Conceptual spaces: Naturalness or cognitive sparseness?

Sebastian Scholz & Gottfried Vosgerau

Suggested Citation:
Scholz, S., & Vosgerau, G. (2025). Conceptual spaces: Naturalness or cognitive sparseness? Synthese, 205(3), Article 129. https://doi.org/10.1007/s11229-025-04941-8

Synthese (2025) 205:129
https://doi.org/10.1007/s11229-025-04941-8

ORIGINAL RESEARCH

Conceptual spaces: Naturalness or cognitive sparseness?

Sebastian Scholz[1] · **Gottfried Vosgerau**[1]

Received: 19 May 2024 / Accepted: 23 January 2025

Abstract
The conceptual spaces framework posits that conceptual content is structured geometrically, and is equipped with cognitive criteria of naturalness (namely, convexity and principles of cognitive economy). Its proponents suggest that cognitive naturalness is naturalness simpliciter, a novel move in a debate that is traditionally focused on how the world, and not the mind, is structured. We argue that "cognitive naturalness" is a misnomer and that the framework describes cognitive sparseness instead. To demonstrate this, we explore the approach's shortcomings across various branches of the naturalness debate, most notably its failure to distinguish natural kinds from fictional kinds. Our diagnosis is that the evolutionary pragmatism employed by its proponents fails to establish a connection to the real world, thus failing to secure the ontological and epistemic objectivity required for a theory of naturalness. We propose an alternative view, ecological empiricism, which posits that natural concepts or properties are those revealed through interaction with the real world.

Keywords Concepts · Naturalness · Objectivity · Ecological empiricism · Semantic Externalism · Metaphysics of Science

1 What is cognitive naturalness?

Traditionally, naturalness is *ontological*. It concerns the world having objective or mind-independent joints to be discovered, or being carved with "conceptual cutlery" (Slater & Borghini, 2013: 25). This rather gruesome metaphor harks back to Plato's

✉ Sebastian Scholz
s.scholz@hhu.de

Gottfried Vosgerau
vosgerau@hhu.de

[1] Institut für Philosophie, Heinrich-Heine-University Düsseldorf, Düsseldorf, Germany

Published online: 12 March 2025

Phaedrus, but the underlying idea has carried over into contemporary analytic philosophy, where the semantic externalists repopularized it in the 1970s.

Cognitive naturalness is quite different. As a first approximation, whether an entity is natural here depends on properties of the mind instead of the world. Again, the idea has been around for some time. It relates to John Locke's *nominal essences*, which "[…] are the abstract ideas that constitute the definitions of species or genera" (Jones, 2023: 1)– but in its modern form emphasizes cognitive psychology over naming and language. Here is David Lewis, a contemporary metaphysician with considerable influence on the naturalness debate, with a disparaging comment:

> Nor should it be said […] that as a contingent psychological fact we turn out to have states whose content involves some properties rather than others, and that is what makes it so that the former properties are more natural. (This would be a psychologistic theory of naturalness). (Lewis, 1983: 377).

Such psychologistic theories have a natural tendency towards *conventionalism*, which is "[…] the view natural kinds don't exist independently of the scientists and others who talk about them." (Bird & Tobin, 2024: 1.1.2). If *strong* conventionalism is true, then there simply are no joints of nature and our classifications are all relative, if not arbitrary. This squares very poorly with Lewis' traditional approach, which will be fleshed out in some detail below. The traditional theories will be called *strongly realistic*, where "realism" shall not be read as an endorsement of abstracta or universals, but of objective or mind-independent joints of nature. The account of cognitive naturalness we discuss, however, combines *weak* realism with *weak* conventionalism, i.e., naturalness is partly about the world and partly about how we think of it. It was penned by Peter Gärdenfors, the first and to date only researcher to develop cognitive naturalness *criteria*. We will therefore take a closer look at his theory of Conceptual Spaces (CS) and examine whether it is relevant for the metaphysical problems that are pertinent to the debate. Next, we will explore what the account says about various distinctions that are important to philosophers, including discussions on mind-independence and objectivity. We conclude that CS' criteria primarily reflect cognitive sparseness, and we propose an alternative view to address its shortcomings. quer

1.1 Conceptual spaces: the framework

Conceptual Spaces are theoretical entities that are postulated within cognitive science to explain and predict cognitive phenomena, or to construct artificial agents. They are mathematically defined, as they are based on axiomatic geometry, metrics, dimensions, and the like. Think of them as models of cognition which epitomize the proposition that conceptual content is structured geometrically. Crucially, these models represent the *similarity* of objects by their spatial proximity in CS– which is also how naturalness enters the scene: According to Quine (1969), similarity and naturalness are variants of the same notion. It is their similarity in relevant respects that makes members of a natural kind members of that kind. However, he did not find a satisfying analysis of either notion, which Gärdenfors ascribes to his reliance on the meth-

odology of logical positivism. By contrast, the geometric analysis seeks to "go below language" (Gärdenfors, 2003: 95) and is presented as a remedy to Quine's problems. As similar objects form clusters in CS, the degree of cohesion of the respective regions can be taken to indicate their degree of naturalness. A necessary condition of naturalness is derived, namely *convexity*, which expresses a very high degree of cohesion. A region is said to be convex, iff for any points A and B from within the region, there is no point C between A and B that does not lie in the region as well.[1]

Let's illustrate the approach with Gärdenfors' (2000) own paradigm example, the *color domain*. "Domain" means that the relevant dimensions come as a package. The color domain consists of three dimensions, namely hue, saturation, and brightness, where hue is a circular dimension accounting for complementary colors. This "color spindle" may be calculated from psychological data such as similarity judgments about visual stimuli by regression analyses such as *multidimensional scaling*. Visual properties such as *being green* are defined as regions in the domain. *Being green* is convex and a good candidate for a natural property. By contrast, the gerrymandered Goodman (1955) property of *being grue* (defined as: *"applies to all things examined before t just in case they are green but to other things just in case they are blue"*, *p. 74*) turns out to be non-convex in a visual space plus time-dimension because the corresponding region disconnects at time *t* (cf. Gärdenfors, 1990). It is evident the author devises his criterion not merely for descriptive cognitive research, but for solving philosophical problems such as Goodman's New Riddle of Induction. Accordingly, the approach has implications for a whole array of sub-disciplines that will be explored in the upcoming sections.

1.2 Cognitive semantics and pragmatism

Putnam (1975) coined the slogan of semantic externalism: *Meanings just ain't in the head!* More specifically, he opposed the idea that the extensions of natural kind terms such as "water" are determined by their intensions, which are cashed out by internal psychological states. Internalism can culminate in the notion that adherents of different scientific paradigms talk past each other, which leads to radical epistemic relativism (cf. Schrenk, 2016: 238–239). By contrast, Putnam adopted a version of strong realism in which natural kind terms are thought to refer directly by being causally linked with essences such as (presumably) H_2O, so that diverging theorists are still able to talk about the same 'stuff'.

