_placeholders |
| 02-03c_fasttext_cosine_between_IA2_and_EA1_slot_fillers_without_placeholders |
| 02d-vector-features-fasttext-with-placeholder: |
| 02-01d_fasttext_cosine_between_EA1_and_IA2_slot_fillers_with_placeholders |
| 02-02d_fasttext_cosine_between_EA1_and_EA1_slot_fillers_with_placeholders |
| 02-03d_fasttext_cosine_between_IA2_and_EA1_slot_fillers_with_placeholders |

for a verb's EA1 slot fillers and the verb's EA2 slot fillers, but the centroids involved in feature 02-01 are further apart, this is an indicator that the verb does not alternate. This is the case for the slot fillers shown for *eat* in the example vector space in Figure 75 on page 242.

FEATURE 02-03: *Cosine distance between IA2 and EA2 slot fillers.* This feature expresses the similarity between the IA2 slot fillers and the EA2 slot fillers of a verb. For the causative-inchoative alternation and the instrument subject alternation, observing similar slot fillers for these slots is a likely indicator for the verb imposing similar selectional preferences on its different semantic arguments.

For the third alternation under investigation here, the together reciprocal alternation (intransitive), a different behavior is expected. Since themes are expressed in all three slots of alternating verbs, all three vector features should express a high similarity of the slot fillers in order to indicate alternation participation. For verbs that do not alternate, each of the slots may express a non-theme role: the EA1 slot and the EA2 slot may express a (set of) agent(s) instead of a (set of) theme(s), and the IA2 slot may express, for instance, a goal or manner instead of a theme.

Candidate verbs for the together reciprocal alternation (intransitive) should be labeled as participating in the alternation if all three features indicate similar slot fillers, and as not participating if the slot fillers are less similar. Whether the distinction can be learned based on the corpora and small verb sets available for this alternation will be discussed in Section 10.4.4.

Table 15 lists all vector features for the classification.

Table 16: Scores (accuracy, precision, recall, $F_1$) for the vector-based feature sets (top: GloVe, bottom: fastText).

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| | glove | | | | |
| CIA | 02a-vector-features-glove-without-placeholder | 59.34 | 60.82 | 55.01 | 57.68 |
| | 02b-vector-features-glove-with-placeholder | 58.25 | 59.50 | 54.21 | 56.51 |
| ISA | 02a-vector-features-glove-without-placeholder | 53.37 | 53.57 | 57.25 | 54.70 |
| | 02b-vector-features-glove-with-placeholder | 53.14 | 53.26 | 54.05 | 53.31 |
| TRAI | 02a-vector-features-glove-without-placeholder | 31.83 | 15.83 | 24.50 | 19.83 |
| | 02b-vector-features-glove-with-placeholder | 36.00 | 28.83 | 42.50 | 27.43 |
| | fasttext | | | | |
| CIA | 02c-vector-features-fasttext-without-placeholder | 61.74 | 64.79 | 53.45 | 58.16 |
| | 02d-vector-features-fasttext-with-placeholder | 61.67 | 65.39 | 52.60 | 57.83 |
| ISA | 02c-vector-features-fasttext-without-placeholder | 53.61 | 53.58 | 52.93 | 52.51 |
| | 02d-vector-features-fasttext-with-placeholder | 53.42 | 53.60 | 54.46 | 54.45 |
| TRAI | 02c-vector-features-fasttext-without-placeholder | 41.67 | 29.00 | 53.50 | 37.40 |
| | 02d-vector-features-fasttext-with-placeholder | 33.67 | 22.83 | 37.50 | 27.47 |

### 10.4.4 *Results*

The results of the vector-based feature sets on the full verb sets for each of the three alternations are shown in Table 16.

For the causative-inchoative alternation and the instrument subject alternation, including a slot-specific placeholder vector in the case of unknown slot filler tokens does not impact the classification scores by much. For the together reciprocal alternation (intransitive), adding a placeholder improves the results somewhat in the GloVe setting, but when fastText vectors are being used, including the placeholder leads to worse classification scores.

The different vector sources, GloVe and fastText, yield similar classification scores for the causative-inchoative alternation and the instrument subject alternation. For the together reciprocal alternation (intransitive), the GloVe vectors seem to cover the slot fillers slightly more exhaustively, so that the feature set without placeholders leads to higher accuracy and $F_1$ scores than the parallel feature set using GloVe vectors.

The alternation whose verbs are most accurately identified by these feature sets is the causative-inchoative alternation. This may be because the semantic arguments expressed in the EA1, EA2 and IA2 slots belong to sufficiently distinct semantic categories. For the instrument subject alternation, the slot fillers in slots expressing the instrument role may be more heterogeneous, which results in the centroid for one set of slot fillers possibly being located far away from the centroid for the other set of slot fillers, even if both slots encode an instrument.

Table 17 on page 247 shows how the vector-based feature sets perform on the reduced verb sets with a minimum number of 20 attestations per verb.

Discarding infrequent verbs improves the classification accuracy and $F_1$ scores for the causative-inchoative alternation by a few points in the setting using GloVe vectors. When fastText vectors are used, the $F_1$ score for the causative-inchoative alternation is

Table 17: Scores (accuracy, precision, recall, $F_1$) for the vector-based feature sets (top: GloVe, bottom: fastText), with a minimum alternation candidate frequency of 20.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| | glove | | | | |
| CIA | 02a-vector-features-glove-without-placeholder | 62.19 | 62.99 | 62.46 | 63.13 |
| | 02b-vector-features-glove-with-placeholder | 61.31 | 62.42 | 60.64 | 60.81 |
| ISA | 02a-vector-features-glove-without-placeholder | 51.68 | 52.47 | 55.38 | 53.27 |
| | 02b-vector-features-glove-with-placeholder | 52.95 | 52.74 | 58.03 | 54.16 |
| TRAI | 02a-vector-features-glove-without-placeholder | 26.50 | 20.00 | 28.00 | 23.33 |
| | 02b-vector-features-glove-with-placeholder | 27.00 | 17.00 | 34.00 | 26.00 |
| | fasttext | | | | |
| CIA | 02c-vector-features-fasttext-without-placeholder | 61.50 | 62.89 | 58.96 | 60.21 |
| | 02d-vector-features-fasttext-with-placeholder | 62.34 | 63.08 | 58.70 | 60.58 |
| ISA | 02c-vector-features-fasttext-without-placeholder | 52.88 | 52.11 | 59.20 | 54.58 |
| | 02d-vector-features-fasttext-with-placeholder | 50.75 | 51.13 | 56.56 | 53.08 |
| TRAI | 02c-vector-features-fasttext-without-placeholder | 33.50 | 18.00 | 41.00 | 30.00 |
| | 02d-vector-features-fasttext-with-placeholder | 48.00 | 46.50 | 65.00 | 49.00 |

also improved slightly. For the instrument subject alternation and the together reciprocal alternation (intransitive), none of the vector-based feature sets benefit noticeably from the minimum frequency cutoff.

There is no clear trend across alternations, vector sources and frequency filtering conditions regarding the impact of using placeholders. Presumably, providing more and less noisy training data to the classifier would be more beneficial for the classification scores than the inclusion of the placeholder vectors.

In the condition without a minimum candidate frequency filter, the fastText vectors slightly outperform the GloVe vectors in both settings (with/without placeholder vectors) for the causative-inchoative alternation. For the together reciprocal alternation (intransitive), fastText vectors are better than GloVe vectors in the setting without placeholders, but perform on par in the setting using placeholders. Finally, for the instrument subject alternation, no clear advantage of one vector source over the other can be observed. In the condition with a minimum candidate frequency threshold, fastText vectors outperform GloVe vectors for the together reciprocal alternation (intransitive), but the scores are comparable for the other two alternation. As expected, the subword information encoded in fastText vectors is not a critical advantage for this classification task, since all words are lemmatized.

Overall, the vector feature sets do not seem to confirm the expectation that centroids of slot filler lemmas are sufficiently representative of different selectional preferences of a verb for its different semantic arguments. This may be an effect of the size and quality of the subcorpora used for these experiments. When only a handful of vectors are available to characterize each slot for a verb in an alternation-specific pattern, the generalization power provided by the distributional approach does not lead to strong gains in classification accuracy or $F_1$ scores.

Both the lemma-based and the vector-based classifiers suffer from sparsity, which impacts the results for the together reciprocal alternation (intransitive) the most. The following section is concerned with implementing a set of perplexity-based features, for which the observed instances of all candidate verbs are augmented with generated alternate versions instantiating the other pattern(s) for the relevant alternation. This may systematically improve the classification scores.

## 10.5 PERPLEXITY-BASED FEATURES

### 10.5.1 *Introduction*

The previously-discussed approaches to alternation identification relied on various features that were derived from observed corpus occurrences of each verb in each relevant pattern. This section is concerned with the idea of directly testing the central assumption of verb alternation participation: if a verb participates in an alternation, it should be grammatical and acceptable in all constructions involved in the alternation.

Thus, for any sentence from the corpus that instantiates one of the syntactic pattern types for an alternation, it should be possible to create one or more alternate sentences (depending on the number of patterns involved in the alternation) that instantiate the other pattern(s) licensed by the alternation's constructions – provided that the sentence's root verb participates in the alternation. A rule-based transformation script can be used to generate alternate sentences. The last step is to determine whether the resulting alternate sentence(s) is or are acceptable.

The syntactic patterns for each alternation were discussed and presented in Section 9.4. Mapping them to each other is mainly a matter of correctly aligning the different slots with each other. The resulting sentences still need to be grammatical, for instance in terms of case and number agreement between verbs and their arguments.

Once the alternate sentences have been generated, some measure of acceptability is necessary in order to determine whether the alternate sentences for a verb tend to be "good" (indicating participation in the alternation) or "bad" (indicating non-participation). As with the previously-discussed approaches, the data points here are expected to be distributed on a spectrum, but they may exhibit certain trends that can be used to distinguish between alternating and non-alternating verbs.

In the context of the alternation identification task, a binary measure of acceptability is not necessary or useful. The alternate sentences that are being generated by the transformation script are expected to typically be "less good" than original sentences. They may include some pragmatic or information-theoretic oddities that are introduced by the transformation script, whether or not the verb participates in the alternation at hand.

Examples of these effects are provided by Bresnan et al. (2007), who attempt to determine which of the different constructions of the English dative alternation is preferred for a given verb and its arguments (see also Section 9.1 of this thesis). They use web search results for specific instances of the double object construction and the prepositional object construction to show that intuitions of grammaticality can be overriden by various factors, making individual construction usages possible even if they are not "supposed to be grammatically possible" (Bresnan et al. 2007: 6). Bresnan et al. present examples that illustrate that the acceptability of each of the dative alternation's constructions is influenced by factors such as pronominality, givenness, heaviness, definiteness, and length of the constituents in the different slots.

This is an important aspect to keep in mind, since it puts into question the assumption that a given verb either can or cannot appear in the different constructions licensed by an alternation. A verb may participate in the relevant alternation, but be observed in a sentence whose constituents can only be expressed in a somewhat awkward manner in one of the relevant constructions. This is why the perplexity-based features introduced in this section are not based on an expectation that all alternation-specific constructions are acceptable to the same degree for a given verb and set of arguments. Instead, the features aim to approximate trends over all usages of a verb.

Romain (2018) is concerned with the observation that some verbs that participate in alternations have a strong preference for one of the relevant constructions, instead of being distributed more equally. Additionally, some combinations of alternating verbs and theme arguments are only valid in one of the constructions, as illustrated in (74) (taken from Romain 2018: 21).

(74) a. I have never *broken* a *law* in my life.
b. * A *law* has never *broken*.

Using collostructional analysis (Stefanowitsch & Gries 2003, 2005), Romain (2018: 99) measures how likely a given lexeme is to appear in a certain slot of an alternation-specific construction. In other words, the measure determines which lexemes are particularly attracted to a construction. The analysis also sheds light on subtle semantic differences between alternating constructions, since verbs that prefer one of the constructions may have more in common semantically with each other than they do with verbs without a preference or verbs that prefer another construction of the same alternation.

Romain conducts this analysis for 29 verbs that are known to participate in the causative-inchoative alternation. 26 of these verbs are found to have a preference for one or the other construction relevant to the alternation. These preferences are found based on observed construction and verb frequencies in a subset of the CoCA corpus (Davies 2010). The approach taken in this thesis augments observed instances of alternation-specific constructions with generated alternate counterparts; this would also be a possible avenue for future research into alternation strength and construction preferences of alternating verbs. However, this is beyond the scope of this thesis.

The findings by Romain do not invalidate the perplexity-based features for the identification of alternating verbs proposed in this section. Even if a particular known alternating verb is more likely to appear in one of the alternation-specific constructions than in the other relevant construction(s), the perplexity-based features are still expected to be able to distinguish between these cases and cases where the verb cannot appear in one or more of the constructions at all. Special cases like the one illustrated in (74) are expected to be in the minority. Overall, the features proposed in this section aim to draw on the general trends that can be observed across instances of verbs in the relevant constructions.

In order to implement these features, a measure of acceptability for the generated alternate sentences is required. According to Lau, Clark & Lappin (2015, 2017), the term *grammaticality* is typically associated with binary judgments that distinguish sequences that are possible in a language from those that are not (Lau, Clark & Lappin 2017: 1204), while *acceptability* judgments are graded (Lau, Clark & Lappin 2015: 1618). They also contrast acceptability with *probability*, which describes how likely a given sequence is to appear in a text in the language. Probability is sensitive to factors like sequence length and token frequency to a larger extent than acceptability. Lau, Clark & Lappin propose several measures that can be used to derive acceptability judgments from probability scores. For the purposes of this chapter, such a measure is useful if it can accurately distinguish between generated alternate sentences that are good examples of the alternation, and those that are bad examples or are simply uninterpretable. Some existing work on predicting acceptability scores for sentences will be summarized in Section 10.5.2.

Since the generated alternate sentences are generally expected to be less acceptable than their original counterparts, the classifier should take into account not an absolute acceptability score, but instead the delta between scores for original sentences and their generated alternate counterparts, and assign a label to each verb based on whether that delta for that verb tends to be lower (meaning that alternate sentences are often fine, and the verb participates in the alternation) or higher (meaning that alternate sentences are often bad, and the verb does not participate in the alternation).

Implementing the features based on these ideas requires choosing a source for acceptability scores. Ideally, the scores would be based on human-sourced ratings; as habitual users of language, they would presumably be able to assign acceptability scores that lead to a successful distinction between alternating verbs and non-alternating ones. However, collecting human judgments for datasets the size of the alternation-specific corpora used here would be expensive and time-consuming.

One of the goals of this thesis chapter is to find approaches to alternating verb identification that minimally rely on manually-created resources and ratings. As an alternative to human-sourced acceptability judgments, the features implemented in this section will instead rely on an unsupervised setup that can use unannotated data as a basis on which to decide whether new input is acceptable or not. One possibility to source such scores is by consulting a large language model and using the score assigned to a sequence by the language model as an approximation of acceptability.

*Perplexity* is a measure that encodes how well a given input to a language model can be predicted using the probability distributions learned by that model. This allows intrinsic evaluations of language models using held-out test data: given a language model that has been trained on the available training data, how well can the sequences in the test data be predicted by that model? Comparisons between different language models or different architectures can also be made based on perplexity (see e.g. Krause et al. 2018, 2019): which of two competing systems is more successful at learning the probability distributions that allow it to recognize that an unseen sequence belongs to the same language as its training data?

Perplexity is defined as the inverse of the probability of the test set, normalized by the number of tokens in that set. In other words, a language model that assigns a large probability to the test set will have a small perplexity value for that test set. Normalizing by the number of tokens in the given sequence reduces the impact of the issues raised by Lau, Clark & Lappin (2015, 2017) about approximating acceptability based on language model scores: scores for sequences of different lengths can yield comparable perplexity values as indicators of the per-token "acceptability" of the sequence.

Typically, perplexity is used to determine whether a language model can recognize its test set as belonging to the same language as its training set. For the experiments to be conducted in this thesis, the focus is on a different aspect: perplexity is used to differentiate between verbs whose observed sentences instantiating one pattern of an alternation can be transformed to yield acceptable alternate sentences, and verbs for which the alternate sentences are less acceptable. The underlying expectation is that an unacceptable alternate sentence should lead to a much higher perplexity value than an acceptable one, in relation to the original sentences on which the alternate sentences are based. To a certain extent, the behavior of verbs with regard to verb alternations is observable in the training data; beyond these cases, a language model may be able to generalize over latent classes of verbs and constructions, and thus recognize that a verb that has never been observed in a particular context is nevertheless acceptable in that context.

In the context of this thesis, since the classifier assigns a label to each verb based on the verb lemma, not on individual tokens, all instances of the verb in all relevant patterns will be taken into account as the basis for the perplexity-based features. If sentences with the verb that instantiate one syntactic pattern of the alternation in question tend to lead to a large perplexity delta when the sentences are transformed to their alternate counterparts, then it is unlikely that the verb participates in the alternation. If the delta tends to be smaller, the verb is more likely to alternate. The relevant observation here is the relative difference in perplexity between the original sentence and its alternate version (not the absolute score, since that depends too much on vocabulary choice and other properties of the sentence as a whole that are not of interest here) as an average over all observed instances.

### 10.5.2 *Related work on acceptability prediction*

Automatically assigning acceptability scores to individual sentences is not a trivial task. Lau, Clark & Lappin (2015, 2017) propose several different acceptability measures that attempt to normalize for sequence length and token frequency. Even then, the scores are not necessarily informative by themselves. Fortunately, the classifier features to be developed here do not require acceptability scores that are informative in isolation. Instead, the features will be designed to reflect trends in acceptability differences from original sentences to generated alternate sentences, for individual verbs.

Lau, Clark & Lappin (2015) present an attempt to predict acceptability scores for sentences using probabilistic language models, and show that their approach leads to reasonable correlation with gold judgments from crowd workers on a dataset of 2500 sentences. The annotations are collected in three different degrees of gradedness – one binary, one consisting of a scale with four steps, one consisting of a sliding scale. The authors find a high correlation of mean ratings between the different modes and decide to use the four-step scale for the gold labels.

They point out that acceptability as a graded property of linguistic sequences is distinct from the probability of a sequence as predicted by a language model. In general, probability is more strongly impacted by the length of a sequence and word frequency than acceptability is. To normalize the probability score for a sequence with regard to properties like sequence length and token frequency, and thus generate acceptability scores more closely reflecting human judgments, Lau, Clark & Lappin design several acceptability measures whose values can be derived from the log probability of a sentence as determined by a language model.

The language models that they use as the source for probability values for sentences are trained on 100 million tokens from the BNC corpus. The architectures range from simple n-gram models to Bayesian and neural network models of the time.

To evaluate the predicted acceptability judgments, Lau, Clark & Lappin determine the correlation between predictions and human judgments using Pearson's $\rho$ for the sentences in their test set. The best-performing language model is a recurrent neural network (RNN) model (Mikolov et al. 2011, Mikolov 2012), with a correlation of 0.53. However, the authors also note that a perfect correlation is difficult to achieve even among human annotators, and that the upper bound for the task is therefore somewhere below 1.0.

The acceptability measures proposed by Lau, Clark & Lappin (2015, 2017) all rely on various types of language models whose output for a given input sequence is a probability score that is determined by the model with respect to its available training data. While the authors find that the RNN model leads to the best correlation with human acceptability judgments, there have been some notable advances in language modeling since then, so a more recent architecture is likely to yield better results for making the generalizations that are necessary to distinguish sequences that are expected in the language from sequences that are not expected, or less acceptable.

Warstadt, Singh & Bowman (2019) present a corpus of linguistic acceptability that includes, among others, acceptable and unacceptable sentences in the patterns involved in the causative-inchoative alternation. They use this corpus to train an LSTM to classify input sequences into the binary categories acceptable/not acceptable. As mentioned above, binary labels are not sufficient for the experiments conducted in this thesis, so the approach by Warstadt, Singh & Bowman will not be pursued further here.

In addition to the approaches of Lau, Clark & Lappin (2017) to using language model scores to approximate human-assigned acceptability scores, there have also been studies on whether language models exhibit comparable reactions to stimuli as human brains, e.g. by Hashemzadeh et al. (2020). They expose an LSTM that was trained on acceptable language to nonsensical, out-of-distribution stimuli, and compare its activations with those of human brains when exposed to the same type of stimuli. They find that there is a statistically significant relationship between an LSTM's representation of nonsensical stimuli and the representation created by human brains. In the context of determining whether generated alternate sentences are acceptable or not, this finding supports the approach of approximating acceptability based on a language model's score for the given input sequence.

### 10.5.3 *Related work on alternation identification based on acceptability measures*

In the approach by Kann et al. (2019) to classifying verbs with regard to their acceptability with different alternation-related frames, classifiers are trained based on verb embeddings in one setting, and on sentence embeddings in another. The dataset created by the authors for this task, as well as the setting using verb embeddings, are summarized in Section 10.4.2 of this thesis.

In their second approach, Kann et al. use sentence embeddings from neural networks to predict whether a verb licenses a certain syntactic frame or not. Their goal is to determine to what extent the representations learned by artificial neural networks resemble the linguistic competence of humans. They train one classifier per alternation in their dataset, implemented as a multi-layer perceptron with a single hidden layer which takes a fixed-length sentence embedding for the given sentence as its input, and predicts whether the combination of the verb and syntactic frame observed in that sentence is acceptable or not. If the sentence embeddings encode enough information about a sentence to accurately predict the acceptability of the given frame for the given verb, then verbs that participate in a given alternation should be predicted to be acceptable with both frames associated with that alternation, while non-alternating verbs should be predicted to be acceptable with only one or none of the relevant frames.