Gärdenfors strongly disagrees with Putnam. His approach runs under the label of *cognitive semantics*, which means that the referents of natural kind terms– and all other terms for that matter– are thought to be cognitive entities. Accordingly, the term "water" is thought to refer to the concept WATER, which is a convex region

[1] Douven & Gärdenfors (2019) augment the convexity criterion with principles of cognitive economy and efficiency. Convex regions are seen as the result of an optimization procedure in which limited cognitive creatures try to strike an optimal balance between informativeness, parsimony, and other desiderata in order to survive and adapt to their environments. By appealing to optimal design, the authors move towards postulating not only necessary, but sufficient criteria of naturalness. The details of this approach are included where relevant to our argument. Where they are not, we will refer to convexity as a shorthand.

of some conceptual space.[2] At first glance, this looks like a paradigm instance of semantic internalism, but there are some caveats to point out. First, he acknowledges that language has a social dimension, which Putnam argues for by appealing to the "linguistic division of labor" between experts and laypeople. The process in which conceptual spaces are formed is conceived as a "meeting of minds" (Warglien & Gärdenfors, 2013), so meaning is treated as a communal phenomenon, not as a matter of individual psychological states. Second, internalism about linguistic meaning does not necessitate internalism about mental content: *Does Semantics Need Reality?*

> […] "not directly." Once we accept the conceptual structure of an individual as given, the semantic mapping […] can be described without any recourse to the external world. But a second part of the cognitivistic answer is "indirectly," since the conceptual structure is built up in an individual in interaction with reality. (Gärdenfors, 1999: 14).

Some affinity to realism is clearly present. But does that mean our natural concepts get their meanings from latching onto essences? Absolutely not.

Gärdenfors does not want to subscribe to any view in which the concepts correspond to or represent the world, which would be typical of essentialism or other forms of strong realism. But since both conventionalism and internalism invite relativism, he does want the world to play *some* role in shaping our concepts.[3] This balancing act is undertaken by an appeal to evolution: "Via successful and less successful interactions with the world, the conceptual structure of an individual will adapt to the structure of reality." (Gärdenfors, 2000: 156). It is important to note that he does not talk about reality in and of itself here, but reality as opposing and satisfying our evolutionary needs. This, however, is a pragmatist account akin to Quine's. To put it in a nutshell: Our concepts do not refer, but they *work* because we would have died out if they did not. The 'primordial' CS are thought to provide the basis for more sophisticated, or scientific endeavors to operate on.

2 Shortcomings

We saw how the CS approach invites criticism for leading to problematic relativism, as it makes naturalness dependent on the mind. To address this concern, its proponents adopt a weak version of conventionalism (there is some leeway for conven-

[2] More precisely, it is a "[…] set of convex regions across multiple domains together with a prominence assignment to the domains and information about how the regions in different domains are correlated." (Gärdenfors, 2008: 310).

[3] To get a clear sense of how relativism affects the account, consider again how the CS come about. They are not divinely ordained but rather chosen by scientists who map and model cognition from data on similarity judgments, for instance. This begs the question: Can't we just "pick [our] own conceptual space and in this way make [our] favorite properties come out natural" (Gärdenfors, 1990, p. 91)? If even contrived concepts such as CAT-OR-ELECTRON could be rendered natural, this would lead into full-blown conventionalism– which Gärdenfors seeks to avoid. For an in-depth discussion of contrived (or "gerrymandered") concepts, see Sect. 3.

tions, but we cannot choose the conceptual division of the world at will), while at the same espousing a weak version of realism (we have no epistemic access to the world in itself, we only deal with it successfully). Natural concepts are said to be mind-dependent but still objective. Thereby, proponents of the CS approach suggest cognitive naturalness is naturalness *simpliciter*. By contrast, we will argue that cognitive naturalness is no naturalness at all.

The fourth section will take a closer look at the concepts of mind-independence and objectivity. But first we will show that cognitive naturalness cannot solve the metaphysical problems central to the philosophical debate. Crucially, it cannot account for various distinctions that are important to philosophers. The shortcomings of the approach manifest differently in different branches of the debate, so each will be considered in turn.

According to Judith Crane, "[p]hilosophical treatments of natural kinds are embedded in two distinct projects. […] The kinds studied in the philosophy of science approach are projectible categories that can ground inductive inferences and scientific explanation. The kinds studied in the philosophy of language approach are the referential objects of a special linguistic category—natural kind terms—thought to refer directly." (Crane, 2021: 12177). Her distinction will be taken up and supplemented by a third branch, which is strangely absent from her discussion: The kinds studied in the metaphysics approach are primitives that are postulated to address problems with laws of nature and causality. Note, however, that we introduce the distinction for heuristic reasons, and do not claim it reflects a deep systematic divide. The problems of naturalness can hardly be discussed in isolation.

2.1 Philosophy of language

The philosophy of language approach was shaped by semantic externalism, and is therefore closely tied to the aforementioned concept of direct reference. Proper names are thought to refer directly to the individuals they denote. Analogously, natural kind terms are taken to work as proper names for natural kinds. Thus, their meaning is not a matter of descriptive content, but of their extensions, or the external 'stuff' they rigidly designate across worlds. As with Putnam (1975) and Kripke (1980), reference is secured by causal contact with samples. Accordingly, the theories within this branch deal with the nature of the samples: What must they be like in order for direct reference to work?

Crane (2021) discusses three contemporary approaches, all of which indulge in some form of essentialism.[4] This is because essences are required for rigid designation: The essence of a natural kind is what fundamentally defines it– independently of any particular properties that may be observable in any given world. Natural kind terms would fail to pick out the same stuff across worlds, if they would latch onto the surface properties that are varied. A good deal of her discussion is about the *qua problem* of reference fixing, which has been characterized as follows:

[4] The three forms of essentialism she discusses are represented by Salmon (1981), LaPorte (2003), and Cook (1980), respectively.

> When a speaker grounds a term through perceptual experience, in virtue of what is the term grounded in the cause of that experience qua-one-kind and not qua-another? For that cause will always be an instance of many kinds; e.g. the one object may be an echidna, a monotreme, a mammal, a vertebrate, and so on. (Devitt & Sterelny, 1999: 311).

However, if "any object can belong to at most one natural kind", i.e., if *uniqueness* holds, then "natural kinds […] provide a unique partitioning of the world, carving nature at a unique set of joints" (Crane, 2021: 12191). As a result, the qua problem does not arise because there is only one right kind to ground a given term. Unfortunately, such a view is difficult to maintain because it "severely limits the kinds that count as natural, and it surely discounts the vast majority of scientific kinds" (ibid.).

You would expect the CS account to be silent on all these issues because of its deep-rooted anti-essentialism. Douven goes out of his way to dismantle essentialist intuitions, e.g. by problematizing the micro-essentialist sentiment that water is H_2O (cf. Douven, 2022: 11), and Gärdenfors suggests psychological essentialism is bad metaphysics (cf. 2000: 4.2.2). Moreover, Gärdenfors (1999) criticizes traditional semantics such as direct reference theories for not explaining how cognitive agents are able to "grasp" meanings. For cognitive semantics, this is a matter of establishing neural connections between linguistic expressions ("green") and cognitive structures (GREEN-region), where the latter are at least partly perceptually grounded. Since this view is committed to the idea that "language represents a conceptual structure, but it does not directly represent the world" (Gärdenfors, 2000: 154)– which obviously precludes direct reference theories– cognitive naturalness is simply irrelevant to the metaphysical problems discussed in the philosophy of language branch of the naturalness debate.

Interestingly enough, however, both authors are concerned with uniqueness. For instance: "In standard realist thinking, there could never be more than one conceptual scheme capturing the natural kinds. And it is not clear that the optimal design account guarantees satisfaction of this uniqueness condition." (Douven, 2022: 10). There are a couple of things to unravel here. First, note that this depiction of realist thinking is a straw man. As Crane points out, few philosophers accept uniqueness, and subsequent sections will demonstrate most contemporary realist thinkers accept some form of pluralism about organizing the natural world. Second, it is not obvious to us why the account would need uniqueness in the first place. The reason cannot be to address the qua problem, as this issue arises specifically within causal theories of direct reference, which the CS account does not employ.[5] Most probably, however, Crane and Douven speak of two different kinds of uniqueness: While Crane talks about the idea that for each object there is only one natural kind which it belongs to, Douven addresses the relativistic worry that there might be more than one conceptual scheme. If spaces can be partitioned in multiple optimal ways, e.g. by rotating

[5] We may still ask how the right conceptual region is assigned to a term like "echidna", which would be the cognitive analogue of the qua problem. Presumably, the answer is that linguistic expressions are annexed to cognitive entities in a more dynamic way than to essences. Reference might even change depending on the context, etc.– one does quite certainly not need uniqueness to solve this one.

an optimally partitioned space, then naturalness may look like an arbitrary business, independently of how many kinds are assigned to each object. For this reason, Douven (2022) claims that the optimal partitions are just a very small, *sparse* subset of the possible partitions. However, this does not relate to the discussion about naturalness in the philosophy of language.

2.2 Metaphysics

The discussion about naturalness in metaphysics is mainly inspired by the philosophy of Lewis (1983, 1986), who introduces the notion to solve several interrelated problems. Some of these problems are rather problems within the philosophy of language. We concentrate here on the metaphysical problems (which are, of course, not unrelated to language).

In order to formulate a theory of laws, Lewis introduces his "best system analysis (BSA)". Based on the doctrine of Humean Supervenience (Lewis, 1986), his idea is roughly that laws are the axioms of those deductive systems that best describe the Humean mosaic. The trouble, however, is that there might be an abundance of such deductive systems that turn out equally well but differ in their vocabulary (Loewer, 1996). Moreover, not restricting the possible vocabulary might also lead to the technical difficulty resulting in one single law (Loewer, 2007).

In order to give a metaphysical account of causation, Lewis (2000) presents a counterfactual analysis of causal influence. Combined with his modal realism (Lewis, 1983), this leads to the problem that for every object a proper counterpart or duplicate in other possible words must be found.

Both problems can be solved with "perfectly natural properties" (Lewis, 1983): they determine the vocabulary to be used in the best system and they determine duplicates, which are those objects that share (almost) all perfectly natural properties. As we saw in the quote at the beginning of the paper, Lewis himself thinks that any psychologistic account of natural properties would be a non-starter. However, proponents of the "better best systems account" (BBSA) like Schrenk (2017) and Cohen and Callender (2010) have argued that at least the technical difficulty could be solved with any random set of predicates, including "psychologistic" ones (they advocate to take basic predicates of the special sciences as starting points). Nevertheless, the rest of the job– namely to refer to the right properties– is done by the world. The idea here is that if the predicates are not referring to the right properties, the according system will never be a good one, let alone the best (or a best, according to BBSA).

This kind of thought seems to be what Gärdenfors (2000) has in mind when he says that our concepts are based on our interaction with the world. Douven (2024) explicitly relates CS to Lewis's BSA, suggesting that CS could provide a solution to the problems described above. Nevertheless, we are not convinced that this could work for a definition of natural properties. The reason is, as we have argued, that the world itself must play a pivotal role in the determination of which predicates (or concepts for that matter) can be called natural.[6] Thus, any account that purely concentrates on

[6] As Sider (2011: 27) puts it: "Lewis's constraint on reference is "externalist"; reference is not determined merely by us."

cognitive criteria could maybe establish suitable systems of predicates. However, in order to solve the metaphysical problems, we additionally need the requirement that they refer to the right kind of properties, and this can never be determined by cognitive criteria alone. Merely gesturing at some evolutionary successful interaction is not enough, because this pragmatic move is far too unconstrained.[7]

2.3 Philosophy of science

This branch of the naturalness debate revolves specifically around grounding induction and scientific explanation, as envisaged by Richard Boyd.[8] Boyd's pivotal insight is that the natural world is somewhat messier than the essentialists would have us believe. Traditional essentialism assumes that kinds are defined by necessary and jointly sufficient conditions. By contrast, modern evolutionary biology recognizes that species and taxa are not immutable, uniformly characterized groups. Instead, they are dynamic entities, subject to the forces of genetic drift, natural selection, and gene flow, which can lead to the emergence of new traits and the loss of others without necessarily leading to the formation of a new species. Thus, Boyd conceives of kinds as *homeostatic property clusters*. The idea is roughly that some properties like to 'stick' together qua causal structure to form stable patterns and ground our scientific endeavors. These patterns are objective and make the kinds what they are, while their historicity and vagueness are fully acknowledged.

We are not going to discuss the strengths and weaknesses of Boyd's and his successor's accounts here. What is remarkable, however, is that his thinking, as well as the broader debate he initiated, have been undergoing a pragmatic turn in the sense that the role of scientific practice has been increasingly strengthened, while realism has been losing significance. The early Boyd was basically an essentialist and a direct reference theorist, while the later Boyd relativizes naturalness to *disciplinary matrices* and embraces pluralism in the sense that there are multiple legitimate ways to carve up nature. Although Crane lists realism as one of the key characteristics of the philosophy of science accounts, the realism some of them employ is quite minimal:

> Scientific realism requires that true scientific theories, and the concepts and categories they employ, are responsive to real aspects of the world. […] The more pragmatic accounts require only that a scientific kind be useful for scientific practices, while less pragmatic accounts require that they be connected to causal structures or mechanisms. (Crane, 2021: 12194).