For their classification setup involving sentence embeddings, a bidirectional LSTM encoder taking ELMo-style contextualized embeddings (Peters et al. 2018) as input is trained on a task that involves distinguishing between 12 million real and fake sentences from the BNC. The acceptability classifier uses the training data generated for the Kann et al. paper as well as 10,000 examples from the corpus of linguistic acceptability (Warstadt, Singh & Bowman 2019).

In the discussion of the results of the acceptability classifier, the authors observe that the classifiers are most successful on alternations that involve at least one intransitive frame, which applies to the understood object alternation, the causative-inchoative alternation, and the *there*-insertion alternation. Their explanation for this trend is that "[...] [i]ntransitive verb frames are the simplest syntactic frames possible, and it might be expected that they are easiest to recognize" (Kann et al. 2019: 294). However, these different intransitive patterns all express different semantic arguments of the verb. The classifiers are not set up to make a distinction between e.g. the inchoative frame of a verb in the causative-inchoative alternation, and the intransitive frame of the understood object alternation. This is in contrast to the experiments conducted in this thesis, where the availability of an intransitive frame for a verb is seen as a necessary, but not sufficient condition for that verb's participation in the causative-inchoative alternation. A more detailed analysis of the predictions from the classifiers from Kann et al. may lead to additional insights with regard to whether the classifier assigns the positive label because the verb appears with an acceptable *number* of arguments, or whether the semantic roles of those arguments are also taken into account. This distinction is not one of the main goals of Kann et al., as their focus is on predicting the acceptability of individual (alternation-related) syntactic frames for verbs; for this thesis, on the other hand, it is highly relevant since all positive and negative instances in the datasets for this thesis are observed with the same syntactic patterns, i.e., with the same number of arguments, and the challenge lies in recognizing which constructions are instantiated by these syntactic patterns for each verb.

As noted in Section 10.4.2, a comparison between the results presented in this thesis and the results from Kann et al. is difficult due to their unbalanced datasets and the fact that only one alternation, the causative-inchoative alternation, is discussed both in their paper and in this thesis. The accuracy of 85.4 reported for this alternation from the classifier based on sentence embeddings, in combination with the 77.6 accuracy from a majority-class baseline, indicates that the sentence embeddings do contribute to the classification of these verbs, but are not sufficient to reliably distinguish between alternating and non-alternating verbs.

Warstadt et al. (2020) present the BLiMP dataset, which is meant to facilitate research on acceptability prediction for pairs of minimally distinct sentences. Among other properties, argument structure is covered in the dataset, with 1,000 sentence pairs instantiating the patterns involved in the causative-inchoative alternation. In each pair, one sentence is labeled as acceptable and one is labeled as unacceptable, and the authors report the performance of n-gram, LSTM, and transformer language models on the task of predicting the correct binary label.

This acceptability dataset is unsuitable for the experiments to be conducted here for several reasons. First, it follows a binary perspective of acceptability, which is not desired here, as discussed above. Second, the dataset contains only 1,000 sentence pairs in the relevant patterns for the causative-inchoative alternation; the sentences in a pair typically share some or all arguments, but contain different verbs. This means that a prediction on the level of verb lemmas based on this dataset is likely to be unsuccessful, since there is not enough training data to learn the alternation properties of individual verb lemmas.

### 10.5.4 *Method*

Two preprocessing steps are necessary in order to determine feature values for the perplexity-based classification features. First, a series of scripts will be implemented that are responsible for generating each observed sentence's counterpart(s) in the other construction(s) associated with the relevant alternation. Next, each original sentence and its generated counterpart will be assigned a perplexity score as an approximation of acceptability. These perplexity scores will be the basis of the feature values.

#### 10.5.4.1 *Rule-based reordering scripts*

In order to compare the perplexity scores of original sentences with those of their alternation-specific counterparts, each sentence's (set of) counterpart(s) must first be generated. Since dependency trees for all original sentences are available and the syntactic patterns involved in the different alternations can be described in terms of dependency structure, the generation of alternate sentences will be implemented with a set of rule-based scripts.

Unlike the approach by Kann et al. (2019), the point of departure in this implementation are observed instances of alternating and non-alternating verbs from a corpus. Kann et al. generate simple example sentences for the different combinations of verbs and frames; in the approach presented here, the original sentences are examples of each verb's usage patterns in the corpus, and the generated counterparts for each sentence should resemble the original sentence as much as possible, i.e., if the verb appears with an unusual argument, then that argument will also appear in the generated alternate sentence(s).

Each script first identifies for each given sentence instantiating a particular pattern the position of the relevant syntactic arguments. For instance, the sentence shown in Figure 76 on page 255 contains a `root` verb at index 1, the `nsubj` phrase contains the token at index 0, and the `obj` phrase contains the tokens from index 2 to index 3.

In the next step, the arguments are reordered to form the alternate sentence for the current original sentence, depending on which alternation and which patterns are currently being examined. The sentence displayed in Figure 76 contains a transitive sentence, so its intransitive counterpart may be informative as to whether the root verb

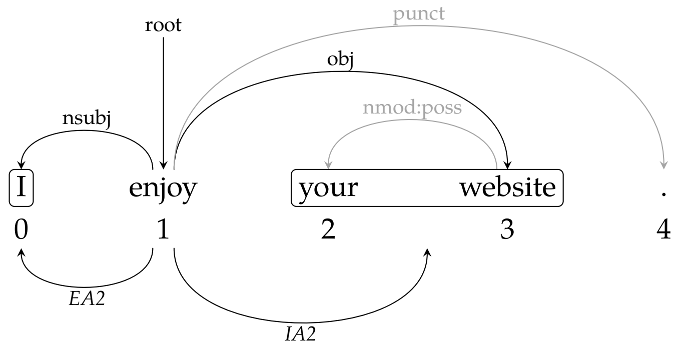


Figure 76: Dependency tree for the original transitive sentence #*ENCOW-01-00000114*.

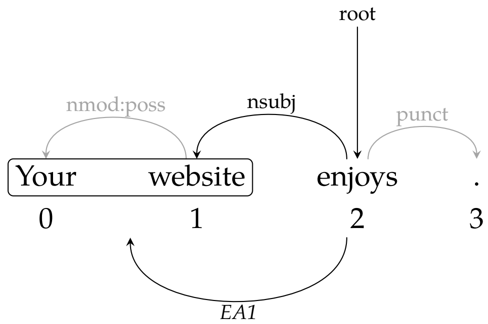


Figure 77: Dependency tree for the generated alternate form of sentence #*ENCOW-01-00000114* with respect to the causative-inchoative alternation.

*enjoy* participates in the causative-inchoative alternation. If it does, then the entity in the IA2 slot in the original sentence expresses the theme of the verb, and the same role would be expressed by the entity in the EA1 slot in the sentence's intransitive counterpart.

The available arguments from the original sentence are incorporated into the generated sentence under the assumption that both sentences instantiate the relevant constructions of the alternation. For the example sentence given above, this is achieved by deleting the original `nsubj` phrase, and promoting the original `obj` phrase to the `nsubj` position. Following these steps leads to the resulting sentence shown in Figure 77. Some rules to force agreement between the new `nsubj` element and the `root` verb are applied to make sure the resulting sentences are grammatical.

While the verb *enjoy* does appear intransitively in the original ENCOW corpus – otherwise, it would not be a part of the alternation-specific subcorpus for the causative-inchoative alternation –, the argument appearing here as the new subject, *your website*, is an unexpected choice for this argument slot for the verb *enjoy*. Human speakers of English would expect the subject of *enjoy* to correspond to an experiencer, which would typically be an animate entity, not a website. If the language model has learned this expectation as well, it will assign a higher perplexity score to the generated alternate sentence than it would for sentences whose verbs participate in the alternation.

Transforming sentences instantiating the 2ARSA pattern of the causative-inchoative alternation or the instrument subject alternation to alternate sentences that instantiate the corresponding 1ARSA pattern is a matter of deleting some argument and reordering the remaining ones. In the other direction, an additional argument has

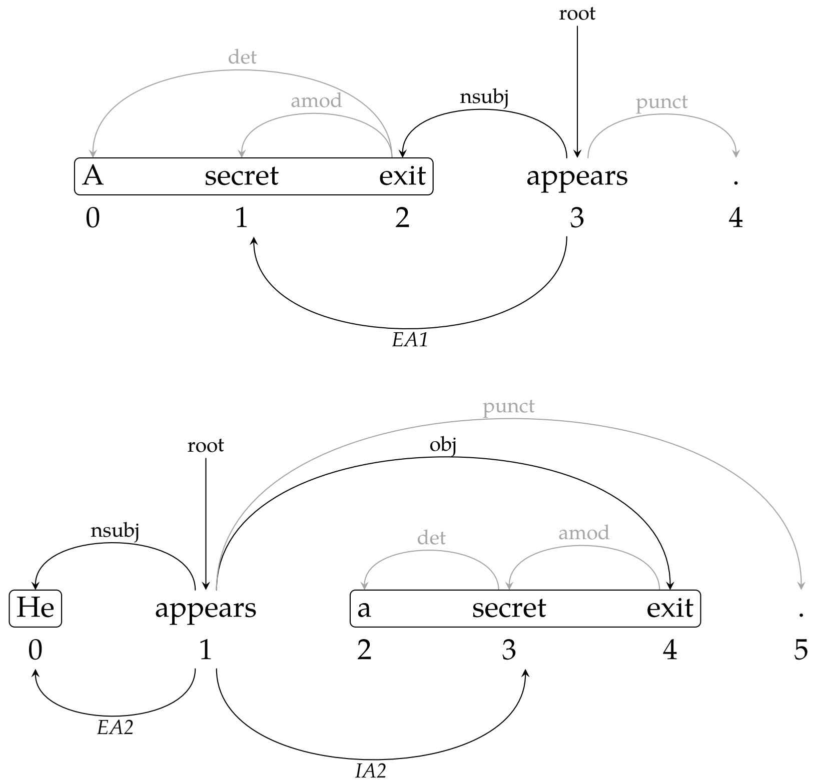


Figure 78: Dependency trees for the original sentence *#ENCOW-01-00023623* (top) and its reordered alternate version with respect to the causative-inchoative alternation (bottom).

to be inserted in the appropriate position. The scripts that generate alternate sentences all insert pronouns in those positions. For instance, for the original sentence displayed in Figure 78, the original `nsubj` phrase is moved to the `obj` position, and a pronoun is inserted as the new `nsubj` argument. Only pronouns that grammatically agree with the `root` verb are taken into consideration. If the verb is observed with pronoun subjects in the original transitive subcorpus, the most frequent pronoun subject observed for this verb is selected for the new `nsubj` position. Otherwise, the pronoun is chosen at random. Note that the reordering step is performed independently of any assumption that the resulting sentences will actually be interpretable. In fact, sentences with non-alternating verbs may be completely uninterpretable in their reordered form, as in the case illustrated in Figure 78.

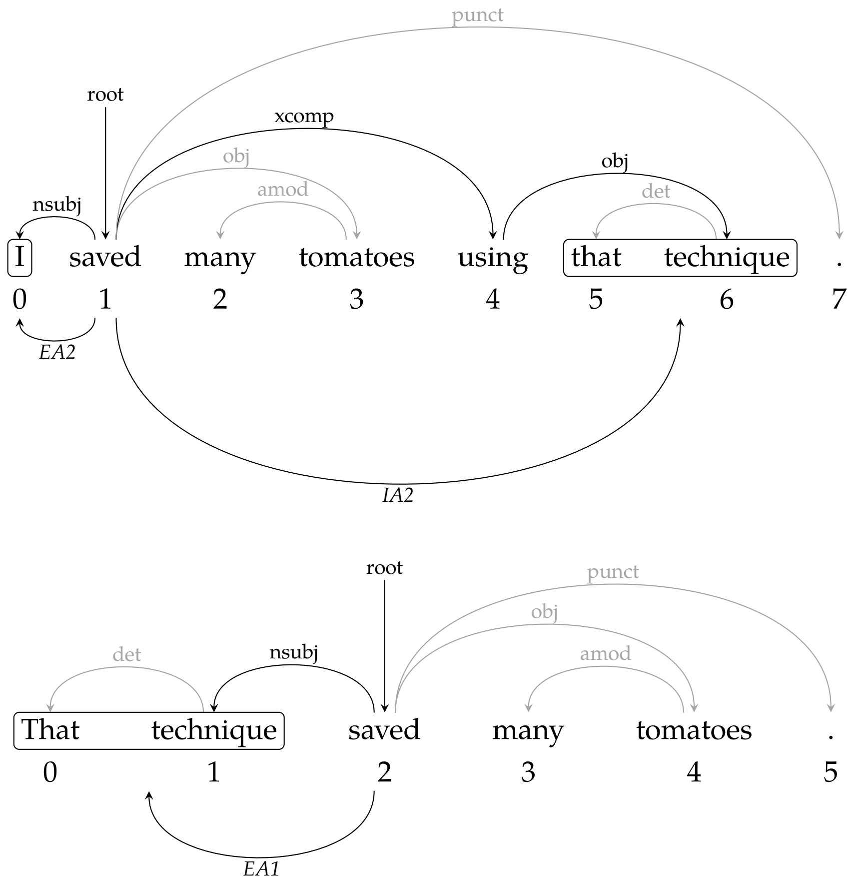


Figure 79: Dependency trees for the original sentence *#ENCOW-01-02486815* (top) and its reordered alternate version with respect to the instrument subject alternation (bottom).

Sentences instantiating the 2ARSA pattern of the instrument subject alternation can be transformed to the 1ARSA pattern by moving the IA2 from the `advcl/xcomp` or `obl` phrase to the `nsubj` position, as illustrated in Figure 79. In the other direction, the `nsubj` phrase is moved to a new *with* or *using* phrase, and a new subject is inserted, following the same strategy described above.

Transforming sentences from the subcorpora for the together reciprocal alternation (intransitive) does not involve deleting or adding arguments. Instead, the observed arguments are moved to different positions in the sentence. Original sentences that instantiate the 1ARSA pattern contain a conjunct subject and the adverbial modifier *together*. They are transformed to instantiate the 2ARSA pattern by keeping one of the original `nsubj` arguments in the new `nsubj` position, and moving the other one into a new prepositional phrase. The 2ARSA pattern of the together reciprocal alternation (intransitive) can be realized with the prepositions *into*, *to* and *with*. In order to avoid making assumptions about which verbs will work well with which preposition, all patterns are generated, and the features for the classifier will be able to take all patterns into account.

Since the patterns for the together reciprocal alternation (intransitive) are more rare than those associated with the other alternations, the subcorpus for this alternation

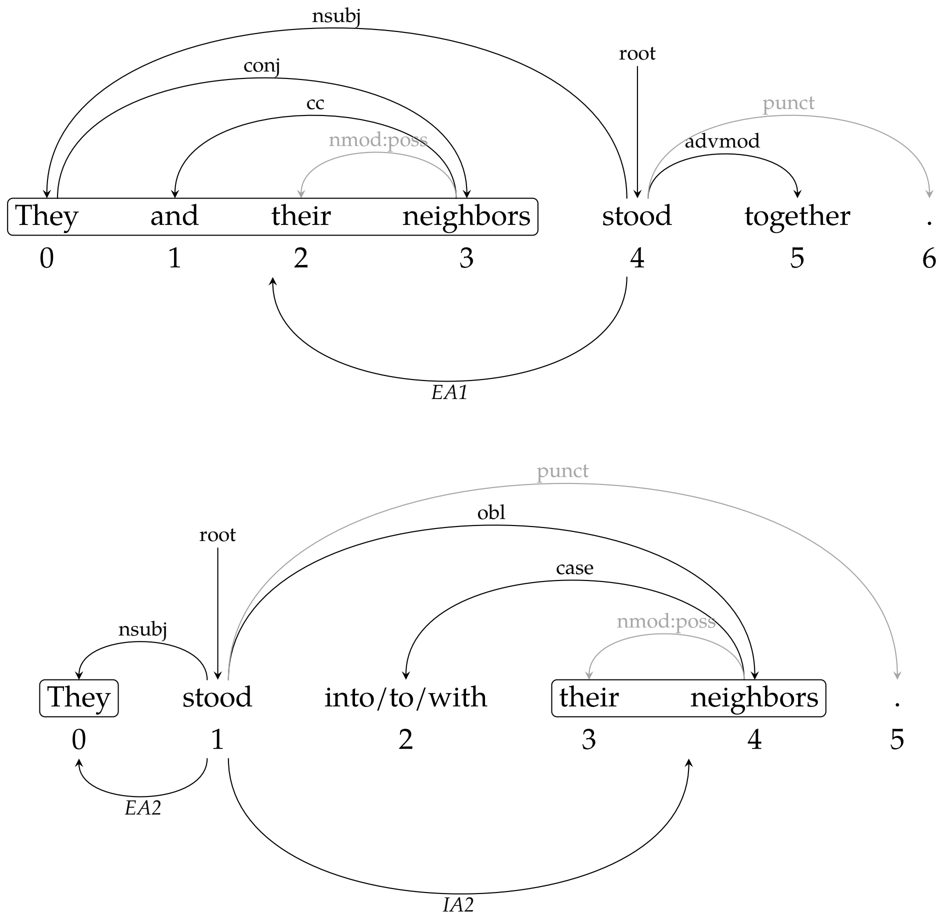


Figure 80: Dependency trees for the original sentence #*ENCOW-01-06803226* (top) and its reordered alternate version with respect to the together reciprocal alternation (intransitive) (bottom).

contains some variations with regard to the 1ARSA pattern. Instead of requiring all sentences in the *together* subcorpus to have a conjunct `nsubj`, all sentences that involve a *together* token in the required position are included, independently of the contents of their `nsubj` phrase. This leads to the inclusion of three main types of sentences: those including a conjunct `nsubj`; those including a single `nsubj` that is filled by an entity in plural form; and those including a single `nsubj` that is filled by an entity in singular form.

The first type, sentences with conjunct `nsubj` phrases, are reordered as in the examples given for the alternation by Levin (1993), as illustrated in Figure 80. Sentences with a plural `nsubj` token are reordered by keeping the `nsubj` phrase in the same position and adding *each other* in the prepositional phrase, as illustrated in Figure 81 on page 259. Sentences with a singular `nsubj` token are reordered by keeping the `nsubj` phrase in the same position and selecting a random pronominal prepositional object slot filler based on corpus observations, as illustrated in Figure 82 on page 259.

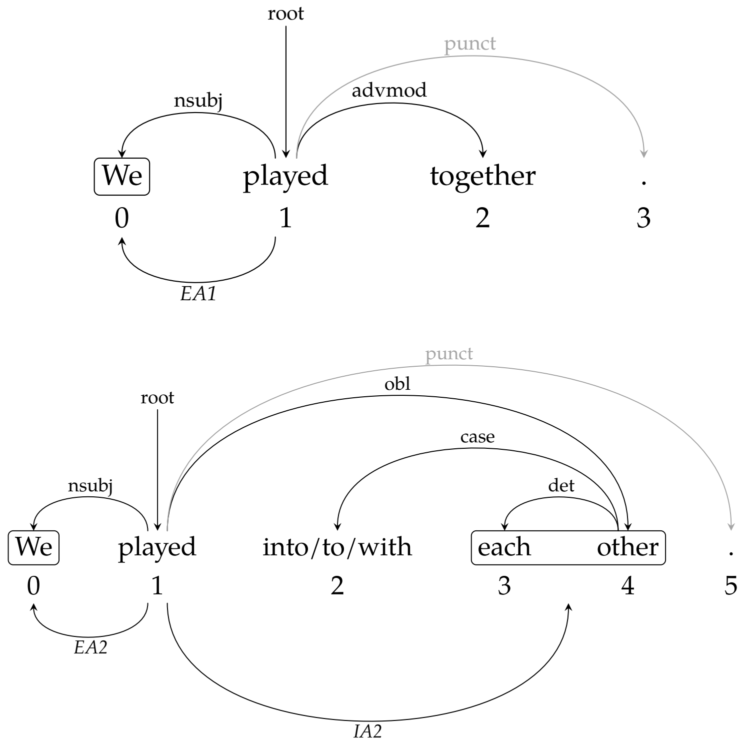


Figure 81: Dependency trees for the original sentence *#ENCOW-01-00169685* (top) and its reordered alternate version with respect to the together reciprocal alternation (intransitive) (bottom).

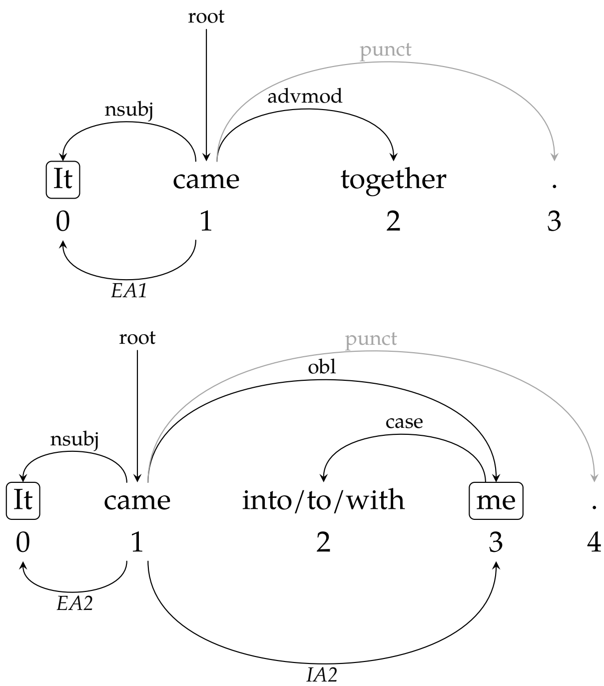


Figure 82: Dependency trees for the original sentence *#ENCOW-01-00923158* (top) and its reordered alternate version with respect to the together reciprocal alternation (intransitive) (bottom).

Table 18: Properties of ENCOW sentences that lead to the exclusion of sentences from the experiments conducted here.

| Alternation | Pattern | Reasons for exclusion |
|---|---|---|
| (all) | (all) | contains a `:`<br>contains a `?`<br>contains punctuation tokens that have dependencies directly under them<br>contains punctuation tokens that do not depend directly on `root` |
| CIA | transitive | `obj` has a direct dependency under it of the type `case`, `nsubj`, `obj`, `aux`, `nsubj:pass`, or `cop` |
| | intransitive | `nsubj` has a direct dependency under it of the type `case`, `nsubj`, `obj`, `aux`, `nsubj:pass`, or `cop` |
| ISA | using | – |
| | with | – |
| | instrument subject | `nsubj` has a direct dependency under it of the type `case`, `nsubj`, `obj`, `aux`, `nsubj:pass`, `acl`, or `cop` |
| TRAI | into | – |
| | to | – |
| | with | – |
| | together | – |

Some types of sentences are more challenging to reorder than others. For instance, if a sentence has an `nsubj` phrase that contains a relative clause, this leads to more tokens whose agreement features need to be checked and possibly repaired. In order to avoid defining an overly complicated set of rules for the generation of alternate sentences, original sentences were filtered during the subcorpus collection step to ensure only "reorderable" sentences are being processed. Table 18 details the properties that lead to sentences being discarded. This corpus reduction step is conducted prior to all alternation identification experiments, in order to make the performance of the different approaches comparable; the corpus sizes given in Table 4 on page 216 already take reorderability into account. While filtering for reorderability in some cases leads to the loss of more than 50% of sentences, the goal is to work with a smaller but more informative set of sentences.

#### 10.5.4.2 *Perplexity scores for individual sentences*

Perplexity values are derived from log-likelihood values for the given sequence. This measure can only be applied to autoregressive language models whose objective is to predict the next token in a sequence (Radford et al. 2018), and not to masked language models like BERT (Devlin et al. 2019). Popular language models in this category at the time of writing are Transformer-XL (Dai et al. 2019), XL-NET (Yang et al. 2019), and GPT-2 (Radford et al. 2019). A key strength of Transformer-XL lies in its ability to learn long-distance dependencies and to generate long, reasonably coherent texts; these factors are not the main issue for the experiments conducted in this thesis, since only relatively short, isolated sentences will be processed. XLNet is an autoregressive architecture that builds on Transformer-XL and learns bidirectional contexts. GPT-2, which is trained with a causal language modeling objective, has the key strength of generating syntactically coherent text.

The ability of GPT-2 to learn syntactic coherence promises useful insights for the distinction between "good" generated sentences and "bad" ones. Therefore, the language model to be used for the perplexity-based classification features will be GPT-2.

The underlying assumption for this approach to alternation identification is that verbs that participate in a given alternation will lead to more acceptable generated alternate sentences than their non-alternating counterparts. Once all alternate versions for the observed sentences in the alternation-specific corpora are generated, they are processed with GPT-2, and the perplexity scores of each original sentence and its reordered alternate version are stored.

The perplexity scores are collected using the pretrained `GPT2LMHead Model` class from the `transformers` module for Python by *Huggingface*. The scripts to collect the scores use version 4.60.0 of `tqdm`, version 1.9.0 of `torch`, 4.8.2 of `transformers`, and version 3.8.3 of Python.

#### 10.5.4.3 *Feature definition*

Each sentence's perplexity score by itself is not informative as to whether the sentence's root verb participates in the alternation or not. The perplexity of originally observed sentences from the corpus is expected to vary within a certain range, so the perplexity of a generated alternate sentence should always be interpreted with regard to its difference to the original sentence's perplexity. Additionally, even though the perplexity measure normalizes for the number of tokens in the given sequence, longer sequences or more complex phrases may lead to disproportionately smaller log-likelihood scores from the language model, which result in higher perplexity values. Finally, for the alternations in which one pattern type contains one argument more than the other (the causative-inchoative alternation and the instrument subject alternation), the process of deleting a phrase or adding one during the transformation step will systematically lead to a certain difference in perplexity in addition to the differences that are due to the verb's alternation behavior.

Instead of the individual, sentence-level perplexity score, what is relevant for this approach to alternation identification are the trends in perplexity deltas that can be observed across different sentences involving the same root verb. Perplexity essentially expresses how "surprised" a language model is about a given sequence. If a language model has access to sufficient data from the language under investigation, a generated alternate sentence with a non-alternating verb should "surprise" the model more than a generated alternate sentence with an alternating verb. The generated sentence with the alternating verb should more closely resemble what the model expects a sequence to look like than a sentence that is ungrammatical or that involves an unusual argument in a certain position.

FEATURES 03-01 AND 03-02: *Average perplexity of 1ARSA and 2ARSA pattern sentences.* The first two perplexity-based features for the classification simply express the average perplexity of all observed sentences in the 1ARSA/2ARSA pattern type for the current alternation and the current verb. If a verb appears in both pattern types, but all instances of one of the patterns are assigned high perplexity values, this may be sufficient evidence to label that verb as not participating in the given alternation.

FEATURE 03-03: *Average perplexity difference between original 1ARSA pattern type and generated 2ARSA pattern type.* This feature expresses the average difference in perplexity between observed instances of a given verb in the 1ARSA pattern type of an alternation, and the generated counterparts for the observed sentences that instantiate the 2ARSA pattern. While differences in perplexity are unavoidable, the underlying assumption is that the generated sentences for non-alternating verbs should lead to higher average differences in perplexity, since the resulting generated sentences are expected to be more "surprising" than the generated sentences for verbs that do participate in the current alternation.

In order to reduce the impact of word choice or other properties of each sentence on the feature values, the feature value will represent the relative difference in perplexity values between the original sentence and its counterpart.

FEATURE 03-04: *Average perplexity difference between original 2ARSA pattern type and generated 1ARSA pattern type* Since neither the 1ARSA nor the 2ARSA pattern for an alternation are assumed to be "underlying" forms, the average perplexity difference feature is implemented in both directions. There is no underlying assumption as to which of the patterns is more likely to be acceptable, so each transformation of an observed sentence to its respective alternate pattern(s) is a possible relevant indicator for the ability of the verb to participate in the alternation. This approach may be particularly informative for alternations with fewer attestations of the different patterns; if the perplexity features do approximate alternatability as described above, then each observed sentence in one of the patterns is a useful data point, even if no direct counterpart is observed in the corpus – the counterparts used for these features are generated, not collected from the corpora.

For the causative-inchoative alternation, each pattern type covers exactly one pattern, so the calculation of values for the perplexity features as described above is straightforward. For the instrument subject alternation and the together reciprocal alternation (intransitive), there are pattern types that include more than one pattern. In the case of the instrument subject alternation, the average perplexity difference between original and generated sentences should take all subpatterns that belong to the current pattern type into account: in order to determine the value of feature 03-03 for candidate verbs for the instrument subject alternation, all generated alternate sentences that belong to the *with*-pattern are being taken into account, as well as all generated alternate sentences that belong to the *using*-pattern. In the other direction, to determine how perplex observed agent subject sentences for the instrument subject alternation are, all *with* sentences and all *using* sentences are viewed as instances of that pattern type, and the generated instrument-subject alternate versions for each of those original sentences are all used to determine the average perplexity difference.

This treatment is ideal for the instrument subject alternation because the constructions that license the *using* subpattern and the *with* subpattern are equally able to express an instrument in the IA2 slot, although the polysemy of *with* means that elements that appear in the *with* slot are not necessarily instruments.

For the together reciprocal alternation (intransitive), on the other hand, each verb that participates in the alternation does so with a specific preposition or set of prepositions. For instance, Levin (1993: 64) lists the verb *blend* as participating in the together reciprocal alternation (intransitive) with the preposition *into* and the preposition *with*, while the verb *fuse* participates with the prepositions *to* and *with*, and the verb *collect* is only listed as participating in the alternation with the preposition *into*. This means that the classifier should not expect a single verb to be able to appear with all three prepositions in a corpus, or to lead to acceptable sentences when observed sentences of the *together* pattern are reordered to the three different prepositional alternate patterns.

Which subpatterns should be taken into account to determine for an alternation candidate whether it is likely to participate in the together reciprocal alternation (intransitive) or not? If only the most frequent prepositional pattern is taken into account, this approach neglects the verbs that are able to appear in the 2ARSA pattern with more than one preposition, such as *blend* or *fuse*. At the same time, only taking into account the prepositional pattern whose average perplexity is nearest to the average perplexity of the alternate sentences in the 1ARSA pattern is an undesirable simplification.

Instead, all generated 2ARSA pattern sentences for a candidate verb will be pooled together to determine the value for feature 03-03, and all original 2ARSA pattern sentences for a candidate verb will be pooled together to determine the value for feature 03-04. While this will include sentences that exhibit the prepositional pattern for reasons unrelated to the alternation, it is the best way to use the available data as a basis for the classification, without requiring individual verbs to show a tendency as to which preposition fits them the best. Moreover, the data may show that individual verbs can actually appear in more prepositional patterns than indicated by Levin; for instance, the verb *bond* is listed there as taking the preposition *to* and none of the other prepositions, but it is observed in the corpus data both in the *to* subpattern and in the *with* subpattern. Thus, taking all prepositional subpatterns into account is the way the values for the perplexity features for this alternation will be derived.

Table 19 on page 264 lists all perplexity-based features for the classification. A perplexity value of 0 would mean that the given sequence is maximally probable according to the language model, that is, that the sequence is exactly what the language model would predict based on its training data. Since this is an unlikely scenario, only perplexity values above 0 are expected. The feature values have no upper limit: a generated alternate sentence may be so improbable that the language model assigns a perplexity value to it that is orders of magnitude higher than that of the original sentence. For the classification step, all feature values for the current set of verb candidates will be normalized to a range between 0 and 1.

### 10.5.5 *Results*

The results of the perplexity-based feature sets on the full verb sets for each of the three alternations are shown in Table 20 on page 264.

Table 19: Overview of all perplexity-based features for the classification of candidate verbs for each of the three alternations. Feature sets 03a and 03b will be evaluated separately.

| Perplexity-based feature sets |
|---|
| 03a-perplexity-features-original |
| 03-01_average_pplx_1ARSA |
| 03-02_average_pplx_2ARSA |
| 03b-perplexity-features-reordered-pplx-difference |
| 03-03_average_pplx_difference_original_1ARSA_turned_to_2ARSA |
| 03-04_average_pplx_difference_original_2ARSA_turned_to_1ARSA |

Table 20: Scores (accuracy, precision, recall, $F_1$) for the perplexity-based feature sets.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 03a-perplexity-features-original | 51.74 | 50.74 | 82.23 | 62.49 |
| ISA | 03a-perplexity-features-original | 54.14 | 55.46 | 41.37 | 47.39 |
| TRAI | 03a-perplexity-features-original | 43.83 | 35.00 | 49.00 | 32.97 |
| CIA | 03b-perplexity-features-reordered-pplx-difference | 55.11 | 53.28 | 80.12 | 64.02 |
| ISA | 03b-perplexity-features-reordered-pplx-difference | 48.42 | 44.82 | 35.28 | 34.51 |
| TRAI | 03b-perplexity-features-reordered-pplx-difference | 62.50 | 55.00 | 69.50 | 58.23 |

Once more, the alternation with the best $F_1$ score for each of the perplexity-based feature sets is the causative-inchoative alternation.

Surprisingly, the feature set comparing each original sentence's perplexity score to the perplexity score of its reordered counterpart performs worse for the instrument subject alternation than the feature set containing merely the original perplexity values for observed sentences. In other words, reordering sentences and scoring the resulting sentences does not provide any additional information to the classifier that helps it to determine whether a verb participates in the instrument subject alternation.

Expressing an instrument that is involved in an event is often optional, especially when the verb used to denote the event is associated with one or more typical instruments. At the same time, the preposition *with* is highly polysemous. As a result, many of the sentences collected from ENCOW that presumably instantiate the *with* construction of the instrument subject alternation may in fact be unrelated to that construction, even for verbs that are known to participate in the alternation. In other words, the entity expressed in the IA2 slot of such sentences can fill roles like theme or patient, as well as instrument. Then, moving these entities to the EA1 position yields sentences with high perplexity values, even if the verb participates in the instrument subject alternation.

Another effect that impacts the feature values occurs when the entity in the IA2 slot refers to an entity of a similar type as entities in the EA1 or EA2 slot, for instance because the two arguments are co-themes. This is licensed, for instance, by the construction associated with the 2ARSA pattern type of the together reciprocal alternation (intransitive). In this scenario, reordering a sentence instantiating the 2ARSA pattern of the instrument subject alternation into the 1ARSA pattern of the instrument subject al-

Table 21: Scores (accuracy, precision, recall, $F_1$) for the perplexity-based feature sets, with a minimum alternation candidate frequency of 20.

| | Feature set | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA | 03a-perplexity-features-original | 47.71 | 47.78 | 49.80 | 48.76 |
| ISA | 03a-perplexity-features-original | 53.43 | 54.08 | 46.09 | 49.68 |
| TRAI | 03a-perplexity-features-original | 63.50 | 60.00 | 90.00 | 72.33 |
| CIA | 03b-perplexity-features-reordered-pplx-difference | 52.71 | 51.96 | 83.72 | 63.93 |
| ISA | 03b-perplexity-features-reordered-pplx-difference | 51.64 | 50.85 | 80.81 | 62.48 |
| TRAI | 03b-perplexity-features-reordered-pplx-difference | 72.50 | 63.50 | 81.00 | 71.00 |

ternation will result in a low-perplexity sentence even if the verb does not participate in this alternation.

For the together reciprocal alternation (intransitive), the perplexity-based features lead to surprisingly strong accuracy and $F_1$ scores. It is possible that the perplexity features remedy some of the sparsity issues mentioned in the discussion of the results in the previous two sections, which is particularly relevant for an alternation with as small a candidate set and subcorpora as the together reciprocal alternation (intransitive).

Table 21 shows how the perplexity-based features perform on the reduced verb sets with a minimum number of 20 attestations per verb.

Discarding infrequent verbs has a minor effect on the classifier scores for the causative-inchoative alternation. For the instrument subject alternation, the feature set encoding the perplexity difference between original and reordered sentences yields slightly higher scores when the verb candidate set only includes verbs above the minimum frequency threshold. Finally, for the together reciprocal alternation (intransitive), reducing the verb set to only those that appear at least 20 times causes a steep climb in the classification score for all perplexity-based feature sets.

Overall, the perplexity features are most beneficial for the identification of verbs that participate in the alternation with the smallest number of participating verbs and with the smallest subcorpora available.

## 10.6 DISCUSSION

The previous sections of this chapter discussed a baseline and three types of linguistically motivated feature sets for the identification of verbs that participate in three English verb alternations. In this section, the results will be compared and interpreted with a focus on each alternation individually. A discussion of the results with respect to the hypotheses formulated in Section 9.1 will follow in Section 10.6.2.

### 10.6.1 *Alternation-specific discussion*

Table 22 on page 266 shows the results of the baseline and the feature sets on the full verb set for the causative-inchoative alternation. Table 23 on page 266 presents the scores for the same alternation with a minimum candidate frequency threshold of 20.

Table 22: Scores (accuracy, precision, recall, $F_1$) for CIA from each feature set.

| Scores for CIA | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 57.56 | 57.35 | 60.74 | 58.63 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 61.13 | 67.20 | 42.84 | 52.46 |
| 01b-lemma-features-token-based | 60.47 | 68.09 | 40.61 | 50.83 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 59.34 | 60.82 | 55.01 | 57.68 |
| 02b-vector-features-glove-with-placeholder | 58.25 | 59.50 | 54.21 | 56.51 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 61.74 | 64.79 | 53.45 | 58.16 |
| 02d-vector-features-fasttext-with-placeholder | 61.67 | 65.39 | 52.60 | 57.83 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 51.74 | 50.74 | 82.23 | 62.49 |
| 03b-perplexity-features-reordered-pplx-difference | 55.11 | 53.28 | 80.12 | 64.02 |

Table 23: Scores (accuracy, precision, recall, $F_1$) for CIA from each feature set, with a minimum alternation candidate frequency of 20.

| Scores for CIA (minfreq 20) | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 63.79 | 64.70 | 63.10 | 63.21 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 62.06 | 64.63 | 53.66 | 58.09 |
| 01b-lemma-features-token-based | 59.29 | 62.80 | 49.51 | 54.58 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 62.19 | 62.99 | 62.46 | 63.13 |
| 02b-vector-features-glove-with-placeholder | 61.31 | 62.42 | 60.64 | 60.81 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 61.50 | 62.89 | 58.96 | 60.21 |
| 02d-vector-features-fasttext-with-placeholder | 62.34 | 63.08 | 58.70 | 60.58 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 47.71 | 47.78 | 49.80 | 48.76 |
| 03b-perplexity-features-reordered-pplx-difference | 52.71 | 51.96 | 83.72 | 63.93 |

For both conditions, the feature set with the highest $F_1$ score on the classification task for the causative-inchoative alternation is feature set 03b, where each sentence's original perplexity is compared to the perplexity of its generated alternate form instantiating the other pattern of the alternation. This method augments observed data for each verb from the corpus with a generated counterpart for each relevant instance of a pattern.

The $F_1$ score for the frequency-based baseline is almost as strong as the $F_1$ score for feature set 03b for the condition with a minimum frequency threshold. The baseline also achieves the strongest accuracy score in this setting.

The strong performance of the baseline for this alternation may be due to the way the alternation-specific subcorpora have been collected. For this alternation, all sentences were selected which contain either only an `nsubj` relation, or an `nsubj` relation and an `obj` relation at the top level of the dependency tree. A subset of the negative verbs that appear in both patterns are object-drop verbs. However, due to the nature of the ENCOW corpus, there are also sentences that are incomplete, for instance, sentences whose verb theoretically requires a direct object, but which contain no object because they are cut off in the middle. On the other hand, there are also very long sentences which are misparsed, such that the dependency tree incorrectly contains a direct object. For these cases, one of the relevant patterns will be attested much less frequently than for verbs participating in the causative-inchoative alternation or object-drop verbs. In other words, the baseline seems to be successful at filtering out verbs for which the syntactic patterns are linguistically motivated *somehow*, even though this motivation is not always founded in the verb's participation in the causative-inchoative alternation.

Almost all feature sets benefit from the minimum candidate frequency threshold of 20 instances. The only exception is feature set 03a, where the average perplexity of all original sentences in each pattern of the alternation is used directly as a feature value. Since there are multiple competing reasons for sentences to instantiate the 1ARSA or 2ARSA pattern of the causative-inchoative alternation, it is not surprising that simply approximating the acceptability of the patterns for a verb is not informative enough to decide whether that verb participates in the alternation.

Feature set 03b achieves a particularly high recall in both conditions, compared to the other features. It is thus more successful at identifying positive verbs correctly than other feature sets, whereas its precision tends to be lower than that of other feature sets. However, in the context of these experiments, the negative set may contain verbs that can participate in the alternation and should be moved to the positive set. For the purposes of identifying these new alternation candidates, a higher recall and lower precision is actually desirable, since it yields a certain number of "false positives" which can then be relabeled manually. For more on the usefulness of false positives for the identification of new alternation candidates, see Section 10.7.

Tables 24 and 25 on page 268 present the scores of each feature set for the instrument subject alternation.

This alternation also tends to benefit from the minimum candidate frequency threshold of 20 instances, but only by a small margin for most feature sets. The biggest difference between the two settings is observed for feature set 03b, which is fairly weak for the full candidate set but achieves the highest $F_1$ score out of all feature sets in the frequency-filtered condition.

Table 24: Scores (accuracy, precision, recall, $F_1$) for ISA from each feature set.

| Scores for ISA | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 52.70 | 51.90 | 54.97 | 53.51 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 52.31 | 51.62 | 63.05 | 55.49 |
| 01b-lemma-features-token-based | 54.90 | 53.95 | 71.70 | 61.22 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 53.37 | 53.57 | 57.25 | 54.70 |
| 02b-vector-features-glove-with-placeholder | 53.14 | 53.26 | 54.05 | 53.31 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 53.61 | 53.58 | 52.93 | 52.51 |
| 02d-vector-features-fasttext-with-placeholder | 53.42 | 53.60 | 54.46 | 54.45 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 54.14 | 55.46 | 41.37 | 47.39 |
| 03b-perplexity-features-reordered-pplx-difference | 48.42 | 44.82 | 35.28 | 34.51 |

Table 25: Scores (accuracy, precision, recall, $F_1$) for ISA from each feature set, with a minimum alternation candidate frequency of 20.

| Scores for ISA (minfreq 20) | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 56.84 | 55.90 | 63.65 | 59.13 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 55.73 | 56.37 | 54.07 | 53.34 |
| 01b-lemma-features-token-based | 58.24 | 58.36 | 58.45 | 57.70 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 51.68 | 52.47 | 55.38 | 53.27 |
| 02b-vector-features-glove-with-placeholder | 52.95 | 52.74 | 58.03 | 54.16 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 52.88 | 52.11 | 59.20 | 54.58 |
| 02d-vector-features-fasttext-with-placeholder | 50.75 | 51.13 | 56.56 | 53.08 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 53.43 | 54.08 | 46.09 | 49.68 |
| 03b-perplexity-features-reordered-pplx-difference | 51.64 | 50.85 | 80.81 | 62.48 |

On the full candidate set, the strongest $F_1$ score is achieved by feature set 01b, which only uses slot filler lemma overlap as an indicator for shared selectional preferences. Surprisingly, the vector-based feature sets, which are meant to generalize over individual observed lemmas in each slot, do not improve over the scores of the lemma-based feature sets. This may be due to the heterogeneous nature of the slot filler sets for the slots expressing an instrument.