You would think this is where the CS account fits in and can finally shine. Gärdenfors is clearly concerned with grounding induction in naturalness– recall Goodman's Riddle. The account resonates with other theories that make epistemic subjects part of the naturalness equation, e.g. Boyd's disciplinary matrices, Chang's (2022) opera-

[7] In contrast, the requirement of success of a scientific theory might fare much better since scientific methodology is constrained in a way that aims to secure a truth-conducive relation to reality. See also Sect. 4.

[8] "It is a truism that the philosophical theory of *natural* kinds is about how classificatory schemes come to contribute to the epistemic reliability of inductive and explanatory practices." (Boyd, 1999: 146).

tional coherence, or Massimi's (2022) perspectival realism. It might even have an edge over its competitors because it uniquely comes with empirically testable naturalness criteria. Yes, its notion of naturalness is purely cognitive and contains descriptivist elements, but due to its pragmatism is still "responsive to real aspects of the world"– or is it?

We see at least two reasons to doubt that the account can offer meaningful solutions to the problems of the philosophy of science branch of the naturalness debate: Its reliance on evolution and its focus on perception. The first point relates to Goodman's Riddle, which is one of the central problems of the debate. Scholz (2024) argues that Gärdenfors' solution strategy amounts to grounding the convex concepts' instrumental success in their evolutionary success, thereby replacing one relativism with another. Even if this move can be spelled out in a non-circular way, it invites evolutionarily motivated epistemic skepticism because *poor reasoning can be highly adaptive*. In other words, there is an insurmountable gap between what is evolutionarily useful and what the world is really like– just think of the contrast between the color concepts and the physical light spectrum. This example is also the transition to our second point, namely that the CS account of naturalness works well only for cases that are close to perception. Its standard examples involve color, taste, sound, etc., but the further removed from perceptual domains the concepts are, the more difficult it is to find interpretable dimensions and plausible naturalness criteria. Scientific concepts, however, are frequently very abstract, hard to learn, and cannot easily be traced to perceptual states– as also acknowledged by Douven (2024: 8–9). Thus, Boyd and his successors would likely regard the CS account as irrelevant to their theoretical ends. Note that Gärdenfors and Zenker (2013, 2014, 2016) have published a number of papers on how scientific theory structure and change can be represented in the CS framework, but these papers do not mention naturalness criteria at all– which is quite revealing from our point of view.

The deeper issue underlying both points is that the account is not only responsive to real aspects of the world, after all. It is responsive to how we evolved to perceive and think about the world. The world as described by physics, for instance, is quite different from what everyday experience tells us. Science can help us overcome our evolutionary misconceptions and biases precisely because it picks up on what the world is really like.[9]

3 Kinds of kinds

Over two thousand years of naturalness debate have brought about a whole zoo of putative kinds of kinds. Philosophers have typically introduced them as contrasting classes to those considered as natural, serving diverse theoretical purposes. This section is devoted to exploring some of these distinctions in order to review whether and how the CS criteria trace the ontological differences that are at stake. To spoil

[9] Of course, our perception is also able to track some causal relations in the world, but they will always depend on the causal mechanisms of perception and will thus be potentially mind-dependent in ways that make them unsuitable for the needs of the Philosophy of Science project.

the discussion a bit, it turns out that the account can only really trace one distinction, namely the one between sparse and gerrymandered kinds. For most cases, this is not necessarily a problem, but we do think there is one highly problematic case, namely fictional kinds.

Gerrymandered kinds are exemplified by grue emeralds, non-ravens, cats-or-dogs, fungi-that-grow-within-a-100 m-radius-of-my-house, and so forth. They figure prominently in the philosophy of science debate because in contrast to the natural kinds they are thought *not* to support scientific reasoning and explanation. For instance, non-black non-ravens are generally taken *not* to confirm the hypothesis that all ravens are black, although, technically speaking, they do (cf. Hempel, 1945). The trouble is this: When we permit contrived predicates or concepts that denote gerrymandered kinds and properties to be projected in reasoning, we face a scenario where anything can seemingly confirm anything else. Consequently, the standard instantial model of confirmation, according to which a "positive instance of a generalization lends some support to that generalization" (Slater & Borghini, 2013: 5), breaks down. (This is the main reason why Goodman's Riddle is such a big deal). Ruling out the undesirable predicates, however, is not straightforward. Many philosophers think the good predicates are precisely those that carve nature at its joints, while the bad predicates refer to gerrymandered kinds, if anything at all. Within the metaphysics branch of the naturalness debate, the gerrymandered kinds are a thorn in the side of *set nominalism*. Set nominalists such as David Lewis take any set, or arbitrary combination of objects to count as a full-blown property or kind. This "abundance" of properties and kinds leads to the theoretical problems discussed above. Thus, Lewis delineates among all the groupings of things in nature an elite, or *sparse* subset of natural groupings of similar things to serve his theoretical ends.

Gärdenfors seems to achieve a similar result by different means. As was previously mentioned, convexity successfully precludes the Goodman property of *being grue*. Other cases are tricky, e.g. the property of *being non-black* is presumably a convex region of the color domain. (This is not detrimental, of course, as convexity is meant to be necessary, but not sufficient.) Whether the relevant examples are represented by non-convex regions is an empirical concern. However, we consider it plausible that the gerrymandered kinds can be eliminated by *some* criteria of cognitive economy. Why is that? Precisely because they do not depend on the structure of reality (or do so only in a derivative sense when some of their components do), but are artifacts of our own mental "contortions". For instance, we can always conjure up contrived concepts by performing logical operations, e.g. by combining natural ones into random alternations. It goes without saying that such concepts strike a poor balance between informativeness and parsimony, and are difficult to process. Thus, it should come as no surprise that the CS account is well-suited for dealing with these cases. The fact remains, however, that this makes the ontological difference between sparse and gerrymandered concepts a matter of them being represented in different ways. Although feasible, realistically minded philosophers will not be content with this way of drawing the distinction.

Functional kinds are important especially to the essentialists of the philosophy of language branch because their nature is a matter of function, rather than any deep underlying features. One and the same function can be fulfilled by vastly differing

properties, which is to say functions are multiply realizable. Thus, they are puzzling for direct reference theorists because it is unclear what the functional kind terms latch onto. For instance, Schwartz (1978)– who is otherwise an adherent of semantic externalism– denies that artifact terms such as "pen" refer directly. Instead, he speaks of *"*nominal kinds*"* to associate with Locke's nominal essences.