Additionally, as discussed previously, the 1ARSA pattern and 2ARSA pattern associated with the instrument subject alternation may each instantiate constructions unrelated to the alternation, even when they appear with a verb that participates in the alternation. The 1ARSA pattern may express an agent in the subject position instead of an instrument, and the 2ARSA pattern may involve a sense of *with* that does not encode an instrument. Because the classification takes all observed occurrences of each pattern into account, the classifiers suffer from conflicting signals, which explains the relatively low classification scores for the instrument subject alternation across all feature sets.

Again, the best-scoring feature sets in both conditions achieve a relatively high recall. As mentioned above for the causative-inchoative alternation, this can be useful for the purposes of identifying new alternation candidates.

Tables 26 and 27 on page 270 present the scores of each feature set for the together reciprocal alternation (intransitive).

For most feature sets, the classification results for the together reciprocal alternation (intransitive) are much weaker than for the previous two alternations. The baseline is associated with the lowest $F_1$ score in both settings.

The lemma-based feature sets achieve better scores than the vector-based ones. This is counter-intuitive, since the vector features are meant to group slot filler lemmas together to approximate their semantic category in a way that is more robust against sparsity than the lemma-based features. However, the set of candidate verbs for the together reciprocal alternation (intransitive) is fairly small, and the verbs in the set are each attested fairly infrequently in the corpus. The numbers of observed slot fillers for the different slots may be too small to benefit from the generalization power of the distributional approach.

The strongest result for the together reciprocal alternation (intransitive) in both settings is achieved by feature set 03b. For this alternation, this feature set achieves both the highest accuracy score and the highest $F_1$ score in both conditions. The strength of this feature set is most likely due to the way it augments observed data points with generated alternate sentences, which somewhat reduces the sparsity issues.

The syntactic patterns of the constructions associated with the together reciprocal alternation (intransitive) may appear for a range of reasons, since prepositional phrases headed by *with*, *into* or *to* can simply be added to sentences with alternating or non-alternating verbs as adjuncts, as can the adverb *together*. In such cases, the entities in the presumed co-theme slots are not necessarily themes. Thus, similarly to the instrument subject alternation, the syntactic patterns associated with the together reciprocal alternation (intransitive) may frequently be realizations of unrelated constructions, even for verbs that do participate in the alternation. This is one possible explanation for the low classification scores achieved by most feature sets for this alternation.

Table 26: Scores (accuracy, precision, recall, $F_1$) for TRAI from each feature set.

| Scores for TRAI | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 30.33 | 8.67 | 23.00 | 10.17 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 51.67 | 42.00 | 50.00 | 40.70 |
| 01b-lemma-features-token-based | 36.17 | 31.83 | 51.00 | 40.23 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 31.83 | 15.83 | 24.50 | 19.83 |
| 02b-vector-features-glove-with-placeholder | 36.00 | 28.83 | 42.50 | 27.43 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 41.67 | 29.00 | 53.50 | 37.40 |
| 02d-vector-features-fasttext-with-placeholder | 33.67 | 22.83 | 37.50 | 27.47 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 43.83 | 35.00 | 49.00 | 32.97 |
| 03b-perplexity-features-reordered-pplx-difference | 62.50 | 55.00 | 69.50 | 58.23 |

Table 27: Scores (accuracy, precision, recall, $F_1$) for TRAI from each feature set, with a minimum alternation candidate frequency of 20.

| Scores for TRAI (minfreq 20) | A | P | R | $F_1$ |
|---|---|---|---|---|
| 00-baseline | 39.50 | 5.00 | 11.00 | 6.67 |
| Lemma-based | | | | |
| 01a-lemma-features-type-based | 63.50 | 49.50 | 63.00 | 55.33 |
| 01b-lemma-features-token-based | 46.00 | 36.50 | 54.00 | 43.67 |
| Vector-based | | | | |
| glove | | | | |
| 02a-vector-features-glove-without-placeholder | 26.50 | 20.00 | 28.00 | 23.33 |
| 02b-vector-features-glove-with-placeholder | 27.00 | 17.00 | 34.00 | 26.00 |
| fasttext | | | | |
| 02c-vector-features-fasttext-without-placeholder | 33.50 | 18.00 | 41.00 | 30.00 |
| 02d-vector-features-fasttext-with-placeholder | 48.00 | 46.50 | 65.00 | 49.00 |
| Perplexity-based | | | | |
| 03a-perplexity-features-original | 63.50 | 60.00 | 90.00 | 72.33 |
| 03b-perplexity-features-reordered-pplx-difference | 72.50 | 63.50 | 81.00 | 71.00 |

Table 28: Scores (accuracy, precision, recall, $F_1$) for the best-performing feature set for each alternation, using the best-performing hyperparameter values for each alternation.

| Alternation, parameters | Features | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA, C=100, $\gamma$=0.1 | 02c | 62.67 | 65.22 | 55.51 | 60.15 |
| ISA, C=10,000, $\gamma$=0.01 | 02c | 55.34 | 54.50 | 58.60 | 56.33 |
| TRAI, C=100, $\gamma$=10 | 03a | 68.5 | 58.5 | 70.00 | 62.73 |

Table 29: Scores (accuracy, precision, recall, $F_1$) for the best-performing feature set for each alternation, using the best-performing hyperparameter values for each alternation, with a minimum alternation candidate frequency of 20.

| Alternation, parameters | Features | A | P | R | $F_1$ |
|---|---|---|---|---|---|
| CIA, C=100, $\gamma$=0.1 | 02c | 64.33 | 65.47 | 60.58 | 62.11 |
| ISA, C=10, $\gamma$=10 | 01a | 56.92 | 59.73 | 46.66 | 50.89 |
| TRAI, C=100, $\gamma$=10 | 03a | 74.50 | 60.00 | 74.00 | 64.33 |

The classification scores for each alternation can be improved further by tuning the hyperparameters (C and $\gamma$). The scores reported above are all achieved with the default settings for these parameters. Setting the values for the hyperparameters per-alternation leads to additional gains in accuracy. Table 28 reports the best scores for each alternation with optimal hyperparameters (chosen via grid search, with a range of C values between 0.001 and 10,000 and a range of $\gamma$ values between 0.001 and 1,000, selected for highest classification accuracy).

Table 29 reports the classification scores for each alternation after hyperparameter tuning on the verb set with a minimum frequency threshold of 20.

In contrast to the findings reported by Merlo & Stevenson (2001), combining the different feature types does not improve the classification results here. Apparently, the strengths of the strongest feature set for each alternation are canceled out by the weaknesses of the other feature sets, leading to overall scores below those of the best feature set when used in isolation.

Concerning possible improvements for the classification results, several directions for future work seem promising. As mentioned multiple times throughout this chapter, there are some issues with the quality of the available dependency annotations in the ENCOW corpus. The corpus contains user-generated and unedited text and is therefore not "clean"; using a less noisy, dependency-annotated corpus containing, for instance, news text or published literature may change the feature values for each verb to some degree, making it easier for the classifier to label candidates correctly. Additionally, having access to more instances of the relevant constructions per verb would reduce the impact of sparsity and make the classifiers more robust.

The classification settings with a minimum candidate frequency threshold of 20 tend to lead to better classification results. This illustrates that infrequent verbs are particularly difficult to label correctly. Depending on the goal of the classification, infrequent verbs can either be discarded or left in the candidate set.

However, as discussed by Merlo & Stevenson (2001), it is unlikely for any approach to alternating verb identification to achieve perfect agreement with the gold standard. This is because, as Merlo & Stevenson (2001: 395) report, even linguistically-trained experts do not always agree on alternation participation of individual verbs. Ambiguity and the lack of discourse context in this setup complicate this issue further. For more on these issues, see Section 10.7 of this thesis.

Recall that some related work on alternating verb identification used verbs that cannot appear in the relevant syntactic patterns as the negative examples for the classifier training (see Section 9.2 of this thesis). The task setup implemented in this thesis is much more challenging, since only verbs that are observed at least once in each relevant pattern type are used as candidates. The goal of the classification was to determine to what extent information about argument slot fillers can be used to distinguish alternating verbs from non-alternating verbs.

### 10.6.2 *Discussion of hypotheses H1 through H4*

This section will go through the four hypotheses presented in Section 9.1 and determine whether they are supported by the findings from the classification experiments.

H1. *Verbs that participate in alternations whose constructions have generally frequent syntactic patterns are harder to identify than verbs that participate in alternations with less frequent syntactic patterns.* This hypothesis is not supported by the classification results reported in this chapter. The identification of verbs in the alternation with the least frequent syntactic patterns, the together reciprocal alternation (intransitive), suffered from sparsity to a greater extent than the other two alternations. This alternation received the lowest classification scores for the majority of feature sets. The alternation with the most frequent syntactic patterns, the causative-inchoative alternation, often received the best or second-best scores for each feature set. Even though the patterns licensed by the constructions of the causative-inchoative alternation can appear for several other reasons that are unrelated to the alternation, the sheer number of observed attestations of these patterns for each positive and negative verb seemingly compensated for the ambiguity of the patterns with respect to different possible constructions.

H2. *Using only the number of attestations of certain syntactic patterns for a verb in a corpus is a weak strategy to identify verbs in any alternation; it may be more successful for alternations with less frequent patterns than for alternations with frequent syntactic patterns.* This hypothesis was tested with the frequency-based baseline, which used the ratio of 1ARSA pattern instances to 2ARSA pattern instances as the sole indicator for or against alternation participation. Surprisingly, the baseline achieves strong scores for the causative-inchoative alternation, which is an alternation with fairly frequent syntactic patterns, while the baseline scores for the together reciprocal alternation (intransitive), which is associated with fairly infrequent patterns, are very weak. Thus, hypothesis H2 is only partially supported by the findings reported in this thesis.

H3. *Alternating verb identification is improved by features that approximate the selectional preferences of each candidate verb for its argument slots in the relevant syntactic patterns.* The lemma-based and vector-based features fall into this category. For the causative-inchoative alternation, the vector-based features yield better $F_1$ scores than the lemma-based ones, but neither of them outperform the frequency-based baseline. For the instrument subject alternation, the $F_1$ score of the vector-based features is similar to that of the frequency-based baseline, but the lemma-based features perform slightly better. The best $F_1$ score for the instrument subject alternation on the unfiltered verb set is achieved by the lemma-based feature set 01b. On the verb set with a minimum candidate frequency of 20, the baseline outperforms both lemma-based and vector-based features in terms of $F_1$ score for this alternation. For the together reciprocal alternation (intransitive), both lemma-based and vector-based features outperform the frequency-based baseline, with the lemma-based features achieving higher $F_1$ scores than the vector-based ones. In sum, the features approximating selectional preferences are beneficial for the instrument subject alternation (on the full candidate set) and the together reciprocal alternation (intransitive), but do not outperform the frequency-based baseline for the causative-inchoative alternation. Hypothesis H3 is partially supported by the classification results.

H4. *Features that approximate the acceptability of an alternation's constructions for a candidate verb can lead to better classification results, especially for alternations whose patterns are attested infrequently.* The perplexity-based features fall into this category. In terms of $F_1$ score, these features yield the best results for the causative-inchoative alternation on the unfiltered candidate set. For the instrument subject alternation, they lead to low $F_1$ scores on the unfiltered candidate set but to the highest score for the condition involving a minimum frequency threshold for candidates. For the together reciprocal alternation (intransitive), the perplexity-based features lead to the highest $F_1$ scores in both settings. The verbs in the candidate set for the together reciprocal alternation (intransitive) are attested very infrequently compared to the other two alternations. Augmenting observed instances with their alternate counterparts thus leads to the best classification gains for the least frequent alternation. Thus, hypothesis H4 is supported by the classification results.

Overall, these results show that different verb alternations benefit from different approaches to the classification task. None of the feature sets was consistently accurate or consistently weak across all alternations. The feature set that yielded good classification scores in most settings was 03b, where an approximation of acceptability for observed sentences is compared to an approximation of acceptability for their generated alternate forms.

The perplexity-based features are easy to apply to additional alternations. For instance, the spray-load alternation licenses sentences like those shown in (75). The goal of the *loading* event (the *bucket*) is expressed as the direct object in (75a), with the theme (the *spermaceti oil*) being expressed as a prepositional object. In (75b), the positions are swapped. A simple rule-based script can be created to transform sentences like (75a) into their counterparts as in (75b) and vice versa. Then, perplexity scores can be used as a source for feature values as implemented in this chapter.

(75) a. Tashtego loads the bucket with spermaceti oil.
b. Tashtego loads spermaceti oil into the bucket.

Although some alternations covered in this chapter had strong baseline scores, the perplexity-based features are a better choice for the identification of verbs in additional alternations not yet covered here, especially since they proved to be useful for infrequent candidates.

## 10.7 EXTENDING THE POSITIVE VERB SETS BASED ON FALSE POSITIVES

So far, the sets of alternating verbs have been discussed as if they were fixed and completely correct and comprehensive, and the sets of negative sets have been treated as if they were reliably examples of verbs that do not participate in the relevant alternation. However, as mentioned in Section 9.6, the resource from which the positive verb sets were sourced may not be comprehensive. In other words, there may be more verbs that participate in each alternation than are contained in the available positive set.

This invites a new perspective on the classification results reported in this chapter. Instead of attempting to predict the assigned positive or negative label for each candidate for an alternation, the classifiers can be used to search for new likely alternating verbs that are incorrectly included in the negative set.

In this context, the positive verb sets are still assumed to be correct, that is, all members of the positive set for an alternation are assumed to truly participate in that alternation. The positive sets, unlike the negative ones, are sourced from a hand-crafted resource and have been checked by multiple linguists throughout the creation of that resource. The negative sets were created based on the subcorpora collected for the purposes of the experiments presented in this chapter, and only by virtue of these verbs not being listed in the positive sets, but appearing in the corpus in the relevant syntactic patterns for their alternation. Thus, there is a possibility for each verb in the presumed negative set to actually be a valid member of the positive set instead. This section presents a process by which presumed-negative verbs can be relabeled to the positive category based on the classifier predictions.

The general idea is as follows. The alternation candidates that are of special interest for the goal of extending the positive set are the verbs that appear in the presumed negative set. Out of the negative verbs, the ones that are most likely to be good alternation candidates are those that are "mislabeled" by a good classifier as positive instances. After all, if a classifier recognizes that a given verb behaves very similarly to known alternating verbs, then the positive label it assigns to the current candidate may turn out to be correct.

In the proposed setup for the extension of the positive verb sets, a manual annotation step is introduced in order to determine whether each "false positive" verb is in fact a good alternation candidate. False positives are data points whose gold label is negative, but which receive the positive class label from the classifier. Human annotators are shown the observed instances of false positive candidates from the corpus in all relevant patterns of the alternation, and decide how well the sentences in the patterns fit the constructions associated with the alternation. If annotators find that the majority of instances of the patterns do in fact instantiate the relevant alternation-specific constructions, the verb is moved from the negative set to the positive set.

As verbs are moved from the negative set to the positive set, the two classes are not balanced anymore in terms of contained verbs. Recall that the negative set was originally created based on observed verbs that appear in the relevant patterns, but are not included in the known positive set. Each positive verb received a corresponding negative verb with roughly the same overall frequency in the corpus. After moving verbs from the negative set to the positive set, new negative verbs can be selected from the corpus as counterparts for the freshly-added positive ones and for the existing positive verbs that lost their counterparts. Then, the classifier for that alternation can be retrained based on the new positive and negative sets, and new predictions can be made.

The predicted labels from this second round of predictions can again contain false positives, which are again annotated manually and possibly moved to the positive set. This process is repeated until no new additions are made to the positive set.

The SVM classifiers presented in this chapter use observable features of each verb as indicators for whether the given verb should receive the positive or negative label. Each verb's label is derived from a continuous score that is in turn based on the combination of the various feature values for that verb. That score indicates a positive label if it is above zero, and a negative label if it is below zero. The score itself represents the distance of the data point from the separating hyperplane whose position and angle were learned during training. If the hyperplane is a good approximation of the data that is being modeled, then instances that are less strictly linked to one of the classes will be located closer to the hyperplane, and instances whose properties are more strongly associated with a class will be located further away from the hyperplane.

In order to find new alternating verbs in the data, the most promising candidates are those in the set of false positives that received a strong positive-label score from the classifier. These are verbs that were not initially members of the positive set, but whose properties are so similar to known positive instances that the classifier confidently places them in the positive category. False positives with weaker scores are less promising, because their properties do not align as strongly with the positive class.

False negatives are not relevant here, since the setup relies on the assumption that all verbs in the positive set are truly alternating verbs. If a classifier mislabels a positive verb as belonging to the negative class, the verb should not be moved to the negative class based on that classifier label.

This iterative approach to the extension of the set of participating verbs for an alternation has been published in an earlier form as Seyffarth & Kallmeyer (2020). In the context of this chapter, no new implementation will be performed, as the repeated annotation task and the manual labor required for it are beyond the scope of this thesis. The remainder of this section will be dedicated to summarizing and expanding upon the findings from Seyffarth & Kallmeyer.

The experiments reported by Seyffarth & Kallmeyer focus on the instrument subject alternation and the causative-inchoative alternation. The initial classification results are different from those reported in this thesis because a different set of features in each feature category are used and the corpora are less strictly filtered than they are here. The perplexity features in the paper are based on Transformer-XL, while this thesis uses GPT-2. In Seyffarth & Kallmeyer, the perplexity-based features achieve the strongest score for the causative-inchoative alternation, as in this thesis, while the vector-based features achieve the strongest score for the instrument subject alternation.

The annotation step is described by Seyffarth & Kallmeyer (2020: 4047) as follows. After each classification round, the 40 highest-scoring false positives are presented to two annotators. Each of these candidate verbs is annotated on the lemma level, taking into account the observed instances from the relevant subcorpora for the verb.

Choosing to annotate only the 40 highest-scoring candidates per annotation round reduces the manual annotation effort that would come with annotating all false positives. At the same time, annotating 40 verbs instead of just one or a handful of candidates makes it more likely that good new candidates are actually identified in each round.

Annotators choose between four labels judging the extent to which the observed instances of the relevant patterns instantiate the relevant alternation-specific constructions. Labels are assigned to the whole set of sentences presented for each verb and each pattern. The labels for the instrument subject alternation are shown below.

1. 1ARSA pattern (ISA):
   a) Sentences have a strong instrument candidate in the EA1 slot.
   b) Sentences have a good instrument candidate in the EA1 slot.
   c) Sentences sometimes have an instrument candidate in the EA1 slot.
   d) Sentences do not have a good instrument candidate in the EA1 slot.
2. 2ARSA pattern (ISA):
   a) Sentences have a strong instrument candidate in the IA2 slot.
   b) Sentences have a good instrument candidate in the IA2 slot.
   c) Sentences sometimes have an instrument candidate in the IA2 slot.
   d) Sentences do not have a good instrument candidate in the IA2 slot.

The labels for the causative-inchoative alternation candidates are shown below.

1. 1ARSA pattern (CIA):
   a) The majority of subjects are inchoatively impacted by the verb.
   b) Some subjects are inchoatively impacted by the verb.
   c) Subjects are rarely inchoatively impacted by the verb.
   d) Subjects are not inchoatively impacted by the verb.
2. 2ARSA pattern (CIA):
   a) The majority of subjects have a causative impact on their objects.
   b) Some subjects have a causative impact on their objects.
   c) Subjects rarely have a causative impact on their objects.
   d) Subjects do not have a causative impact on their objects.

The results of each annotation round are then used to prepare the next classification and annotation round for that annotator. Verbs that receive an (a) or (b) label from the annotator are moved to the positive set, and new negative verbs are added to the negative set, using the same negative verb collection strategy as in the initial creation of the verb sets, to ensure class balance for the next round.

Additionally, verbs that receive the label (d) for the relevant patterns are discarded from the negative set and replaced with new, presumably-negative candidates for the next classification round. This is done because the strong rejection of the alternation-specific constructions on the part of the annotator makes it clear that these candidates are atypical, but correct examples of the negative class; human annotators clearly recognize them as negative verbs, but their feature values lead the classifier to label them as positive verbs instead.

For the next classification round, the classifier is trained on the updated positive and negative sets and is then used to predict labels for the same set of instances. While testing on training data is not usually good practice, it is done by Seyffarth & Kallmeyer (2020: 4049) because the goal of the classifier at this stage is not to generalize to new instances, but instead to identify which verbs in its training data are most likely to belong to another class.

Seyffarth & Kallmeyer (2020: 4049) stress that the goal for the classifier at this stage is not to achieve perfect accuracy, since the classes are not regarded as completely fixed. Instead, the goal is to identify presumably-negative verbs that behave similarly enough to verbs in the positive class to warrant human annotation.

In subsequent annotation rounds, candidate verbs that were already seen by the annotator are not presented again. Each time an annotator labels a false positive as not belonging to the positive class, it is stored as a known non-alternating verb.

The results of each annotation round directly influence the contents of the next annotation round. Since annotators work independently of each other and are not synchronized, it is possible for verbs to only ever be seen by one annotator. After the series of annotation and classification rounds is finished by each annotator (yielding no new candidates for the positive set after each annotator's final annotation round), a final annotation round is conducted in which each annotator assigns labels to verbs they have not annotated before, but which have received a label from another annotator. This ensures that each verb is eventually labeled by all annotators.

Seyffarth & Kallmeyer present the scores of the classifiers before and after the manual annotation rounds and the resulting updates to the verb sets. They note (Seyffarth & Kallmeyer 2020: 4049) that the reported classifier scores should be viewed as a lower bound, since even at the end, the negative set is not guaranteed to contain exclusively non-alternating verbs. High recall scores are thus more important than high precision scores.

Seyffarth & Kallmeyer 2020: 4051 report that the iterative alternating verb identification approach involving multiple rounds of classification and manual annotation from two annotators yields new alternating verbs for both the instrument subject alternation and the causative-inchoative alternation. The authors treat verbs that receive (a) or (b) labels from both annotators as strong alternation candidates, and verbs that receive these labels from only one annotator as weak alternation candidates. The distinction shows that alternation participation is not always easy to determine objectively, even for trained linguists and with the help of corpus examples.

29 strong candidates and 21 weak candidates for the instrument subject alternation are found among the presumed-negative verbs for that alternation. 117 additional verbs are annotated, but not relabeled to the positive category. For the causative-inchoative alternation, 6 strong candidates and 12 weak candidates are found, and 64 additional verbs are annotated as real negative verbs. Seyffarth & Kallmeyer also present results for an experimental setting in which only sentences with a length up to 10 tokens are taken into account; this leads to the identification of fewer new candidates

for the instrument subject alternation and roughly the same number of new candidates for the causative-inchoative alternation. For the instrument subject alternation, 29.9% of annotated verbs are relabeled to the positive category (16.8% in the setting filtering corpus examples by sentence length). For the causative-inchoative alternation, 22% of annotated verbs are relabeled to the positive category (24.3% in the setting filtering corpus examples by sentence length).

The annotation reveals that some of the false positives for the instrument subject alternation are verbs that participate in the locatum subject alternation (LSA), where an entity whose location is described by the verb can appear either in the subject position or in a *with* phrase. In other words, the constructions involved in the locatum subject alternation license the same syntactic patterns as the constructions involved in the instrument subject alternation. As in the instrument subject alternation, the constructions involve a semantic role (the locatum) that can be expressed in the EA1 slot or in the IA2 slot. The features defined by Seyffarth & Kallmeyer and in this thesis do not capture which role is expressed in these slots, only that slots are likely to express the same role.

Seyffarth & Kallmeyer (2020: 4052) cite the argument by Levin & Rappaport (1988: 1073) that locatum *with* phrases may be interpreted as instrument phrases, but that this role assignment is unavailable if an additional *with* phrase is present that expresses an instrument. Verbs in the locatum subject alternation may receive a stative reading, in which case no instrument slot is provided by the verb; however, these verbs can also have a change-of-state reading so that the interpretation licensed by the instrument subject alternation is available. An example from Seyffarth & Kallmeyer (2020: 4052) is given in (76).

(76) a. *Trees* decorate the garden.
b. He decorates the garden with *trees*.

In (76a), the dominant reading seems to be the stative one: the sentence describes a state in which the garden has the property of containing decorative trees. However, in (76b), a change-of-state reading is easily accessible to readers such that the sentence describes an event in which the garden is being changed to a state in which it contains decorative trees. This is one case where annotators may diverge in their opinions on the extent to which *trees* can be viewed as an instrument of *decorate*.

A related situation arises for psych verbs that can appear in the syntactic patterns associated with the constructions specific to the instrument subject alternation. An example sentence pair from Seyffarth & Kallmeyer (2020: 4051) is given in (77).

(77) a. *Her stories* amused me.
b. She amused me with *her stories*.

In these sentences, the referent of *her stories* may be regarded as an instrument in a scenario where the agent, *she*, actively performs an act of *amusing* someone and uses *her stories* to achieve this goal. On the other hand, experiencer verbs like this can also appear in these patterns in unintentional scenarios, where *her stories* are not a means to an intended goal but simply happen to have an *amusing* effect on someone.

Levin (1993) presents sentences similar to those in (77) as examples of the possessor subject (transitive) alternation. She also notes that

> It is possible to draw a parallel between the instrument subject alternation (sec. 3.3) and this alternation, with the *with* phrase playing the role of the instrument. From this perspective, the cause of the psychological state can be expressed either as an "oblique" subject or as a prepositional phrase. (Levin 1993: 76)

(Presumed) intentionality impacts annotators' decisions about whether or not something is used as an instrument in the event described by a sentence. Since sentences are presented to annotators in isolation in the procedure described by Seyffarth & Kallmeyer (2020: 4051), and ENCOW is generally only available in a sentence-shuffled form, there is no discourse context available to annotators beyond the sentence itself. As a result, cases like these are likely to be judged differently by different annotators because the available information in the corpus leaves some things open to individual interpretation.

Additionally, for some verbs, the annotators only have access to very few instances of a pattern that can be linked to the relevant alternation-specific construction, particularly when instruments were not obligatory for a verb. For example, the verb *infect* is observed with and without instruments; attestations of the 1ARSA pattern express either an agent or an instrument in the subject position. The annotators' proficiency in English sometimes influences them to select (a) or (b) labels in such cases, even if the attestations instantiating the instrument subject construction are actually in the minority. This is once again a manifestation of the problem that the absence of evidence of felicity is not evidence of infelicity. Constructions may be felicitous for a verb without necessarily being attested in the corpus. In future implementations of the iterative alternating verb identification setup involving manual annotation, special care should go into ensuring that annotators follow a clear strategy in cases where the required constructions are easily imaginable, but not attested in the available corpus data.

The annotation also reveals that the polysemy of *with* is a major factor for the incorrect classification of verbs participating in the instrument subject alternation. Seyffarth & Kallmeyer report that most true negative verbs that are labeled by annotators do in fact occur predominantly with *with* phrases that denote something other than instrument use, such as accompaniment (Seyffarth & Kallmeyer 2020: 4052).

Concerning the new alternation candidates for the causative-inchoative alternation, the set of annotated verbs for this alternation is smaller because the classifiers more accurately predict the initial assigned labels for candidate verbs. This leads to annotation rounds in which fewer than 40 false positives are available for annotation.

Seyffarth & Kallmeyer (2020: 4052) identify, among others, three new verbs participating in the causative-inchoative alternation that belong to the VerbNet class *begin-55-1*. Levin (1993: 275) also states that some members of the *Begin* class participate in the causative-inchoative alternation. These verbs were not originally included in the positive set for the experiments because no explicitly causative frame is given for them in VerbNet 3.3. The classifier correctly identifies these verbs as behaving similarly to the known positive verbs, and the labels given by annotators are consistent with the literature.

Furthermore, the annotators also identify verbs that denote physical state changes or location changes, like *spread*, *circulate*, *relocate*, *transfer*, *pass*, *bounce*, *spring* as new candidates for the causative-inchoative alternation. The classifier apparently recognizes that these verbs behave similarly enough to known positive verbs that they may also belong to the alternation, and the positive label is then confirmed by the annotators.

After all annotation rounds for an alternation are finished, the classifiers are evaluated again on the final positive and negative verb sets. The $F_1$ scores of the best feature set improve slightly for the instrument subject alternation and stay at the same level for the causative-inchoative alternation.

Two main conclusions offer themselves based on the work by Seyffarth & Kallmeyer (2020). First, as in this chapter, verbs participating in different alternations are not necessarily recognized with the same accuracy when basing the classifier on the same set of features. Second, the findings from the manual annotation procedure underline that alternation participation or non-participation is not always easy for humans to objectively classify, and for some verbs, annotators may be certain that the relevant constructions are valid even if they do not appear in the corpus.

Concerning the first point, the role-switching behavior of alternations like the causative-inchoative alternation seems to be less difficult to capture than that of alternations like the instrument subject alternation. The instrument subject alternation is centered around the syntactic position of the instrument role, which does not obligatorily need to be explicitly mentioned. Together with the polysemy of the signal word of the 2ARSA pattern, *with*, this means that the observed instances of the 1ARSA or 2ARSA pattern will not necessarily instantiate the relevant construction, even if the verb participates in the alternation. The underlying assumption of the features implemented in Seyffarth & Kallmeyer and in this thesis is that instances of an alternation-specific syntactic pattern predominantly instantiate the relevant construction if the verb participates in the alternation, and predominantly instantiate unrelated constructions if the verb does not alternate. In the scenario outlined above, this assumption does not necessarily hold.

Concerning the second point, a more refined annotation scheme may be helpful in future implementations of the annotation procedure, but disagreements about individual verbs may remain. This is particularly challenging for alternations like the instrument subject alternation, because the instrument role can be filled by entities belonging to a range of semantic types, and the (un)availability of discourse context impacts whether an annotator views entities in a specific slot as instruments or not. Phenomena like the locatum subject alternation and the possessor subject (transitive) alternation license the same patterns as the instrument subject alternation, but annotators may disagree on whether they are subtypes of the instrument subject alternation or not. Finally, the annotation scheme should provide a mechanism for annotator judgments in cases where relevant constructions are not attested in the corpus, but would be valid according to the language intuition of annotators. A mechanism like this is not available to the classifiers, although one could argue that the perplexity-based features approximate this by "imagining" sentences instantiating the relevant constructions and labeling alternation candidates based on how acceptable these generated sentences are in comparison to the original attested sentences.

## 10.8 conclusion

This part of the thesis revolved around the implementation of several approaches to automatically identifying verbs that participate in three English alternations. Chapter 9 introduced the goals of this task and described the selection of alternations to cover in the experiments, the creation of alternation-specific subcorpora, and the collection of positive and negative verb sets for each alternation. The three chosen alternations – the causative-inchoative alternation, the instrument subject alternation, and the transitive reciprocal alternation (intransitive) – were characterized in terms of their constructions, which are realized with different syntactic patterns that were referred to as the 1ARSA pattern type and the 2ARSA pattern type. Chapter 10 presented a frequency-based baseline for the classification, followed by a set of lemma-based features, a set of vector-based features, and a set of perplexity-based features. The performance of each feature set varied across alternations, although the perplexity-based features were generally among the best-scoring ones. Section 10.7 described a procedure by which "false positives" from the classification were annotated manually by experts in an attempt to extend the known positive sets for each alternation with new alternation candidates.

In this thesis, the alternation identification task is set up in a particularly challenging way, as all candidates, positive as well as negative, are chosen from the pool of verbs that appear at least once in each alternation-specific syntactic pattern in a corpus. This is unlike approaches such as the one taken by Baroni & Lenci (2009), where the negative set contains verbs that never appear in one or more of the relevant syntactic patterns.

Verbs may participate in an alternation in one of their senses, yet be unable to alternate in another. For instance, Levin (1993: 28–29) lists *advance* as a member of a set of "other alternating verbs of change of state" that participate in the causative-inchoative alternation, but also explicitly lists the same verb as a member of the non-alternating set "verbs of future having". This illustrates that alternation participation should ideally be predicted per verb sense. However, word sense disambiguation typically relies on document or sentence context (Raganato, Camacho-Collados & Navigli 2017: 100), both of which are unavailable in ENCOW.

A necessary precondition for each feature set implemented here is the availability of dependency-parsed sentences containing the candidate verbs. Each alternation-specific construction can be described in terms of the syntactic pattern that is associated with it, and the available dependency information for each attestation of the relevant pattern for a candidate verb makes it possible to gather feature values based on the various argument slot fillers for that verb. This also means that the approach can easily be extended to additional alternations.

Merlo & Stevenson (2001) implement various linguistically motivated features that are specifically aimed at estimating the likelihood that an individual argument head for a verb functions as an agent or a theme. This is relevant for their classification task because the alternations they investigate exhibit different patterns with respect to assigning these semantic roles to particular syntactic positions. In contrast to this, the classifier features developed in this thesis are all "role-agnostic", as they merely approximate the semantic similarity of observed entities in different syntactic slots, without explicitly estimating each argument's likelihood to express any specific semantic role. This approach was chosen because it can easily be applied to all three alternations under investigation in this thesis: for the causative-inchoative alternation, the roles that are being contrasted are agents and themes, while for the instrument subject alternation, the relevant roles are agents and instruments, and for the together reciprocal alternation (intransitive), the relevant roles are themes and non-themes.

The assumption that a particular observed syntactic pattern is always an instantiation of one particular construction is a simplification. Because constructions with different semantics can be associated with the same syntactic structure, the feature values can be inaccurate in cases where alternating or non-alternating verbs are observed in superficially similar constructions that are incorrectly interpreted as instances of alternation-specific constructions. A similar issue arises with prepositions, which are highly polysemous. This impacts the performance of the classifiers for the instrument subject alternation and the together reciprocal alternation (intransitive).

The feature sets require different additional resources. The lemma-based features are exclusively based on observed lemmas in each relevant slot for each candidate. This approach is particularly sensitive to sparsity issues arising from infrequent verbs, since it is likely in these cases that no lemma overlap between slots is observed, even if the verb alternates. The vector-based features rely on an external source of pre-trained vectors, although vectors could also be learned from the same corpus that provides the training data. These features aim to alleviate the sparsity issues somewhat, but may still be unsuccessful in cases where no vector is available for a specific slot filler. Finally, the perplexity-based features require a pre-trained language model to score input sentences for perplexity. This method is the best strategy against sparsity because it augments sentences in the available corpus data with "hypothetical", generated alternate sentences instantiating the other construction(s) of the relevant alternation. Note that none of the feature sets require manual annotation or any other handcrafted resource, which makes them easy to apply to alternations in other languages.

While some improvements to the classification setup are possible, it does not seem realistic for any classifier to achieve perfect accuracy. The first issue is one of unattested verbs or verb-pattern combinations. Lapata (1999) shows that some verbs that are listed as alternating by Levin (1993) are unattested in at least one relevant syntactic pattern in a corpus, and that good alternation candidates are attested in the corpus in the relevant patterns, but do not appear in Levin (1993). Classifiers like those presented in this thesis cannot predict labels for items for which no data is available. A second obstacle for achieving perfect accuracy is the fact that even human annotators do not easily agree whether a given verb participates in an alternation or not. Merlo & Stevenson (2001: 395) estimate the expert-based upper bound for a similar task to be around 86.5% agreement with the gold standard.

The two biggest disadvantages of the approaches taken in this thesis are related to the source of the training data. Firstly, the text contained in ENCOW is noisy, and the available dependency annotations in the corpus are not entirely reliable. Secondly, some verbs from the alternation-specific candidate sets are attested so infrequently in this corpus that it is difficult to label them correctly based on their observed occurrences. The available training data is reduced further in the interest of only keeping "reorderable" sentences, which is necessary in order to correctly implement the perplexity-based features. In future work, re-implementing the classification with the feature sets presented here on a corpus that is less noisy and contains more attestations of the relevant verbs may raise the classification scores.

Part V

# CONCLUSION

This part concludes the thesis and gives an outlook on future work.

# 11 CONCLUSION

The work presented in this thesis revolves around certain aspects of verbal alternation behavior that pose challenges for a successful computational treatment of such verbs. The argument structure constructions that are involved in the verb alternations discussed in the thesis are often difficult to distinguish from other constructions, due to the fact that different constructions can share their syntactic structure while differing in their meaning. As an effect, sentences instantiating alternation-specific constructions can only be interpreted correctly if information about the instantiating verbs' alternation behavior is available. The thesis explored this issue from two different angles, the first with a focus on explicitly modeling the interactions between meanings contributed by verbs, their arguments, and the constructions they appear in, the second with a focus on determining for individual verbs whether they participate in specific alternations, based on their attestations in a large corpus.

Part III of this thesis presented a set of metagrammars modeling three English verb alternations and one productive argument structure construction, and sketched how additional alternations can be incorporated into the model. The availability of certain constructions and of certain semantic roles for the observed arguments of a verb was controlled by semantic type constraints and a semantic type hierarchy indicating subtype relationships and explicit incompatibility between types. These constraints were handcrafted, based on theoretical literature about the phenomena at hand. For instance, verbs that can participate in the induced action alternation were required to be compatible with the type *agentive_manner_of_motion*, in accordance with the analysis for such verbs provided by Levin & Rappaport Hovav (1994), and instruments that can function as the ACTOR in a sentence were required to be compatible with the type *autonomous*.

Part IV of the thesis presented a series of approaches to the automatic identification of verbs that participate in the causative-inchoative alternation, the instrument subject alternation, or the together reciprocal alternation (intransitive). Each candidate verb was classified based on the extent to which the observed instances of the verb in the sets of syntactic environments associated with the constructions of a specific alternation were in fact deemed likely to be instances of those alternation-specific constructions. The distinction was made across all attestations for a verb in a syntactic pattern, and no sentence-specific construction classification was performed.

More specifically, the classification features approximated the overlap in the selectional preferences of each verb for the different syntactic argument positions in the relevant syntactic patterns. For all alternations under investigation, the alternation-specific constructions associate a different set of syntactic positions with a shared semantic role than the argument structure constructions outside the alternation whose syntactic structure is identical to that of one of the alternants.

However, as the classification results showed, this approach to identifying alternating verbs is not robust against ambiguity and sparsity, and is not generally reliable for distinguishing between alternating and non-alternating verbs. The distinctions made in the metagrammars in Part III of this thesis did not rely exclusively on selectional preferences for argument slots, but modeled the semantic properties of verbs themselves, the arguments with which they were observed, and properties of the constructions themselves. Providing a set of handcrafted semantic constraints made it possible to identify the correct construction(s) for each input sentence individually, instead of attempting to find trends across all attestations of a verb in a certain syntactic environment throughout a corpus.

In addition, the metagrammar model yielded a grammar in which certain sentences received multiple competing analyses. This was the case, for instance, for sentences instantiating either the induced action construction or the caused motion construction. While these constructions can be realized with the exact same syntactic structure, their semantic representations differ from each other: in the induced action, the verb describes an action performed by the entity in the syntactic object, while the verb in the caused-motion construction expresses an action performed by the entity in the syntactic subject position. Sentences like *Queequeg danced Daggoo to the door* are compatible with both of these constructions, so it is desirable that the parser identifies multiple derivations for them. In contrast to this, the classifiers developed in Part IV relied on the assumption that attestations of a certain syntactic pattern always instantiate an alternation-specific construction in cases where the verb participates in the alternation, and always instantiate a construction unrelated to the alternation if the verb does not alternate.

The metagrammar presented in Part III thus makes more fine-grained distinctions than the classifiers presented in Part IV. It specifies semantic constraints for verbs, arguments and constructions that all combine to form all possible meanings for a given input sentence. While the focus of the metagrammar is on distinguishing certain alternation-specific constructions from other argument structure constructions, alternation behavior is not an explicit lexical property of each verb. Instead, alternation availability follows logically from the sets of constructions that can be instantiated by each verb, based on the compatibility of the verb's semantic type and the types of its arguments with the types required by each construction.

This suggests that a promising direction for alternation classification could be an approach in which such (latent) semantic properties are learned separately, instead of attempting to specify alternation participation as a simple binary label. In this scenario, the construction(s) instantiated by a given sentence with a certain syntactic structure could be identified purely based on these learned properties, analogously to the derivations given by the parser based on the metagrammar implemented in Part III. Levin (1993: 4) notes that "what enables a speaker to determine the behavior of a verb is its meaning". Based on this idea, a computational system processing sentences could determine the plausibility of various competing analyses for a given sentence if it has access to a model specifying the meaning contributions from verbs, arguments and constructions and the ways in which they interact.

As shown by the parsing results in Part III, ambiguous sentences with respect to construction instantiation are possible and expected. A probabilistic ranking of all available analyses from most to least likely may improve the metagrammar further and make it reflect human processing in more detail. Such a probabilistic dimension has not been implemented in the metagrammar so far, but is also suggested by Kallmeyer & Osswald (2013: 322) with respect to modeling the distribution of the alternants associated with the dative alternation.

Merlo & Stevenson (2001) predict alternation participation for three classes of verbs based on a set of linguistically motivated features that specifically take semantic properties of the verbs' arguments into account. Applying their approach to the proposed advanced alternation classifier would involve unsupervised learning of the semantic types of verbs and arguments, the requirements imposed on arguments by a verb's lexicon entry, the type requirements imposed by specific constructions, and a type hierarchy denoting all subtype and incompatibility relations between types.

The proposed approach to argument structure construction identification based on learned (or approximated) latent properties of verbs, their arguments and the constructions in which they appear actually constitutes a departure from the alternation classification task as it was pursued in Part IV (and the numerous related work referenced there). Instead of determining whether a given verb can participate in a specific alternation, the availability of each relevant construction could be determined directly based on the observed lexical items and the learned properties. All benefits of a successful alternating verb classifier would also be achieved by such a system: the correct representation of a sentence's semantic content depends more on the specific instantiated construction than on the alternation licensing that construction. Modeling the "meaning facets" proposed as relevant factors for alternation participation by Levin (1993) benefits not only the analysis of alternating verbs specifically, but in fact the processing of all constructions covered by the learned properties, whether they are associated with alternations or not. Ambiguous constructions would not pose an issue as long as the constraints for each one can be applied to determine the validity (or likelihood) of each one in a specific input sentence.

This proposal is in line with the argument by Goldberg (2002: 349) that constructions are a more useful level of analysis than alternations. As shown in Part III of this thesis, the latent semantic properties of verbs and their arguments that guide construction availability are not necessarily directly linked to alternation participation. For instance, the fact that a verb is compatible with the type *agentive_manner_of_motion* is mainly informative of the fact that this verb can appear in the induced action construction. The induced action alternation is one lens through which the availability of that construction can be categorized – but the *reason* for the availability of that construction is the verb's semantic type, or meaning facet in the term of Levin (1993). To give another example, the agent causation constructions and instrument causation constructions modeled in Chapter 7 are available not exclusively to verbs participating in the instrument subject alternation, such as *slice*, but in fact to all verbs that are compatible with a causative environment, including *change_of_state* verbs like *break*. This overlap in available constructions seems to indicate that a perspective focusing on alternations separately from each other misses out on the construction-level generalizations argued for by Goldberg.