Gärdenfors (2000) acknowledges that "the analysis of functional properties is an enigma for [his] theory" (p. 98) because CS are all about perceptual properties. He briefly explores reducing them to "the actions that the objects 'afford'" (ibid.), and extends this analysis in a later paper (2007). Independently of whether this is feasible, it should be clear that his naturalness criteria are not able to differentiate between functional and other concepts, as natural action concepts are thought to be represented in terms of convex regions as well. On the other hand, one might argue that the account can mimic essential features by placing maximal salience weights on the respective dimensions. However, such maximal salience weights would equally apply to natural kinds (e.g. water having a certain chemical structure) and to functional or "nominal" kinds (e.g. chairs being something to sit on). Since the result of maximal salience weights are not essences but features which are *represented* as being essential, this move cannot provide the ontological distinction in question.

What is so striking about artifacts is that they are *made by us* for specific purposes. This circumstance puts them into dubious ontological standing– from the point of view of traditional realist metaphysics, that is (cf. Khalidi, 2016). Traditional realists tend to stress the importance of mind-independence when it comes to judging how real something is. The next section will discuss this notion in some detail. For the time being, note that many kinds of kinds are mind-dependent in some sense, and that, consequently, naturalness criteria may be evaluated on whether they delineate them from the natural ones. This concerns, for example, categories that are known from social ontology, namely *social kinds* such as money or democracy. Another potentially problematic class are *psychological kinds* such as belief, anger, depression, and so forth. Both might be considered functional alongside artifacts, but there are non-functional cases as well, e.g. *fictional kinds* such as orcs or gods. Moreover, there are a number of intriguing borderline cases, most notably *artificial kinds* such as dog breeds or synthetic chemicals, that appear to be mind-dependent in some sense of that term, but might be taken to have essential features.

How does the CS account deal with these kinds of kinds? Douven (2024) does not make an ontological distinction between natural and social concepts, which is indicative of a larger issue: Although the details are open to empirical inquiry, all of the examples given above can *prima facie* be expected to be represented by cognitively economic concepts, so likely all of them will turn out CS-natural. While problematic from the point of view of traditional realism, this is not necessarily detrimental. One might argue that social, psychological, and artificial kinds are natural phenomena open to scientific inquiry, so no "hard" ontological distinction is required. But this is not an option in the case of fictional kinds: Entities like orcs, wizards, and gods in myths, which do not have a physical existence. They exist primarily in narratives and are products of human imagination and do not causally interact with the physical world. Knowledge about these kinds is derived from narrative contexts and artistic creations. They are not bound by empirical reality but by the internal logic and rules

set by their creators or the traditions of their genres.[10] Qua definition, they have nothing to do with the joints of nature, so any theory of naturalness that treats them as natural must be regarded as highly revisionary. Note that the case of fictional kinds constitutes a great illustration for a point made in the preceding section: Fictional kinds can only turn out natural in an account that is not– or at least not only– "responsive to real aspects of the world".

One may wonder, however, whether Douven's adaptation of the Best Systems account can come to the rescue. After all, part of the motivation behind his move is to make scientific concepts such as those of the social sciences accessible to the CS criteria.[11] One could perhaps argue that fictional kinds– as opposed to social kinds– do not participate in any best system and are therefore not objective, which allows demarcation. However, we will argue that CS-natural kinds are not objective either.

4 Mind-independence and objectivity

Douven (2024) argues that while natural concepts as defined by conceptual spaces "are mind-dependent in *some* sense of that expression, mind-dependence in this sense does not compromise their objectivity" (p. 13). Traditionally, objectivity is closely related to mind-independence in the following sense: Something is said to be objective(ly real) if it is as it is independently of how we conceive of it, i.e. if it is so independently of human minds. When applied to concepts, the idea behind the criterion is that objective concepts "carve up nature at its joints", such that nature itself dictates the boundaries of the concepts. However, both the notion of mind-independence and the notion of objectivity are rather vague and have to be spelled out in more detail to evaluate this claim.

Let us begin by distinguishing two notions of mind-independence which are frequently mixed up– to the detriment of the debate. Firstly, things could be said to be independent of minds if their existence is independent of the existence of (human) minds. This will quite obviously not do for natural concepts: if TIGER and DOG are natural concepts (these are standard examples), so is HOMO SAPIENS. However, the existence of homo sapiens is not independent of the existence of human minds, since humans actually have a human mind (for a similar argument, see Khalidi, 2016).

Much more plausible is to say that, secondly, something is mind-independent if it is as it is independently of how we cognize (perceive, think about, linguistically describe) it. Homo sapiens are the way they are, no matter what theory we have about them. The same is even true for perception and cognitive processes (thinking): Pre-

[10] This is not to say that fictional kinds never relate to *any* aspects of the physical world. However, such relations will not be central to the kinds in the sense that fictional kinds cannot be explored by investigating the world.

[11] "This is proof of principle that we can construct conceptual spaces for scientific concepts– in the above case, concepts pertaining to the social sciences– in basically the same way in which we can construct a space for color concepts […]. […] In short, the question of how we are to demarcate bona fide scientific concepts from gerrymandered ones is not readily provided by Douven and Gärdenfors' optimal design principles. The following proposes an answer to this question that retains the spirit of Gärdenfors' optimality approach by adapting Lewis' Best Systems Account." (Douven, 2024: 8–9).

sumably, humans perceive and think the way they do independently of which theory of perception or thought is currently available.[12]

To see how these ideas connect to the idea of objectivity, we follow the helpful distinction by Sankey (2024) between ontological objectivity, objectivity of truth and epistemic objectivity, which he introduces for his discussion of the objectivity of science. Only the first two notions relate to mind-independence, the third one is concerned with epistemically fruitful methods. Moreover, since we are concerned with concepts here which cannot be true or false, the objectivity of truth is not very important for our discussion and is thus only mentioned alongside with ontological objectivity.

4.1 Ontological objectivity

The notion of an "'objective reality' [expresses] the idea that the world or reality exists in and of itself […] independently of human belief, thought or experience." (Sankey, 2024: 2) Applied to concepts this means that a concept picks out parts of objective reality if what it refers to (its extension) is what it is independently of how we cognize it. TIGER and HUMAN MIND are such objective concepts, since tigers (the kinds or the individuals) and human minds (the kinds or the individuals) are the way they are independently of our cognition about them.[13] MONEY and DEMOCRACY, for example, are not objective because something is only money or a democracy if people have certain beliefs about them.[14]

To say that there is "objective truth" is to say that "[t]ruth is objective in the sense that it does not depend on what we believe." (Sankey, 2024: 4) The connection between ontological objectivity and objectivity of truth is easy if we accept some kind of correspondence theory of truth: If a sentence/thought/proposition is objectively true, it is so because the world is as the sentence/thought/proposition states independently of how (or if) we think about it. This requires that the concepts occurring in the thought pick out parts of objective reality.[15]

Concepts are ontologically objective only in a derivative way, namely if they refer to objective kinds. Thus, the "work" to make them objective is entirely done by the world, the concepts have no role to play in it except referring to the right thing. But how they exactly refer is of no importance. It could be because of some deictic

[12] Note that, unless you have platonic intuitions about concepts, all concepts are at least highly mind-dependent (if not mental themselves). Thus, natural *concepts* could never be mind-independent; rather, concepts can only be said to be mind-independent in a derivative way if they refer to mind-independent kinds.