In this construction-focused view, as modeled in Part III of this thesis, the main challenge lies in identifying those latent semantic properties that are salient for each construction under investigation, whether that construction is exclusive to one alternation or shared between several alternations. For the induced action construction, for instance, the semantic type *intransitive_action* does not allow for the relevant distinctions, and would incorrectly allow the verb *laugh* to instantiate this specific construction. A more specific type, *agentive_manner_of_motion*, is needed. On the other hand, the inchoative construction only requires instantiating verbs to be compatible with the fairly general type *change_of_state*. Learning a type hierarchy of the required granularity is not trivial. In order to implement such a system, high-quality training data with reliable gold examples would be necessary.

Moreover, the requirements of a construction on the arguments in each position must also be accounted for. In the metagrammars developed in Part III of this thesis, requirements from verbs and requirements from constructions applied simultaneously to allow or disallow certain sentences to instantiate a specific construction. As noted by Goldberg (1995: 14), certain constructions may impose requirements on arguments that are in conflict with those imposed by the verb. This scenario did not appear in the context of the phenomena modeled in this thesis, but would need to be handled successfully in order to create a large-coverage metagrammar involving more (and more diverse) constructions. This specific issue is a promising avenue for future work in this area. Another phenomenon that did not appear in the metagrammar, but could be modeled in a possible future extension, is the effect described by Montemagni (1994) by which certain alternation-specific constructions can be blocked for individual verbs with specific arguments, even when the verbs otherwise participate in the alternation.

Of the alternations discussed in Part IV of this thesis with respect to the automatic identification of participating verbs, the causative-inchoative alternation and the instrument subject alternation were also modeled in the metagrammars in Part III. The together reciprocal alternation (intransitive) has not been modeled. This is partially due to the fact that this alternation is less well-studied than the others covered in the metagrammars, little literature exists about it, and Levin (1993) merely notes that

> Most of the verbs found here are drawn from the *mix* verbs and the *talk* verbs. The *amalgamate* verbs do not show this possibility, even if they do allow an intransitive use. (Levin 1993: 64)

In order to manually model the together reciprocal alternation (intransitive) in the same fashion as the other alternations, a theoretical foundation would be required for the relevant semantic properties distinguishing alternating verbs from verbs that can instantiate only one or neither of the alternants.

Implementing a system that learns the semantic properties that were manually modeled in Part III, while covering larger sets of verbs similarly to the classifiers developed in Part IV, is beyond the scope of this thesis. However, it is a promising direction for future work and may shed more light on the interactions between verbs, arguments and alternations that determine sentence meaning.

# Part VI

# APPENDIX

# Bibliography


Abeillé, Anne. 1988. Parsing French with Tree Adjoining Grammar: some linguistic accounts. In *Coling Budapest 1988 volume 1: International Conference on Computational Linguistics*.

Abeillé, Anne, Kathleen Bishop, Sharon Cote & Yves Schabes. 1990. *A Lexicalized Tree Adjoining Grammar for English*. Tech. rep. MS-CIS-90-24. University of Pennsylvania.

Abeillé, Anne & Robert D. Borsley. 2021. Basic properties and elements. In Stefan Müller, Anne Abeillé, Robert D. Borsley & Jean-Pierre Koenig (eds.), *Head-driven phrase structure grammar: the handbook*, vol. 9 (Empirically Oriented Theoretical Morphology and Syntax). Language Science Press. DOI: `10.5281/zenodo.5543318`.

Abeillé, Anne & Owen Rambow (eds.). 2000. *Tree Adjoining Grammars: formalisms, linguistic analyses and processing* (CSLI Lecture Notes 107). Stanford: CSLI Publications.

Alexiadou, Artemis, Elena Anagnostopoulou & Florian Schäfer. 2006. The properties of anticausatives crosslinguistically. In *Phases of Interpretation*. Mara Frascarelli (ed.). Berlin, New York: De Gruyter Mouton. 187–212. DOI: `doi:10.1515/9783110197723.4.187`.

Alsina, Alex. 1993. *Predicate composition: a theory of syntactic function alternations*. Stanford University. (Doctoral dissertation).

Anderson, Stephen Robert. 1971. On the role of deep structure in semantic interpretation. *Foundations of Language* 7(3). 387–396.

Andreas, Jacob & Dan Klein. 2014. How much do word embeddings encode about syntax? In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (volume 2: short papers)*, 822–827. Baltimore, Maryland: Association for Computational Linguistics. DOI: `10.3115/v1/P14-2133`.

Androutsopoulos, Ion & Robert Dale. 2000. Selectional restrictions in HPSG. In *COLING 2000 volume 1: the 18th International Conference on Computational Linguistics*.

Anttila, Arto, Matthew Adams & Michael Speriosu. 2010. The role of prosody in the English dative alternation. *Language and Cognitive Processes* 25(7-9). 946–981. DOI: `10.1080/01690960903525481`.

Aranovich, Raúl & Jeffrey T. Runner. 2001. Diathesis alternations and rule interaction in the lexicon. In Karine Megerdoomian & Leora Anne Bar-el (eds.), *Proceedings of the Twentieth West Coast Conference on Formal Linguistics*. Somerville, USA: Cascadilla Press.

Arnold, Jennifer E., Anthony Losongco, Thomas Wasow & Ryan Ginstrom. 2000. Heaviness vs. newness: the effects of structural complexity and discourse status on constituent ordering. *Language* 76(1). 28–55. DOI: `10.2307/417392`.

Arps, David & Simon Petitjean. 2018. A parser for LTAG and frame semantics. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA).

Asudeh, Ash, Mary Dalrymple & Ida Toivonen. 2008. Constructions with lexical integrity: templates as the lexicon-syntax interface. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG08 conference*, 68–88. CSLI Publications.

Asudeh, Ash, Mary Dalrymple & Ida Toivonen. 2013. Constructions with Lexical Integrity. *Journal of Language Modelling* 1(1). 1–54. DOI: 10.15398/jlm.v1i1.56.

Asudeh, Ash & Ida Toivonen. 2014. With Lexical Integrity. *Theoretical Linguistics* 40(1-2). 175–186. DOI: 10.1515/tl-2014-0008.

Baker, Mark Cleland. 1985. *Incorporation: a theory of grammatical function changing*. Massachusetts Institute of Technology. (Doctoral dissertation).

Baker, Mark Cleland. 1997. Thematic roles and syntactic structure. In Liliane Haegeman (ed.), *Elements of grammar: handbook in generative syntax*, 73–137. Dordrecht: Springer Netherlands. DOI: 10.1007/978-94-011-5420-8_2.

Balogh, Kata. 2016. Verbal fields in Hungarian simple sentences and infinitival clausal complements. In David Chiang & Alexander Koller (eds.), *Proceedings of the 12th International Workshop on Tree Adjoining Grammars and Related Formalisms* (*TAG+12*), 58–66. Düsseldorf, Germany.

Baroni, Marco, Georgiana Dinu & Germán Kruszewski. 2014. Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* (*volume 1: long papers*), 238–247. Baltimore, Maryland: Association for Computational Linguistics. DOI: 10.3115/v1/P14-1023.

Baroni, Marco & Alessandro Lenci. 2009. One distributional memory, many semantic spaces. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics*, 1–8. Athens, Greece: Association for Computational Linguistics.

Barsalou, Lawrence W. 1992. Frames, concepts, and conceptual fields. In Adrienne Lehrer & Eva Feder Kittay (eds.), *Frames, fields, and contrasts: new essays in semantic and lexical organization*, 21–74. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc.

Beavers, John. 2005. Towards a semantic analysis of argument/oblique alternations in HPSG. *Proceedings of the 12th International Conference on Head-Driven Phrase Structure Grammar*. 28–48. DOI: 10.21248/hpsg.2005.2.

Beavers, John. 2011. An aspectual analysis of ditransitive verbs of caused possession in English. *Journal of Semantics* 28(1). 1–54. DOI: 10.1093/jos/ffq014.

Belyaev, Oleg. 2023. Introduction to lfg. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar* (Empirically Oriented Theoretical Morphology and Syntax 13). Language Science Press. DOI: 10.5281/zenodo.10037797.

Ben Khelil, Chérifa, Denys Duchier, Yannick Parmentier, Chiraz Zribi & Fériel Ben Fraj. 2016. ArabTAG: from a handcrafted to a semi-automatically generated TAG. In David Chiang & Alexander Koller (eds.), *Proceedings of the 12th International Workshop on Tree Adjoining Grammars and Related Formalisms* (*TAG+12*), 18–26. Düsseldorf, Germany.

Bhatt, Rajesh, Owen Rambow & Fei Xia. 2012. Creating a Tree Adjoining Grammar from a multilayer treebank. In Giorgio Satta & Chung-Hye Han (eds.), *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms* (*TAG+11*), 162–170. Paris, France.

Blackburn, Patrick, Johan Bos & Kristina Striegnitz. 2006. *Learn Prolog now!* Vol. 7 (Texts in Computing). College Publications.

Bladier, Tatiana, Andreas van Cranenburgh, Younes Samih & Laura Kallmeyer. 2018. German and French neural supertagging experiments for LTAG parsing. In Vered Shwartz, Jeniya Tabassum, Rob Voigt, Wanxiang Che, Marie-Catherine de Marneffe

& Malvina Nissim (eds.), *Proceedings of ACL 2018, Student Research Workshop*, 59–66. Melbourne, Australia: Association for Computational Linguistics. DOI: 10.18653/v1/P18-3009.

Bojanowski, Piotr, Edouard Grave, Armand Joulin & Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5. 135–146. DOI: 10.1162/tacl_a_00051.

Bowerman, Melissa. 1988. The 'no negative evidence' problem: how do children avoid constructing an overly general grammar? In J. Hawkins (ed.), *Explaining language universals*, chap. 4, 73–101. Oxford: Basil Blackwell.

Bresnan, Joan. 1994. Locative inversion and the architecture of Universal Grammar. *Language* 70(1). 72–131. DOI: 10.2307/416741.

Bresnan, Joan, Anna Cueni, Tatiana Nikitina & Harald R. Baayen. 2007. Predicting the dative alternation. In G. Bouma, I. Kraemer & J. Zwarts (eds.), *Cognitive foundations of interpretation*, 69–94. Amsterdam: KNAW.

Bresnan, Joan & Jonni M. Kanerva. 1989. Locative inversion in Chicheŵa: a case study of factorization in grammar. *Linguistic Inquiry* 20(1). 1–50. (30 September, 2024).

Bresnan, Joan & Tatiana Nikitina. 2010. The gradience of the dative alternation. In Linda Ann Uyechi & Lian-Hee Wee (eds.), *Reality exploration and discovery: pattern interaction in language and life*, vol. 197 (CSLI Lecture Notes), 161–84. CSLI Publications.

Briscoe, Ted & John Carroll. 1997. Automatic extraction of subcategorization from corpora. In *Fifth Conference on Applied Natural Language Processing*, 356–363. Washington, DC, USA: Association for Computational Linguistics. DOI: 10.3115/974557.974609.

Bullinaria, John A. & Joseph P. Levy. 2007. Extracting semantic representations from word co-occurrence statistics: a computational study. *Behavior Research Methods* 39(3). 510–526. DOI: 10.3758/bf03193020.

Burkhardt, Benjamin, Timm Lichte & Laura Kallmeyer. 2017. Depictives in English: an LTAG approach. In *Proceedings of the 13th International Workshop on Tree Adjoining Grammars and Related Formalisms*, 21–30. Umeå, Sweden: Association for Computational Linguistics.

Butt, Miriam Jessica. 1993. *The structure of complex predicates in Urdu*. Stanford University, Department of Linguistics. (Doctoral dissertation).

Candito, Marie-Hélène. 1996. A principle-based hierarchical representation of LTAGs. In *COLING 1996 volume 1: the 16th International Conference on Computational Linguistics*.

Candito, Marie-Hélène. 1998. Building parallel LTAG for French and Italian. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, volume 1*, 211–217. Montreal, Quebec, Canada: Association for Computational Linguistics. DOI: 10.3115/980845.980880.

Candito, Marie-Hélène. 1999. *Organisation modulaire et paramétrable de grammaires électroniques lexicalisées. application au français et à l'italien*. Université Paris 7. (Doctoral dissertation).

Carpenter, Robert L. 1992. *The logic of typed feature structures: with applications to unification grammars, logic programs and constraint resolution* (Cambridge Tracts in Theoretical Computer Science). Cambridge University Press. DOI: 10.1017/CBO9780511530098.

Chang, Chih-Chung & Chih-Jen Lin. 2011. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2(3). DOI: 10.1145/1961189.1961199.

Chomsky, Noam. 1957. *Syntactic structures*. Mouton. DOI: 10.1515/9783110218329.

Chomsky, Noam. 1975. *The logical structure of linguistic theory*. Springer New York, NY.

Chomsky, Noam. 1995. Categories and transformations. In *The minimalist program*, 219–394. Cambridge, MA: MIT Press. DOI: https://doi.org/10.7551/mitpress/9780262527347.003.0004.

Clément, Lionel & Alexandra Kinyon. 2003. Generating parallel multilingual LFG-TAG grammars from a MetaGrammar. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, 184–191. Sapporo, Japan: Association for Computational Linguistics. DOI: 10.3115/1075096.1075120.

Collins, Peter. 1995. The indirect object construction in English: an informational approach. *Linguistics* 33. 35–49.

Corley, Steffan, Martin Corley, Frank Keller, Matthew W. Crocker & Shari Trewin. 2001. Finding syntactic structure in unparsed corpora: the Gsearch corpus query system. *Computers and the Humanities* 35(2). 81–94. DOI: 10.1023/a:1002497503122.

Cotterell, Ryan & Hinrich Schütze. 2015. Morphological word-embeddings. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1287–1292. Denver, Colorado: Association for Computational Linguistics. DOI: 10.3115/v1/N15-1140.

Crabbé, Benoît. 2005a. Grammatical development with XMG. In Philippe Blache, Edward Stabler, Joan Busquets & Richard Moot (eds.), *5th International Conference on Logical Aspects of Computational Linguistics - LACL 2005*, vol. 3492 (Lecture notes in computer science), 84–100. Bordeaux, France: Springer. DOI: 10.1007/11422532_6.

Crabbé, Benoît. 2005b. *Représentation informatique de grammaires d'arbres fortement lexicalisées: le cas de la grammaire d'arbres adjoints*. Université Nancy 2. (Doctoral dissertation).

Crabbé, Benoît & Denys Duchier. 2005. Metagrammar redux. In Henning Christiansen, Peter Rossen Skadhauge & Jørgen Villadsen (eds.), *Constraint solving and language processing*, 32–47. Berlin, Heidelberg: Springer. DOI: 10.1007/11424574_3.

Crabbé, Benoît, Denys Duchier, Claire Gardent, Joseph Le Roux & Yannick Parmentier. 2013. XMG: eXtensible MetaGrammar. *Computational Linguistics* 39(3). 591–629. DOI: 10.1162/COLI_a_00144.

Croft, William Albert. 1986. *Categories and relations in syntax: the clause-level organization of information*. Stanford University. (Doctoral dissertation).

Cruse, D. Alan. 1972. A note on English causatives. *Linguistic Inquiry* 3(4). 520–528.

Culy, Christopher. 1985. The complexity of the vocabulary of Bambara. *Linguistics and Philosophy* 8(3). 345–351. DOI: 10.1007/bf00630918.

Dai, Zihang, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le & Ruslan Salakhutdinov. 2019. Transformer-XL: attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2978–2988. Florence, Italy: Association for Computational Linguistics. DOI: 10.18653/v1/P19-1285.

Davies, Mark. 2010. The Corpus of Contemporary American English as the first reliable monitor corpus of English. *Literary and Linguistic Computing* 25(4). 447–464. DOI: 10.1093/llc/fqq018.

de Marneffe, Marie-Catherine, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre & Christopher D. Manning. 2014. Universal Stanford dependencies: a cross-linguistic typology. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation* (*LREC'14*), 4585–4592. Reykjavik, Iceland: European Language Resources Association (ELRA).

de Marneffe, Marie-Catherine, Bill MacCartney & Christopher D. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation* (*LREC'06*). Genoa, Italy: European Language Resources Association (ELRA).

de Marneffe, Marie-Catherine & Christopher D. Manning. 2008. The Stanford typed dependencies representation. In *Coling 2008: proceedings of the Workshop on Cross-Framework and Cross-Domain Parser Evaluation*, 1–8. Manchester, UK: Coling 2008 Organizing Committee.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee & Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, volume 1* (*long and short papers*), 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics. DOI: `10.18653/v1/N19-1423`.

Dixon, Robert M. W. 1994. *Ergativity* (Cambridge Studies in Linguistics). Cambridge University Press. DOI: `10.1017/CBO9780511611896`.

Duan, Huizhong, Yanen Li, ChengXiang Zhai & Dan Roth. 2012. A discriminative model for query spelling correction with latent structural SVM. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 1511–1521. Jeju Island, Korea: Association for Computational Linguistics.

Dunietz, Jesse, Lori Levin & Jaime Carbonell. 2017. Automatically tagging constructions of causation and their slot-fillers. *Transactions of the Association for Computational Linguistics* 5. 117–133. DOI: `10.1162/tacl_a_00050`.

Emonds, Joseph. 1972. Evidence that indirect object movement is a structure-preserving rule. *Foundations of Language* 8(4). 546–561.

Ethayarajh, Kawin. 2019. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (*EMNLP-IJCNLP*), 55–65. Hong Kong, China: Association for Computational Linguistics. DOI: `10.18653/v1/D19-1006`.

Eyigöz, Elif. 2010. TAG analysis of Turkish long distance dependencies. In Srinivas Bangalore, Robert Frank & Maribel Romero (eds.), *Proceedings of the 10th International Workshop on Tree Adjoining Grammar and Related Frameworks* (*TAG+10*), 151–156. Yale University: Linguistic Department, Yale University.

Faulhaber, Susen. 2011. Idiosyncrasy in verb valency patterns. *Zeitschrift für Anglistik und Amerikanistik* 59(4). 331–346.

Ferraresi, Adriano, Eros Zanchetta, Marco Baroni & Silvia Bernardini. 2008. Introducing and evaluating ukWaC, a very large web-derived corpus of English. In Stefan Evert, Adam Kilgarriff & Serge Sharoff (eds.), *Proceedings of the 4th Web as Corpus Workshop* (*WAC-4*). Marrakech, Morocco.

Fillmore, Charles J. 1965. *Indirect object constructions in English and the ordering of transformations*, vol. 1 (Monographs on Linguistic Analysis). The Hague: Mouton & Co.

Fillmore, Charles J. 1970a. Subjects, speakers, and roles. *Synthese* 21(3/4). 251–274. DOI: 10.1007/BF00484800.

Fillmore, Charles J. 1970b. The grammar of hitting and breaking. In R. A. Jacobs & P. A. Rosenbaum (eds.), *Readings in English transformational grammar*, 120–133. Waltham, MA: Ginn.

Fillmore, Charles J. 1975. An alternative to checklist theories of meaning. In *Proceedings of the First Annual Meeting of the Berkeley Linguistics Society*, 123–131. Linguistic Society of America.

Fillmore, Charles J. 1977. Topics in lexical semantics. In R. W. Cole (ed.), *Current issues in linguistic theory*, 79–138. Indiana University Press.

Fillmore, Charles J. 1982. Frame semantics. In The Linguistic Society of Korea (ed.), *Linguistics in the morning calm: selected papers from SICOL-1981*. Seoul, Korea: Hanshin Publishing Company.

Fillmore, Charles J. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6(2). 222–254.

Fillmore, Charles J. & Collin Baker. 2009. A frames approach to semantic analysis. In Bernd Heine & Heiko Narrog (eds.), *The Oxford handbook of linguistic analysis*, 313–340. Oxford University Press. DOI: 10.1093/oxfordhb/9780199544004.013.0013.

Fillmore, Charles J., Christopher R. Johnson & Miriam R.L. Petruck. 2003. Background to FrameNet. *International Journal of Lexicography* 16(3). 235–250. DOI: 10.1093/ijl/16.3.235.

Findlay, Jamie Y. 2023a. Lexical Functional Grammar as a Construction Grammar. *Journal of Language Modelling* 11(2). 197–266. DOI: 10.15398/jlm.v11i2.338.

Findlay, Jamie Y. 2023b. LFG and Tree Adjoining Grammar. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar* (Empirically Oriented Theoretical Morphology and Syntax 13). Language Science Press. DOI: 10.5281/zenodo.10037797.

Findlay, Jamie Y., Roxanne Taylor & Anna Kibort. 2023. Argument structure and mapping theory. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar* (Empirically Oriented Theoretical Morphology and Syntax 13). Language Science Press. DOI: 10.5281/zenodo.10037797.

Firth, John Rupert. 1957. A synopsis of linguistic theory, 1930-55. In *Studies in linguistic analysis* (*special volume of the Philological Society*), vol. Special Volume of the Philological Society. Oxford: Blackwell.

Flickinger, Daniel Paul. 1987. *Lexical rules in the hierarchical lexicon*. Stanford University. (Doctoral dissertation).

Fraj, Fériel Ben, Chiraz Zribi & Mohamed Ben Ahmed. 2008. ArabTAG: a Tree Adjoining Grammar for Arabic syntactic structures. In *Proceedings of the International Arab Conference on Information Technology*. Sfax, Tunisia.

Francis, W. Nelson & Henry Kučera. 1964, 1971, 1979. *A standard corpus of present-day edited American English, for use with digital computers* (*Brown*). Providence, Rhode Island: Brown University.

Frank, Robert. 2002. *Phrase structure composition and syntactic dependencies*. Cambridge, MA: MIT Press.

Frense, Jutta & Paul Bennett. 1996. Verb alternations and semantic classes in English and German. *Language Sciences* 18(1-2). 305–317. DOI: 10.1016/0388-0001(96)00022-8.

Gamerschlag, Thomas, Doris Gerland, Rainer Osswald & Wiebke Petersen. 2014. General introduction. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: applications in language and philosophy*, vol. 94 (Studies in Linguistics and Philosophy), 3–21. Cham: Springer International Publishing. DOI: 10.1007/978-3-319-01541-5_1.