[13] What exactly the extension of a concept is, is left open on purpose: Some might want to say that natural kind concepts refer to natural kinds which have their members or their essences independently of our cognizing about them. Others may prefer a more nominalist perspective and say that they refer to classes of things which have their properties independently of us.

[14] Of course, single individuals could be wrong about whether they live in a democracy or not. However, nothing can be a democracy without some people having beliefs about free voting, for example. In this sense, unless at least some people have certain beliefs, there will be no democracy.

[15] Things might be different with other theories of truth (for a discussion, see Sankey, 2024). However, since the idea behind natural concepts is that they refer to something in objective reality, the very idea of natural kinds and natural concepts strongly favors a correspondence theory of truth.

introduction of the term– a baptism in the sense of Kripke (1980)– because of some descriptive content or whatsoever. None of these mechanisms, however, is able to determine whether the concept is ontologically objective or not. Take for example color concepts, that are most likely introduced by deixis. As Putnam (1987: 4–8) argues, it does not follow that they track objective properties; rather, they track complex properties that are dependent on how we perceive surfaces.[16]

Therefore, a theory of concepts will never be able to determine whether concepts are ontologically objective (in the derivative way). Only a theory of the objective kinds that the concepts refer to could possibly shed light on the status of concepts regarding ontological objectiveness. Moreover, mind-independence does play a crucial role for ontological objectiveness. Thus, Douven (2024) must have a different kind of objectivity in mind when saying that the theory of conceptual spaces renders concepts objective though mind-dependent in some sense.

4.2 Epistemic objectivity

Epistemic objectivity is related to methods in the case of science, according to Sankey (2024). Translated to the case of concepts, epistemic objectivity is related to the processes of concept formation. Douven (2024) is likely to have this kind of objectivity in mind, since his view is "close to Putnam's (1981) internal realism" (p. 4), which he claims to be based on the idea that our conceptual apparatus is made up in a way that ensures objectivity. His characterization of conceptual spaces, he claims, is such that they are governed by rules that are not under our control. Thus, "to say that natural concepts are mind-dependent in the sense of Douven and Gärdenfors (2019) does not imply that it is dependent on our thinking and theorizing what the natural concepts are." (p. 5) This point is very much in line with our argument above. However, the critical point is whether this form of mind-independence ensures epistemic objectivity.

Sankey (2024) discusses two problems of epistemic objectivity, namely the theory-dependence of observation and the variability of the methods of science. He concludes that science is epistemically objective because the methods are "highly reliable truth-conducive tools of inquiry that the sciences lead so regularly to successful interactions with the world" (p. 9). The important point for our discussion here is that epistemic objectivity can only be obtained if a certain connection to the real world is guaranteed by the epistemic tools we use. A different point to the same effect is made by Khalidi (2016: 243): "In distinguishing real from bogus kinds, what concerns us is whether a kind exists in the sense that it has instances that share causal properties." Although he makes no claim about the methods used to investigate the causal properties, his criterion of their being real instances of the kind plays the role of ensuring a stable connection to the actual world.

Applied to theories of concepts, epistemic objectivity can thus be established if the cognitive processes of concept formation ensure a stable connection to the real

[16] To be sure, this is not to say that color concepts are not related to *any* causal structure of the world; it is only to say that they are also dependent on how our perception (i.e. mind) works and are thus not what they are independently of how we cognize them.

world. In other words, the formation processes must be such that they are guaranteed to track the right kinds in the real world. Douven (2024) might be right that his theory ensures one optimal and sparse conceptual scheme; however, since all his criteria for cognitive naturalness are internal, they cannot establish the required connection to the world. Again, the case of color concepts is a formidable example. In short, criteria of concept formation can never be epistemically objective if they stay internal and do not include some mechanism to establish the connection to the real world.

The appeal to evolutionary success will not suffice, since evolutionary success is much too unconstrained to guarantee the tracking of objective reality. As argued for above, evolution is only sensitive to those aspects of the world that are connected to our needs. Thus, we are not tuned to reality "as it is by itself", but to reality as it is supporting or hindering our survival. Epistemic objectivity, however, asks for more: It requires truth-conduciveness, not survival. What CS is lacking is thus cognitive mechanisms behind concept formation that are truth-conducive, i.e. that are able to track causal relations in the world which are as independent of our cognitive apparatus as possible. In the conclusion, we will shortly hint at Ecological Empiricism as a theory to describe such cognitive mechanisms.

5 The problem of abundance and the problem of objectivity

We have argued that the CS account is not able to provide criteria for naturalness, at least not in a way to satisfy the philosophical needs naturalness was designed to address. However, we do not wish to argue its criteria are thus worthless– quite on the contrary, CS is the first and only theory to spell out cognitively adequate criteria of sparseness.

To evaluate this achievement, we must first take a few steps back. The concept of sparseness has been mentioned above in connection with Lewis' set nominalism, namely as a solution to the problem of abundant properties. However, the problem is not specific to set nominalism, but is discussed more broadly. Hempel encounters a version of the problem, as the received theory of confirmation breaks down if abundant predicates such as "… is a non-raven" are allowed into scientific reasoning. It also concerns Armstrong (1978), who "used the traditional doctrine of universals to draw the distinction between genuine and nongenuine features […]: there simply is no universal of 'being either a cow or an electron'." (Sider, 2011: 4). Both Lewis and Armstrong address the issue by invoking naturalness: "Call a predicate "sparse" when it marks a joint in nature. For Armstrong, a predicate is sparse when there exists a corresponding universal; for Lewis, a predicate is sparse when there exists a corresponding natural property or relation." (Sider, 2011: 85). For Gärdenfors, however, a predicate is sparse when it picks out a convex concept.[17] As we have argued, this notion of sparseness cannot be about the objective joints of nature. Rather, we should understand Gärdenfors as demonstrating that there are other means to get a sparse subset of the abundant predicates than by invoking naturalness. Thus, "cognitive

[17] As mentioned above, convexity is not the only criterion (see also Douven and Gärdenfors, 2019).

naturalness" is a misnomer for the criteria he champions. To avoid terminological confusion, we should stick to "cognitive sparseness" instead.