Gardent, Claire & Laura Kallmeyer. 2003. Semantic construction in F-TAG. In *10th Conference of the European Chapter of the Association for Computational Linguistics*. Budapest, Hungary: Association for Computational Linguistics.

Gardent, Claire & Shashi Narayan. 2015. Multiple adjunction in feature-based Tree-Adjoining Grammar. *Computational Linguistics* 41(1). 41–70. DOI: 10.1162/COLI_a_00217.

Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum & Ivan A. Sag. 1985. *Generalized Phrase Structure Grammar*. Harvard University Press.

Gehrke, Berit. 2008. *Ps in motion: on the semantics and syntax of P elements and motion events*. Universiteit Utrecht. (Doctoral dissertation).

Gerdes, Kim. 2002. DTAG? In Robert Frank (ed.), *Proceedings of the Sixth International Workshop on Tree Adjoining Grammar and Related Frameworks* (*TAG+6*), 242–251. Universitá di Venezia: Association for Computational Linguistics.

Godfrey, John J., Edward C. Holliman & Jane McDaniel. 1992. SWITCHBOARD: telephone speech corpus for research and development. In *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 1 (ICASSP'92), 517–520. San Francisco, California: IEEE Computer Society. DOI: 10.1109/ICASSP.1992.225858.

Goldberg, Adele E. 1992. The inherent semantics of argument structure: the case of the english ditransitive construction. *Cognitive Linguistics* 3(1).

Goldberg, Adele E. 1995. *Constructions: A Construction Grammar approach to argument structure*. University of Chicago Press.

Goldberg, Adele E. 2002. Surface generalizations: an alternative to alternations. *Cognitive Linguistics* 13(4). 327–356. DOI: 10.1515/cogl.2002.022.

Goldberg, Adele E. 2013. Constructionist Approaches. In *The Oxford Handbook of Construction Grammar*, 14–31. Oxford University Press. DOI: 10.1093/oxfordhb/9780195396683.013.0002.

Goldberg, Adele E. & Ray Jackendoff. 2004. The English resultative as a family of constructions. *Language* 80(3). 532–568. DOI: https://doi.org/10.1353/lan.2004.0129.

Goldsmith, John A. 1980. Meaning and mechanism in grammar. In Susumu Kuno (ed.), *Harvard studies in syntax and semantics*, 423–449. Cambridge, MA: Harvard University Press.

Green, Georgia M. 1974. *Semantics and syntactic regularity*. Indiana University Press.

Gries, Stefan Th. 2003. Towards a corpus-based identification of prototypical instances of constructions. *Annual Review of Cognitive Linguistics* 1(1). 1–27. DOI: 10.1075/arcl.1.02gri.

Gropen, Jess, Steven Pinker, Michelle Hollander, Richard Goldberg & Ronald Wilson. 1989. The learnability and acquisition of the dative alternation in English. *Language* 65(2). 203–257. DOI: 10.2307/415332.

Guerssel, Mohamed, Kenneth Hale, Mary Laughren, Beth Levin & Josie White Eagle. 1985. A crosslinguistic study of transitivity alternations. In William H. Eilfort, Paul D. Kroeber & Karen L. Peterson (eds.), *Papers from the Parasession on Causatives and Agentivity at the 21st Regional Meeting*, 48–63. Chicago, IL: Chicago Linguistic Society.

Gundel, Jeanette K. 1988. Universals of topic-comment structure. In Michael Hammond, Edith A. Moravcsik & Jessica Wirth (eds.), *Studies in syntactic typology*, vol. 17 (Typological Studies in Language). John Benjamins Publishing Company. DOI: 10.1075/tsl.17.16gun.

Hale, Kenneth Locke & Samuel Jay Keyser. 1986. Some transitivity alternations in English. In *Lexicon Project working papers 7*. Cambridge, MA: Center for Cognitive Science, MIT.

Hale, Kenneth Locke & Samuel Jay Keyser. 1987. A view from the middle. In *Lexicon Project working papers 10*. Cambridge, MA: Center for Cognitive Science, MIT.

Hall, Barbara Corey. 1965. *Subject and object in English*. Cambridge, MA: Massachusetts Institute of Technology. (Doctoral dissertation).

Halliday, Michael Alexander Kirkwood. 1967. Notes on transitivity and theme in English: part 1. *Journal of Linguistics* 3(1). 37–81.

Halliday, Michael Alexander Kirkwood. 1968. Notes on transitivity and theme in English: part 3. *Journal of Linguistics* 4(2). 179–215.

Hampe, Beate & Doris Schönefeld. 2007. Syntactic leaps or lexical variation? – More on "creative syntax". In Stefan Th. Gries & Anatol Stefanowitsch (eds.), *Corpora in cognitive linguistics: corpus-based approaches to syntax and lexis* (Trends in Linguistics). De Gruyter Mouton. DOI: 10.1515/9783110197709.127.

Han, Chung-hye, Juntae Yoon, Nari Kim & Martha Palmer. 2000. *A feature-based Lexicalized Tree Adjoining Grammar for Korean*. Tech. rep. IRCS-00-04. University of Pennsylvania.

Harley, Heidi. 2002. Possession and the double object construction. *Linguistic Variation Yearbook* 2(1). 31–70. DOI: 10.1075/livy.2.04har.

Harris, Zellig S. 1954. Distributional structure. *Word* 10(2-3). 146–162. DOI: 10.1080/00437956.1954.11659520.

Hashemzadeh, Maryam, Greta Kaufeld, Martha White, Andrea E. Martin & Alona Fyshe. 2020. From language to language-ish: how brain-like is an LSTM's representation of nonsensical language stimuli? In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 645–656. Online: Association for Computational Linguistics. DOI: 10.18653/v1/2020.findings-emnlp.57.

Herbelot, Aurelie. 2013. What is in a text, what isn't, and what this has to do with lexical semantics. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – short papers*, 321–327. Potsdam, Germany: Association for Computational Linguistics.

Herbst, Thomas. 2011. The status of generalizations: valency and argument structure constructions. *Zeitschrift für Anglistik und Amerikanistik* 59(4). 347–368. DOI: 10.1515/zaa-2011-0406.

Hochreiter, Sepp & Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9(8). 1735–1780. DOI: 10.1162/neco.1997.9.8.1735.

Hong, Gumwon. 2005. Relation extraction using support vector machine. In *Second International Joint Conference on Natural Language Processing: full papers*. DOI: 10.1007/11562214_33.

Le-Hong, Phuong, Thi Minh Huyen Nguyen, Phuong Thai Nguyen & Azim Roussanaly. 2010. Automated extraction of Tree Adjoining Grammars from a treebank for Vietnamese. In Srinivas Bangalore, Robert Frank & Maribel Romero (eds.), *Proceedings of the 10th International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+10)*, 165–174. Yale University: Linguistic Department, Yale University.

Hopper, Paul J. & Sandra A. Thompson. 1980. Transitivity in grammar and discourse. *Language* 56(2). 251–299. DOI: 10.1353/lan.1980.0017.

Hwang, Jena D. & Martha Palmer. 2015. Identification of caused motion construction. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics*, 51–60. Denver, Colorado: Association for Computational Linguistics. DOI: 10.18653/v1/S15-1006.

Hwang, Jena D., Annie Zaenen & Martha Palmer. 2014. Criteria for identifying and annotating caused motion constructions in corpus data. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, 1297–1304. Reykjavik, Iceland: European Language Resources Association (ELRA).

Jackendoff, Ray S. 1990. *Semantic structures*. MIT press.

Jäger, Gerhard, Johann-Mattis List & Pavel Sofroniev. 2017. Using support vector machines and state-of-the-art algorithms for phonetic alignment to identify cognates in multi-lingual wordlists. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: volume 1, long papers*, 1205–1216. Valencia, Spain: Association for Computational Linguistics.

Joshi, Aravind K. 1985. Tree adjoining grammars: How much context-sensitivity is required to provide reasonable structural descriptions? In *Natural Language Parsing: Psychological, Computational, and Theoretical Perspectives*. David R. Dowty, Lauri Karttunen & Arnold M. Zwicky (eds.). (Studies in Natural Language Processing). Cambridge University Press. 206–250.

Joshi, Aravind K. 1987. An introduction to tree adjoining grammars. In A. Manaster-Ramer (ed.), *Mathematics of language*, 87–114. Amsterdam: John Benjamins.

Joshi, Aravind K. 2004. Starting with complex primitives pays off: complicate locally, simplify globally. *Cognitive Science* 28(5). 637–668. DOI: 10.1207/s15516709cog2805_2.

Joshi, Aravind K. & Yves Schabes. 1997. Tree-Adjoining Grammars. In Grzegorz Rozenberg & Arto Salomaa (eds.), *Handbook of Formal Languages: Volume 3 Beyond Words*, 69–123. Berlin, Heidelberg: Springer Berlin Heidelberg. DOI: 10.1007/978-3-642-59126-6_2.

Jurafsky, Daniel. 1992. *An on-line computational model of human sentence interpretation: a theory of the representation and use of linguistic knowledge*. EECS Department, University of California, Berkeley. (Doctoral dissertation).

Kaeshammer, Miriam & Vera Demberg. 2012. German and English treebanks and lexica for Tree-Adjoining Grammars. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk & Stelios Piperidis (eds.), *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, 1880–1887. Istanbul, Turkey: European Language Resources Association (ELRA).

Kahane, Sylvain, Marie-Hélène Candito & Yannick de Kercadio. 2000. An alternative description of extractions in TAG. In *Proceedings of the Fifth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+5)*, 115–122. Université Paris 7.

Kallmeyer, Laura. 2005. Tree-local multicomponent Tree-Adjoining Grammars with shared nodes. *Computational Linguistics* 31(2). 187–225. DOI: 10.1162/0891201054223968.

Kallmeyer, Laura & Aravind Joshi. 2003. Factoring predicate argument and scope semantics: underspecified semantics with LTAG. *Research on Language and Computation* 1. 3–58. DOI: 10.1023/a:1024564228892.

Kallmeyer, Laura, Timm Lichte, Wolfgang Maier, Yannick Parmentier, Johannes Dellert & Kilian Evang. 2008. TuLiPA: towards a multi-formalism parsing environment for grammar engineering. In *Coling 2008: Proceedings of the Workshop on Grammar Engineering Across Frameworks*, 1–8. Manchester, England: Coling 2008 Organizing Committee.

Kallmeyer, Laura, Timm Lichte, Rainer Osswald & Simon Petitjean. 2016. Argument linking in LTAG: a constraint-based implementation with XMG. In *Proceedings of the 12th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+12)*, 48–57. Düsseldorf, Germany.

Kallmeyer, Laura & Rainer Osswald. 2013. Syntax-driven semantic frame composition in Lexicalized Tree Adjoining Grammars. *Journal of Language Modelling* 1(2). 267–330. DOI: 10.15398/jlm.v1i2.61.

Kallmeyer, Laura & Maribel Romero. 2008. Scope and situation binding in LTAG using semantic unification. *Research on Language and Computation* 6(1). 3–52. DOI: 10.1007/s11168-008-9046-6.

Kallmeyer, Laura & SinWon Yoon. 2004. Tree-local MCTAG with shared nodes: an analysis of word order variation in German and Korean. In Philippe Blache, Noël Nguyen, Nouredine Chenfour & Abdenbi Rajouani (eds.), *Actes de la 11ème conférence sur le Traitement Automatique des Langues Naturelles. Articles longs*, 229–238. Fès, Maroc: ATALA.

Kann, Katharina, Alex Warstadt, Adina Williams & Samuel R. Bowman. 2019. Verb argument structure alternations in word and sentence embeddings. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, 287–297. DOI: 10.7275/q5js-4y86.

Kaplan, Ronald & Joan Bresnan. 1982. Lexical-Functional Grammar: a formal system for grammatical representation. In Joan Bresnan (ed.), *The mental representation of grammatical relations*, 173–281. Cambridge, MA: MIT Press.

Karp, Daniel, Yves Schabes, Martin Zaidel & Dania Egedi. 1992. A freely available wide coverage morphological analyzer for English. In *COLING 1992 volume 3: the 14th International Conference on Computational Linguistics*.

Kasper, Robert, Bernd Kiefer, Klaus Netter & K. Vijay-Shanker. 1995. Compilation of HPSG to TAG. In *33rd Annual Meeting of the Association for Computational Linguistics*, 92–99. Cambridge, Massachusetts, USA: Association for Computational Linguistics. DOI: 10.3115/981658.981671.

Kiela, Douwe & Stephen Clark. 2014. A systematic study of semantic vector space model parameters. In *Proceedings of the 2nd Workshop on Continuous Vector Space Models and their Compositionality (CVSC)*, 21–30. Gothenburg, Sweden: Association for Computational Linguistics. DOI: 10.3115/v1/W14-1503.

Kim, Meesook. 1999. *A cross-linguistic perspective on the acquisition of locative verbs*. University of Delaware. (Doctoral dissertation).

Kipper, Karin, Hoa Trang Dang & Martha Palmer. 2000. Class-based construction of a verb lexicon. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, 691–696. AAAI Press.

Kipper, Karin, Anna Korhonen, Neville Ryant & Martha Palmer. 2006. Extending VerbNet with novel verb classes. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation* (*LREC'06*). Genoa, Italy: European Language Resources Association (ELRA).

Kipper, Karin, Anna Korhonen, Neville Ryant & Martha Palmer. 2007. A large-scale classification of English verbs. *Language Resources and Evaluation* 42(1). 21–40. DOI: 10.1007/s10579-007-9048-2.

Kodama, Kazuhiro. 2004. The English caused-motion construction revisited: a cognitive perspective. *Papers in Linguistic Science* 10. 41–54.

Krause, Ben, Emmanuel Kahembwe, Iain Murray & Steve Renals. 2018. Dynamic evaluation of neural sequence models. In Jennifer Dy & Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, vol. 80 (Proceedings of Machine Learning Research), 2766–2775. PMLR.

Krause, Ben, Emmanuel Kahembwe, Iain Murray & Steve Renals. 2019. Dynamic evaluation of transformer language models. *CoRR* abs/1904.08378.

Krieger, Hans-Ulrich & John Nerbonne. 1993. Feature-Based Inheritance Networks for Computational Lexicons. In *Inheritance, Defaults and the Lexicon*. Ted Briscoe, Ann Copestake & Valeria de Paiva (eds.). (Studies in Natural Language Processing). Cambridge: Cambridge University Press. 90–136.

Krifka, Manfred. 1999. Manner in dative alternation. In Sonya Bird, Andrew Carnie, Jason Haugen & Peter Norquest (eds.), *Proceedings of the 18th West Coast Conference on Formal Linguistics*, 260–271. Somerville, MA: Cascadilla.

Krifka, Manfred. 2004. Semantic and pragmatic conditions for the dative alternation. *Korean Journal of English Language and Linguistics* 4. 1–31.

Kroch, Anthony S. & Aravind K. Joshi. 1985. *The linguistic relevance of Tree Adjoining Grammar*. Tech. rep. MS-CIS-85-16. University of Pennsylvania.

Kroch, Anthony S. & Aravind K. Joshi. 1987. Analyzing extraposition in a Tree Adjoining Grammar. In *Discontinuous constituency*, 107–149. Leiden, The Netherlands: Brill. DOI: 10.1163/9789004373204_006.

Kroch, Anthony S. & Beatrice Santorini. 1991. The derived constituent structure of the West Germanic verb raising construction. In Robert Freidin (ed.), *Principles and parameters in comparative grammar*, 269–338. Cambridge, MA: MIT Press.

Kübler, Sandra & Heike Zinsmeister. 2014. *Corpus linguistics and linguistically annotated corpora*. Bloomsbury Publishing.

Lakoff, George Philip. 1966. *On the nature of syntactic irregularity*. Indiana University. (Doctoral dissertation).

Langacker, Ronald W. 1987. *Foundations of cognitive grammar. Vol. 1.: Theoretical prerequisites*. 2010. Berlin, Boston: De Gruyter Mouton.

Langacker, Ronald W. 1991. *Foundations of cognitive grammar. Vol. 2: Descriptive application*. 2010. Berlin, Boston: De Gruyter Mouton.

Lapata, Maria. 1999. Acquiring lexical generalizations from corpora: a case study for diathesis alternations. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, 397–404. College Park, Maryland, USA: Association for Computational Linguistics. DOI: 10.3115/1034678.1034740.

Lapata, Maria & Chris Brew. 1999. Using subcategorization to resolve verb class ambiguity. In *1999 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*.

Larson, Richard K. 1988. On the double object construction. *Linguistic Inquiry* 19(3). 335–391.

Lau, Jey Han, Alexander Clark & Shalom Lappin. 2015. Unsupervised prediction of acceptability judgements. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (volume 1: long papers)*, 1618–1628. Beijing, China: Association for Computational Linguistics. DOI: 10.3115/v1/P15-1156.

Lau, Jey Han, Alexander Clark & Shalom Lappin. 2017. Grammaticality, acceptability, and probability: a probabilistic view of linguistic knowledge. *Cognitive Science* 41(5). 1202–1241. DOI: 10.1111/cogs.12414.

Laughren, Mary. 1988. Toward a lexical representation of Warlpiri verbs. In *Thematic relations*, vol. 21 (Syntax and Semantics), 215–242. Leiden, The Netherlands: Brill. DOI: 10.1163/9789004373211_013.

Lê Hồng, Phương, Thị Minh Huyền Nguyễn & Azim Roussanaly. 2008. A metagrammar for Vietnamese LTAG. In Claire Gardent & Anoop Sarkar (eds.), *Proceedings of the Ninth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+9)*, 129–132. Tübingen, Germany: Association for Computational Linguistics.

Leech, Geoffrey. 1993. 100 million words of English. *English Today* 9(1). 9–15. DOI: 10.1017/S0266078400006854.

Levin, Beth. 1985. Lexical semantics in review: an introduction. In Beth Levin (ed.), *Lexical semantics in review*. MIT Press.

Levin, Beth. 1993. *English verb classes and alternations: a preliminary investigation*. Chicago: University of Chicago Press.

Levin, Beth. 2006. *English object alternations: a unified account*. Unpublished manuscript.

Levin, Beth. 2015. Semantics and pragmatics of argument alternations. *Annual Review of Linguistics* 1(1). 63–83. DOI: 10.1146/annurev-linguist-030514-125141.

Levin, Beth & Tova R. Rapoport. 1988. Lexical subordination. In *Proceedings of the 24th Annual Meeting of the Chicago Linguistic Society*, 275–289.

Levin, Beth & Malka Rappaport. 1988. Nonevent -er nominals: a probe into argument structure. *Linguistics* 26(6). 1067–1084. DOI: 10.1515/ling.1988.26.6.1067.

Levin, Beth & Malka Rappaport Hovav. 1994. A preliminary analysis of causative verbs in English. *Lingua* 92. 35–77. DOI: 10.1016/0024-3841(94)90337-9.

Levin, Beth & Malka Rappaport Hovav. 2013. Lexicalized meaning and manner/result complementarity. In Boban Arsenijević, Berit Gehrke & Rafael Marín (eds.), *Studies in the composition and decomposition of event predicates*, vol. 93 (Studies in Linguistics and Philosophy), 49–70. Dordrecht: Springer Netherlands. DOI: 10.1007/978-94-007-5983-1_3.

Li, Hang & Naoki Abe. 1998. Generalizing case frames using a thesaurus and the MDL principle. *Computational Linguistics* 24(2). 217–244.

Lichte, Timm. 2007. An MCTAG with tuples for coherent constructions in German. In Gerald Penn Laura Kallmeyer Paola Monachesi & Giorgio Satta (eds.), *Proceedings of the 12th Conference on Formal Grammar*. Dublin, Ireland. DOI: 10.5281/zenodo.6944587.

Lichte, Timm & Laura Kallmeyer. 2006. Licensing German negative polarity items in LTAG. In Tilman Becker & Laura Kallmeyer (eds.), *Proceedings of the Eighth International Workshop on Tree Adjoining Grammar and Related Formalisms*, 81–90. Sydney, Australia: Association for Computational Linguistics.

Lichte, Timm & Laura Kallmeyer. 2017. Tree-adjoining grammar: a tree-based constructionist grammar framework for natural language understanding. In *Papers from the 2017 AAAI Spring Symposium*.

Lichte, Timm & Simon Petitjean. 2015. Implementing semantic frames as typed feature structures with XMG. *Journal of Language Modelling* 3(1). 185–228. DOI: 10.15398/jlm.v3i1.96.

Löbner, Sebastian. 2015. Functional concepts and frames. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Meaning, frames, and conceptual representation*, 15–42. Berlin, Boston: Düsseldorf University Press. DOI: 10.1515/9783110720129-002.

Long, Chen, Laura Kallmeyer & Rainer Osswald. 2022. A frame-based model of inherent polysemy, copredication and argument coercion. In *Proceedings of the Workshop on Cognitive Aspects of the Lexicon*, 58–67. Taipei, Taiwan: Association for Computational Linguistics.

Majewska, Olga, Ivan Vulić, Diana Mccarthy, Yan Huang, Akira Murakami, Veronika Laippala & Anna Korhonen. 2018. Investigating the cross-lingual translatability of VerbNet-style classification. *Language Resources and Evaluation* 52(3). 771–799. DOI: 10.1007/s10579-017-9403-x.

Marcus, Mitchell P., Beatrice Santorini & Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: the Penn Treebank. *Computational Linguistics* 19(2). 313–330.

Matthews, B.W. 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure* 405(2). 442–451. DOI: 10.1016/0005-2795(75)90109-9.

McCarthy, Diana. 2000. Using semantic preferences to identify verbal participation in role switching alternations. In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*.