This result indicates that two problems, which are often treated in one go, must be distinguished: The *problem of abundance*, and the *problem of objectivity*. Technically speaking, any criteria that delineate an elite subset among the abundant groupings of things will count as a solution to the problem of abundance. (This is one of the ideas behind the Better Best System Account; Callender & Cohen, 2010; Schrenk, 2017). What is great about empirically corroborated cognitive criteria, however, is that they reflect how cognitive agents actually think and reason. Consequently, such criteria are helpful in predicting the behavior of cognitive systems, or mimicking their behavior with artificial systems. Crucially, this works independently of 'what holds the world together at its core'.

This is not the case for solutions to the problem of objectivity. This problem shows up in theories of direct reference (Sect. 2.1), Goodman's new riddle of induction (1955), Lewis's (1983) need for defining counterparts of objects in possible worlds (Sect. 2.2), as well as in the philosophy of science branch of the naturalness debate (Sect. 2.3). What all these cases have in common is: They need to make sure that certain predicates, and especially natural kind terms, pick out the– or at least some[18]– *right* objective properties, thereby tracking essences, allowing for induction and establishing real similarity relations. Without addressing the problem of objectivity, essentialists are unable to fix direct reference across worlds and to address the qua problem. Realist metaphysicians are incapable of generating robust accounts of similarity, laws, causation, and so forth. Philosophers of science are powerless to ground or justify induction and scientific explanation.

To see this more clearly, let's zoom in on the latter of these points.[19] Goodman's Riddle, which is pertinent here, cannot be resolved by appealing to the simplicity of the predicates because simplicity is language relative: From within a language that contains "… is grue" as primitive vocabulary, "… is green" appears more contrived. Goodman himself embraced the notion that projectibility is thus relative, and that we just happen to project "… is green" because it got "entrenched" (cf. Cohnitz & Rossberg, 2024: 5.4). Thus, he did not think a solution to the problem of objectivity was available, in the sense that there are no objectively right predicates referring to the objectively right properties, and he left scientific reasoning ungrounded. But he did indeed delineate a sparse subset among the abundant predicates, namely those that are well-entrenched, thereby addressing the abundance problem. This is in stark contrast to the realists of the philosophy of science branch, who postulate objective joints of nature for scientific reasoning to take hold on, thereby establishing a suitable connection to the real world. For these theorists, there are objectively right predicates, namely those that refer to the right kinds or properties. Through this move, which is

[18] We do not want to preclude the possibility of multiple objectively right ways to carve up nature (e.g. Boyd's disciplinary pluralism).

[19] For reasons of length, we cannot go into the details of all these points. Just to gesture at another one: Even if we have sparse grue-predicates, we cannot establish the right similarities (within and across worlds), as the set of grue things contains members that are not similar to each other in any relevant way, which is relevant to questions of laws and causation.

quite common in the naturalness debate, they address both the problem of abundance and the problem of objectivity at once by invoking naturalness.

Scholz's (2024) paper and our discussion show that Douven and Gärdenfors want to follow this very blueprint, only without buying into traditional strong realism. In order for that to work, however, they must somehow establish a suitable connection to the real world– which they try to accomplish by deploying an evolutionary pragmatism. We argued that their approach fails to accomplish the desired goal. However, we do not want to deny the feasibility of pragmatist accounts per se. We see two general ways to ensure the right connection to the world: The first way is to let the world itself take the lead, as it is done e.g. by causal theories of reference à la Kripke (1980), reference magnetism accounts (as adopted by Sider, 2011: 3.2) or by postulating naturalness as an ontological primitive (Lewis, 1983). This way leads to ontological objectivity and objectivity of truth (Sect. 4.1). The second way is to spell out constraints on our epistemic access to the world that ensure that this access is truth-conducive, which leads to epistemic objectivity (Sect. 4.2). The latter way is taken– at least implicitly– by accounts that propose to let science take the lead in identifying the right predicates, as scientific methods are widely taken to be the prime candidate of truth-conducive methods. Thus, the second view is compatible with scientific pragmatism. What we criticize is that the CS brand of evolutionary pragmatism is too unconstrained to generate truth-conducive methods, and solve the problem of objectivity. This is the deeper reason why the CS account does not apply to the metaphysical problems of the naturalness debate. This point, as well as the distinction between the problem of abundance and the problem of objectivity, can be understood as a supplement and a substantive extension of Scholz's criticism.

6 Conclusion

Discussing the core problems that are traditionally discussed in relation to the naturalness of kinds or predicates, we argued that the theory of Conceptual Spaces (CS) is not able to offer suitable solutions. The reason is that CS only addresses one of two problems that have been tried to be solved with naturalness. The first problem is the *problem of the abundance* of predicates and kinds, which shows up in different flavors for different approaches, e.g. for confirmation theories (Hempel, 1945), for a theory of universals (Armstrong, 1978) and for Lewis's (1986) Best Systems Account. This problem is best solved by introducing "sparse" predicates, and it is not very important (for this problem), which predicates exactly are chosen.

In fact, CS proposes criteria that lead to a set of sparse predicates. Moreover, it has the advantage to spell out clear, empirically corroborated criteria and thus to provide a tool to define sparse predicates that are directly related to human cognition. We are convinced that this is an advantage over other proposals, since cognitively adequate criteria are useful for facilitating the explanatory, predictive, and constructive aims of cognitive science. Thus, we argued, CS is a very promising account of sparseness.

However, sparseness provides no solution to the second problem, namely the *problem of objectivity*. The common underlying difficulty here is that we have to ensure that our predicates pick out the right objective properties. This problem has

to be solved by securing the right connection to the world. And exactly this cannot be done by a cognitive theory like CS, since such theories simply do not speak about the world. The rather vague appeal to evolutionary success is far too unconstrained to secure a stable connection to the world.

Naturalness is indeed a solution to both problems at the same time. However, it comes with its own difficulties that have been discussed in the literature. Sparseness, on the other hand, is only a solution to the first problem, not to the second, and therefore cannot replace or substitute naturalness. What is needed for a new conception of naturalness is an approach that provides sparseness *and* the right connection to the world at the same time. In this sense, CS can be viewed as providing one part of such a new conception, but it needs to be supplemented.

There are two ways to ensure the right connection to the world: The first way is to let the world itself take the lead, which leads to ontological objectivity (Sect. 4.1). The second way is to spell out constraints on our epistemic access to the world that ensure that this access is truth-conducive, which leads to epistemic objectivity (Sect. 4.2). Since the second way sits much better with accounts that are somewhat pragmatist in spirit, as CS utterly is, we think that the best supplement for CS will take the second route as well.