McCarthy, Diana. 2001. *Lexical acquisition at the syntax-semantics interface: diathesis alternations, subcategorization frames and selectional preferences*. University of Sussex. (Doctoral dissertation).

McCarthy, Diana & Anna Korhonen. 1998. Detecting verbal participation in diathesis alternations. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, volume 2*, 1493–1495. Montreal, Quebec, Canada: Association for Computational Linguistics. DOI: 10.3115/980691.980817.

Merlo, Paola & Suzanne Stevenson. 2001. Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics* 27(3). 373–408. DOI: 10.1162/089120101317066122.

Meyer, Bertrand. 1997. *Object-oriented software construction, second edition*. Old Tappan, NJ: Prentice Hall.

Michaelis, Laura A. 2015. Sign-Based Construction Grammar. In *The Oxford Handbook of Linguistic Analysis*. Oxford University Press. DOI: 10.1093/oxfordhb/9780199677078.013.0007.

Mikolov, Tomáš. 2012. *Statistical language models based on neural networks*. Brno, CZ: Brno University of Technology, Faculty of Information Technology. (Doctoral dissertation).

Mikolov, Tomas, Stefan Kombrink, Anoop Deoras, Lukas Burget & Jan Honza Cernocky. 2011. RNNLM - recurrent neural network language modeling toolkit. In *IEEE Automatic Speech Recognition and Understanding Workshop*, IEEE Automatic Speech Recognition and Understanding Workshop.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado & Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - volume 2*, 3111–3119. Lake Tahoe, Nevada: Curran Associates Inc.

Miller, George A. 1995. WordNet: a lexical database for English. *Communications of the ACM* 38(11). 39–41. DOI: 10.1145/219717.219748.

Montemagni, Simonetta. 1994. Non alternating argument structures: the causative/inchoative alternation in dictionaries. In Willy Martin, Willem Meijs, Margreet Moerland, Elsemiek ten Pas, Piet van Sterkenburg & Piek Vossen (eds.), *Proceedings of the 6th EURALEX International Congress*, 349–359. Amsterdam, the Netherlands: Euralex.

Müller, Stefan. 2018. *A lexicalist account of argument structure: template-based phrasal LFG approaches and a lexical HPSG alternative* (Conceptual Foundations of Language Science 2). Berlin: Language Science Press. DOI: 10.5281/zenodo.1441351.

Müller, Stefan. 2023. *Grammatical theory: From transformational grammar to constraint-based approaches* (Textbooks in Language Sciences 1). Berlin: Language Science Press. DOI: 10.5281/zenodo.7376662.

Müller, Stefan & Stephen Mark Wechsler. 2014a. Lexical approaches to argument structure. *Theoretical Linguistics* 40(1-2). 1–76. DOI: 10.1515/tl-2014-0001.

Müller, Stefan & Stephen Mark Wechsler. 2014b. Two sides of the same slim Boojum: Further arguments for a lexical approach to argument structure. *Theoretical Linguistics* 40(1–2). 187–224.

Ng, Andrew Y., Michael I. Jordan & Yair Weiss. 2001. On spectral clustering: analysis and an algorithm. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, 849–856. Vancouver, British Columbia, Canada: MIT Press.

Nivre, Joakim. 2003. An efficient algorithm for projective dependency parsing. In *Proceedings of the Eighth International Conference on Parsing Technologies*, 149–160. Nancy, France.

Nivre, Joakim, Johan Hall & Jens Nilsson. 2006. MaltParser: a data-driven parser-generator for dependency parsing. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*. Genoa, Italy: European Language Resources Association (ELRA).

Oehrle, Richard Thomas. 1976. *The grammatical status of the English dative alternation*. Cambridge, MA: MIT. (Doctoral dissertation).

Osswald, Rainer. 2002. *A logic of classification – with applications to linguistic theory*. FernUniversität in Hagen, Department of Computer Science. (Doctoral dissertation).

Osswald, Rainer & Robert D. Van Valin. 2014. FrameNet, frame structure, and the syntax-semantics interface. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: applications in language and philosophy*, 125–156. Cham: Springer International Publishing. DOI: 10.1007/978-3-319-01541-5_6.

Oyón, Alba Luzondo. 2007. Semantic constraints on the caused-motion construction. *Epos: Revista de filología* 0(23). 167–180. DOI: 10.5944/epos.23.2007.10551.

Pedregosa, F. et al. 2011. Scikit-learn: machine learning in Python. *Journal of Machine Learning Research* 12. 2825–2830.

Pennington, Jeffrey, Richard Socher & Christopher Manning. 2014. GloVe: global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Doha, Qatar: Association for Computational Linguistics. DOI: 10.3115/v1/D14-1162.

Peters, Matthew E., Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee & Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, volume 1 (long papers)*, 2227–2237. New Orleans, Louisiana: Association for Computational Linguistics. DOI: 10.18653/v1/N18-1202.

Petersen, Wiebke. 2015. Representation of concepts as frames. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Meaning, frames, and conceptual representation*, 43–68. Berlin, Boston: Düsseldorf University Press. DOI: 10.1515/9783110720129-003.

Petitjean, Simon, Denys Duchier & Yannick Parmentier. 2016. XMG 2: Describing description languages. English. In Maxime Amblard, Philippe de Groote, Sylvain Pogodalla & Christian Retoré (eds.), *Logical Aspects of Computational Linguistics. Celebrating 20 Years of LACL (1996–2016)* (Lecture Notes in Computer Science), 255–272. Springer Berlin Heidelberg.

Pinker, Steven. 2013. *Learnability and cognition: the acquisition of argument structure (1989/2013)*. New edition. MIT Press.

Pollard, Carl & Ivan A Sag. 1994. *Head-driven Phrase Structure Grammar*. University of Chicago Press.

Preiss, Judita, Ted Briscoe & Anna Korhonen. 2007. A system for large-scale acquisition of verbal, nominal and adjectival subcategorization frames from corpora. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 912–919. Prague, Czech Republic: Association for Computational Linguistics.

Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech & Jan Svartvik. 1985. *A comprehensive grammar of the English language* (General Grammar Series). London: Longman.

Radford, Alec, Karthik Narasimhan, Tim Salimans & Ilya Sutskever. 2018. *Improving language understanding by generative pre-training*. Preprint.

Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei & Ilya Sutskever. 2019. *Language models are unsupervised multitask learners*. Preprint.

Raganato, Alessandro, Jose Camacho-Collados & Roberto Navigli. 2017. Word sense disambiguation: A unified evaluation framework and empirical comparison. In *Pro-*

*ceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 99–110. Valencia, Spain: Association for Computational Linguistics.

Rambow, Owen. 2004. *Formal and computational aspects of natural language syntax*. University of Pennsylvania Institute for Research in Cognitive Science. (Doctoral dissertation).

Ransom, Evelyn N. 1979. Definiteness and animacy constraints on passive and double-object constructions in English. *Glossa* 13(2). 215–240.

Rappaport, Malka & Beth Levin. 1988. What to do with theta-roles. In *Lexicon Project working papers 11*, 7–36. Cambridge, MA: Center for Cognitive Science, MIT.

Rappaport Hovav, Malka & Beth Levin. 1998. Building verb meanings. In Miriam Butt & Wilhelm Geuder (eds.), *The projection of arguments: lexical and compositional factors*, 97–134. CSLI Publications.

Rappaport Hovav, Malka & Beth Levin. 2008. The English dative alternation: the case for verb sensitivity. *Journal of Linguistics* 44. 129–167. DOI: 10.1017/S0022226707004975.

Rice, Sally. 1988. Unlikely lexical entries. In *Proceedings of the Fourteenth Annual Meeting of the Berkeley Linguistics Society*, vol. 14, 202–212. Linguistic Society of America. DOI: 10.3765/bls.v14i0.1797.

Romain, Laurence. 2018. *A corpus-based study of the causative alternation in English*. Université Charles de Gaulle - Lille III. (Doctoral dissertation).

Ruppenhofer, Josef, Michael Ellsworth, Miriam R. L. Petruck, Christopher R. Johnson & Jan Scheffczyk. 2006. *FrameNet II: Extended theory and practice*. Working Paper. Berkeley, CA: International Computer Science Institute.

Sag, Ivan A. 1997. English relative clause constructions. *Journal of Linguistics* 33(2). 431–483.

Sag, Ivan A. 2010. English filler-gap constructions. *Language* 86(3). 486–545. DOI: 10.1353/lan.2010.0002.

Sahlgren, Magnus. 2008. The distributional hypothesis. *Rivista di Linguistica* 20. 33–53.

Santorini, Beatrice. 1990. *Part-of-speech tagging guidelines for the Penn Treebank Project*. Tech. rep. MS-CIS-90-47. Department of Computer & Information Science, University of Pennsylvania.

Schäfer, Roland. 2015. Processing and querying large web corpora with the COW14 architecture. In Piotr Bański, Hanno Biber, Evelyn Breiteneder, Marc Kupietz, Harald Lüngen & Andreas Witt (eds.), *Proceedings of Challenges in the Management of Large Corpora 3 (CMLC-3)*. Lancaster: IDS.

Schäfer, Roland & Felix Bildhauer. 2012. Building large corpora from the web using a new efficient tool chain. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, 486–493. Istanbul, Turkey: European Language Resources Association (ELRA).

Schmid, Helmut. 1994. Probabilistic part-of-speech tagging using decision trees. In *Proceedings of International Conference on New Methods in Language Processing*, vol. 12.

Schmid, Helmut. 1999. Improvements in part-of-speech tagging with an application to German. In Susan Armstrong, Kenneth Church, Pierre Isabelle, Sandra Manzi, Evelyne Tzoukermann & David Yarowsky (eds.), *Natural language processing using very large corpora*. Springer.

Schnabel, Tobias, Igor Labutov, David Mimno & Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 298–307. Lisbon, Portugal: Association for Computational Linguistics. DOI: `10.18653/v1/D15-1036`.

Schulte im Walde, Sabine. 2000. Clustering verbs semantically according to their alternation behaviour. In *COLING 2000 volume 2: The 18th International Conference on Computational Linguistics*.

Schulte im Walde, Sabine. 2002. A subcategorisation lexicon for German verbs induced from a lexicalised PCFG. In *Proceedings of the Third International Conference on Language Resources and Evaluation* (*LREC'02*). Las Palmas, Canary Islands - Spain: European Language Resources Association (ELRA).

Schulte im Walde, Sabine. 2006. Experiments on the automatic induction of German semantic verb classes. *Computational Linguistics* 32(2). 159–194. DOI: `10.1162/coli.2006.32.2.159`.

Schuster, Annika, Corina Stroessner, Peter Sutton & Henk Zeevat. 2020. Stochastic frames. In Christine Howes, Stergios Chatzikyriakidis, Adam Ek & Vidya Somashekarappa (eds.), *Proceedings of the Probability and Meaning Conference* (*PaM 2020*), 78–85. Gothenburg: Association for Computational Linguistics.

Seyffarth, Esther. 2018. Verb alternations and their impact on frame induction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, 17–24. New Orleans, Louisiana, USA: Association for Computational Linguistics. DOI: `10.18653/v1/N18-4003`.

Seyffarth, Esther. 2019a. Identifying participation of individual verbs or VerbNet classes in the causative alternation. In *Proceedings of the Society for Computation in Linguistics* (*SCiL*) *2019*, 146–155. DOI: `10.7275/efvz-jy59`.

Seyffarth, Esther. 2019b. Modeling the induced action alternation and the caused-motion construction with Tree Adjoining Grammar (TAG) and semantic frames. In *Proceedings of the IWCS 2019 Workshop on Computing Semantics with Types, Frames and Related Structures*, 19–27. Gothenburg, Sweden: Association for Computational Linguistics. DOI: `10.18653/v1/W19-1003`.

Seyffarth, Esther & Laura Kallmeyer. 2020. Corpus-based identification of verbs participating in verb alternations using classification and manual annotation. In *Proceedings of the 28th International Conference on Computational Linguistics*, 4044–4055. Barcelona, Spain (Online): International Committee on Computational Linguistics. DOI: `10.18653/v1/2020.coling-main.357`.

Silverstein, Michael. 1976. Hierarchy of features and ergativity. In Robert M. W. Dixon (ed.), *Grammatical categories in Australian languages*, 112–171. Canberra: Australian Institute of Aboriginal Studies.

Smyth, Ronald H., Gary D. Prideaux & John T. Hogan. 1979. The effect of context on dative position. *Lingua* 47(1). 27–42. DOI: `10.1016/0024-3841(79)90065-2`.

Snyder, Kieran Margaret. 2003. *The relationship between form and function in ditransitive constructions*. University of Pennsylvania. (Doctoral dissertation).

Speas, Margaret. 1990. *Phrase structure in natural language*, vol. 21 (Studies in Natural Language and Linguistic Theory). Springer.

Stefanowitsch, Anatol. 2008. Negative entrenchment: a usage-based approach to negative evidence. *Cognitive Linguistics* 19(3). 513–531. DOI: `doi:10.1515/COGL.2008.020`.

Stefanowitsch, Anatol & Stefan Th. Gries. 2003. Collostructions: investigating the interaction of words and constructions. *International Journal of Corpus Linguistics* 8(2). 209–243. DOI: doi:10.1075/ijcl.8.2.03ste.

Stefanowitsch, Anatol & Stefan Th. Gries. 2005. Covarying collexemes. *Corpus Linguistics and Linguistic Theory* 1(1). 1–43. DOI: doi:10.1515/cllt.2005.1.1.1.

Storoshenko, Dennis Ryan. 2016. Modelling the ziji blocking effect and constraining bound variable derivations in MC-TAG with delayed locality. In David Chiang & Alexander Koller (eds.), *Proceedings of the 12th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+12)*, 67–76. Düsseldorf, Germany.

Straka, Milan. 2018. UDPipe 2.0 prototype at CoNLL 2018 UD shared task. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 197–207. Brussels, Belgium: Association for Computational Linguistics. DOI: 10.18653/v1/K18-2020.

Straka, Milan, Jan Hajič & Jana Straková. 2016. UDPipe: trainable pipeline for processing CoNLL-U files performing tokenization, morphological analysis, POS tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 4290–4297. Portorož, Slovenia: European Language Resources Association (ELRA).

Sun, Lin & Anna Korhonen. 2009. Improving verb clustering with automatically acquired selectional preferences. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, 638–647. Singapore: Association for Computational Linguistics.

Sun, Lin, Diana McCarthy & Anna Korhonen. 2013. Diathesis alternation approximation for verb clustering. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (volume 2: short papers)*, 736–741. Sofia, Bulgaria: Association for Computational Linguistics.

Talmy, Leonard. 1985. Lexicalization patterns: semantic structure in lexical forms. In Timothy Shopen (ed.), *Language typology and syntactic description*, vol. III: Grammatical categories and the lexicon. Cambridge University Press.

Tesnière, Lucien. 2015. *Elements of structural syntax*. John Benjamins Publishing Company. DOI: 10.1075/z.185.

Thompson, Sandra A. 1995. The iconicity of "dative shift" in English: Considerations from information flow in discourse. In *Syntactic Iconicity and Linguistic Freezes: The Human Dimension*. Marge E. Landsberg (ed.). Berlin, New York: De Gruyter Mouton. 155–176. DOI: 10.1515/9783110882926.155.

Tyler, Andrea & Vyvyan Evans. 2001. Reconsidering prepositional polysemy networks: the case of over. *Language* 77(4). 724–765. DOI: 10.1353/lan.2001.0250.

Van Hooste, Koen. 2018. *Instruments and related concepts at the syntax-semantics interface*. Marge E. Landsberg (ed.). Berlin, New York: Düsseldorf University Press GmbH. 155–176. DOI: 10.1515/9783110882926.155.

Van Roy, Peter Lodewijk. 1990. *Extended DCG notation: a tool for applicative programming in Prolog*. Tech. rep. UCB/CSD-90-583. Berlin, New York: EECS Department, University of California, Berkeley. 155–176. DOI: 10.1515/9783110882926.155.

Van Valin, Robert D., Jr. 2005. *Exploring the syntax-semantics interface*. Marge E. Landsberg (ed.). Berlin, New York: Cambridge University Press. 155–176. DOI: 10.1017/CBO9780511610578.

Van Valin, Robert D., Jr. & Randy J. LaPolla. 1997. *Syntax: structure, meaning, and function*. Marge E. Landsberg (ed.) (Cambridge Textbooks in Linguistics). Berlin, New York: Cambridge University Press. 155–176. DOI: 10.1017/CBO9781139166799.

Vapnik, Vladimir N. 1995. *The nature of statistical learning theory*. Marge E. Landsberg (ed.). Berlin, New York: Springer New York. 155–176. DOI: 10.1007/978-1-4757-2440-0.

Vendler, Zeno. 1972. *Res cogitans: an essay in rational psychology*. Marge E. Landsberg (ed.) (Contemporary philosophy). Berlin, New York: Cornell University Press. 155–176. DOI: 10.1515/9783110882926.155.

Vestergaard, Torben. 1973. A note on objective, instrumental, and affected in English. *Studia Linguistica* 27(1-2). 85–89. DOI: 10.1111/j.1467-9582.1973.tb00597.x.

Vijay-Shanker, K. 1987. *A study of tree adjoining grammars*. Berlin, New York: University of Pennsylvania. (Doctoral dissertation). 155–176. DOI: 10.1515/9783110882926.155.

Vijay-Shanker, K. & Aravind K. Joshi. 1988. Feature structures based Tree Adjoining Grammars. In Marge E. Landsberg (ed.), *Coling Budapest 1988 volume 2: International Conference on Computational Linguistics*, 155–176. Berlin, New York: De Gruyter Mouton. DOI: 10.1515/9783110882926.155.

Vo, Ngoc Phuoc An, Tommaso Caselli & Octavian Popescu. 2014. FBK-TR: applying SVM with multiple linguistic features for cross-level semantic similarity. In Marge E. Landsberg (ed.), *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 284–288. Dublin, Ireland: Association for Computational Linguistics. DOI: 10.3115/v1/S14-2046.

Wang, Alex, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy & Samuel Bowman. 2018. GLUE: a multi-task benchmark and analysis platform for natural language understanding. In Marge E. Landsberg (ed.), *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 353–355. Brussels, Belgium: Association for Computational Linguistics. DOI: 10.18653/v1/W18-5446.

Warstadt, Alex, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang & Samuel R. Bowman. 2020. BLiMP: a benchmark of linguistic minimal pairs for English. In Marge E. Landsberg (ed.), *Proceedings of the Society for Computation in Linguistics 2020*, 409–410. New York, New York: Association for Computational Linguistics. DOI: 10.1515/9783110882926.155.

Warstadt, Alex, Amanpreet Singh & Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics* 7. 625–641. DOI: 10.1162/tacl_a_00290.

Wasow, Thomas. 2002. *Postverbal behavior*. Marge E. Landsberg (ed.) (CSLI Lecture Notes). Berlin, New York: CSLI Publications. 155–176. DOI: 10.1515/9783110882926.155.

Weir, David J. 1988. *Characterizing mildly context-sensitive grammar formalisms*. Berlin, New York: University of Pennsylvania. (Doctoral dissertation). 155–176. DOI: 10.1515/9783110882926.155.

Weischedel, Ralph, Eduard Hovy, Martha Palmer, Mitch Marcus, Robert Belvin, Sameer Pradhan, Lance Ramshaw & Nianwen Xue. 2011. Ontonotes: a large training corpus for enhanced processing. In Marge E. Landsberg (ed.), *Handbook of natural* *language processing and machine translation*, 155–176. Berlin, New York: Springer. DOI: 10.1515/9783110882926.155.

Yang, Zhilin, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov & Quoc V. Le. 2019. XLNet: generalized autoregressive pretraining for language understanding. In Marge E. Landsberg (ed.), *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 5753–5763. Berlin, New York: De Gruyter Mouton. DOI: 10.1515/9783110882926.155.

Zinova, Yulia. 2021. *Russian verbal prefixation: A frame semantic analysis*. Marge E. Landsberg (ed.) (Empirically Oriented Theoretical Morphology and Syntax 7). Berlin: Language Science Press. 155–176. DOI: 10.5281/zenodo.4446717.

Zinova, Yulia & Laura Kallmeyer. 2012. A frame-based semantics of locative alternation in LTAG. In Giorgio Satta & Chung-Hye Han (eds.), *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms* (*TAG+11*), 28–36. Paris, France: De Gruyter Mouton. DOI: 10.1515/9783110882926.155.

# ESTHER SEYFFARTH – CURRICULUM VITAE

## EDUCATION

**Bachelor of Arts**
*2010–2014*
Computational Linguistics and English Studies at Ruhr University Bochum, Germany

**Master of Arts**
*2014–2017*
Information Science and Language Technology at Heinrich Heine University Düsseldorf, Germany
Master Thesis: Vector Semantics and Dependency Parsing: Investigating the Influence of Distributional Information on Parsing Accuracy. Supervised by Laura Kallmeyer.

**Doctor of Philosophy**
*2017–2023*
Linguistics at Heinrich Heine University Düsseldorf, Germany. Supervised by Laura Kallmeyer.

## EMPLOYMENT

**Student assistant**
*2011–2017*
Ruhr University Bochum and Heinrich Heine University Düsseldorf, Germany

**Industry internship**
*2015*
Kauz Semantic Technologies, Düsseldorf, Germany

**Researcher**
*2017–2024*
Member of 2 DFG-funded research projects at Heinrich Heine University Düsseldorf, Germany

# DECLARATION

I affirm in lieu of oath that the dissertation has been prepared by me independently and without unauthorized outside assistance, in compliance with the *Regulations on the Principles for Ensuring Good Scientific Practice at Heinrich Heine University*.

*Düsseldorf, May 8th, 2025*

Esther Seyffarth