As an outlook, we will sketch why Ecological Empiricism (EE; Vosgerau, 2024) is a suitable supplement for CS to yield a full account of naturalness. The basic idea is that concept formation is grounded in the interaction with the world, which provides a stable connection with the real world. In this account, natural concepts are those that refer to "properties that are revealed to us by interacting with the [real] world" (Vosgerau, 2024: 3). In the spirit of (a naturalized version of) affordance theory, EE proposes that our concepts are based on learning stable correlations between our actions and our sense input. Such correlations typically stem from causal features of the world *and* the causal make-up of our cognitive (and corporeal) apparatus. Some correlations, however, are highly dependent on our actions and are thus relatively weak, especially "social affordances" based on conventions (for example that spaghetti are eaten with fork and spoon). Other correlations are much more stable since they depend mainly on the world itself, for example that a rod of a certain length cannot be carried through openings narrower than the rod's length. Vosgerau (2024) argues that scientific measurement is a special and very sophisticated way of interacting with the world that is designed to reveal the most stable correlations. The methods applied in measurement (the drafting of measurement instruments, calibration, modeling the instrument, etc.) aim to diminish the influence of the human agent on the revealed correlations. Scientific measurement thus is truth-conducive as it ensures that the correlations that we find are as "mind-independent" as we can get.

EE is able to supplement CS for the following reasons: Firstly, EE is a cognitive theory as well, giving an account of how concepts emerge on the basis of correlations between actions and sense input (i.e. correlations revealed through interacting with the world). Secondly, it gives the world an essential place in concept formation processes that goes beyond mere sense inputs: the "real" correlations between actions and sense inputs are learned, and these are not just dependent on our perception, but equally dependent on objective reality. Thirdly, it provides criteria for epistemic objectivity, which is the key element that CS is in need of (see 4.2). It does so by iden-

tifying scientific measurement as one human activity that is designed to identify the most stable and thus most mind-independent correlations that we, as human agents, are able to detect. From a very skeptical point of view, this might still not be "reality as it is by itself", but it will be the most objective truth we can get at. EE thus provides an account that is based on pragmatist intuitions (it is based on our interacting with the world) *and* justifies a certain variant of scientific realism, thereby ensuring epistemic objectivity.

To conclude, we have argued that CS alone is not able to provide a new account of naturalness that would be able to address the traditional philosophical problems that naturalness is taken to solve. However, it provides half of a solution by giving a good account of cognitive sparseness. It can be complemented with EE, providing the missing epistemic objectivity, to jointly offer a truly cognitive account of naturalness that does not have to be agnostic to deep metaphysical questions.

Acknowledgements The research was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft DFG) under project number 462340335. We would like to thank our colleagues in Düsseldorf, especially Giacomo Giannini and Markus Schrenk, for most helpful comments and discussions.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

Armstrong, D. M. (1978). *Universals and scientific realism* (Vol. 121 & 2). Cambridge University Press.

Bird, A., & Tobin, E. (2024). Natural kinds. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. Spring Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/spr2024/entries/natural-kinds/

Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 141–85). MIT Press.

Callender, C., & Cohen, J. Special sciences, conspiracy and the better best system account of lawhood. *Erkenntnis*, 73(3), 427–447. https://doi.org/10.1007/s10670-010-9241-3

Chang, H. (2022). *Realism for realistic people: A New Pragmatist Philosophy of Science*. Cambridge University Press. https://doi.org/10.1017/9781108635738

Cohnitz, D., & Rossberg, M. (2024). Nelson Goodman. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. Spring Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/spr2024/entries/goodman/

Cook, M. (1980). If cat is a rigid designator, what does it designate? *Philosophical studies*, 37(1), 61–64. https://doi.org/10.1007/BF00353502

Crane, J. K. (2021). Two approaches to natural kinds. *Synthese*, 199, 5–6. https://doi.org/10.1007/s11229-021-03328-9

Devitt, M. & Sterelny, K. (1999). *Language and reality: An introduction to the Philosophy of Language*. MIT Press.
Douven, I. (2022). Best explanations, natural concepts, and optimal design. In D. Glass & J. Schupbach (Eds.), *Conjunctive Explanations*. Routledge. https://doi.org/10.4324/9781003184324-12
Douven, I. (2024). Naturalness, scientifc concepts, and the substantivity of social metaphysics. *Philosophia*.
Douven, I., & Gärdenfors, P. (2019). What are natural concepts? A design perspective. *Mind and Language*, (3), 313–334. https://doi.org/10.1111/mila.12240
Gärdenfors, P. (1990). Induction, Conceptual Spaces and AI. *Philosophy of Science*, 57(1), 78–95.
Gärdenfors, P. (1999). Does semantics need reality? In A. Riegler, M. Peschl, & A. von Stein (Eds.), *Understanding Representation in the Cognitive Sciences* (pp. 209–217). Springer. https://doi.org/10.1007/978-0-585-29605-0_23
Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.
Gärdenfors, P. (2003). Inductive reasoning: From carnap to cognitive science. In N. C. Singh (Ed.), *Proceedings of the International Conference on Theoretical Neurobiology* (pp. 93–109). National Brain Research Centre.
Gärdenfors, P. (2007). Representing actions and functional properties in conceptual spaces. In T. Ziemke, & J. Zlatev (Eds.), *Body, Language, and Mind* (pp. 167–195).
Gärdenfors, P. (2008). Reasoning in conceptual spaces. In J. E. Adler & L. J. Rips (Eds.), *Reasoning: Studies of Human Inference and Its Foundations* (pp. 302–20). Cambridge University Press. https://doi.org/10.1017/CBO9780511814273.018
Gärdenfors, P., & Zenker, F. (2013). Theory change as dimensional change: Conceptual spaces applied to the dynamics of empirical theories. *Synthese*, 190(6), 1039–58. https://doi.org/10.1007/s11229-011-0060-0
Goodman, N. (1955). *Fact, fiction, and Forecast*. Harvard University Press.
Hempel, C. G. (1945). Studies in the logic of confirmation (I). *Mind*, *54*(213), 1–26.
Jones, J. E. (2023). Locke on Real Essence. In E. N. Zalta, & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*, Summer Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/sum2023/entries/real-essence/
Khalidi, M. A. (2016). Mind-dependent kinds. *Journal of Social Ontology*, 2(21), 223–46. https://doi.org/10.1515/jso-2015-0045
Kripke, S. A. (1980). *Naming and necessity*. Blackwell.
LaPorte, J. (2003). *Natural kinds and conceptual change*. Cambridge University Press.
Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, *61*(4), 343–377. https://doi.org/10.1080/00048408312341131
Lewis, D. (1986). *Philosophical papers, volume II*. Oxford University Press.
Lewis, D. (2000). Causation as influence. *Journal of Philosophy*, *97*(4), 182–197.
Loewer, B. (1996). Humean supervenience. *Philosophical Topics*, *24*(1), 101–127. https://doi.org/10.5840/philtopics199624112
Loewer, B. (2007). Laws and natural properties. *Philosophical Topics*, *35*(1), 313–328. https://doi.org/10.5840/philtopics2007351/214
Massimi, M. (2022). *Perspectival realism. Oxford Studies in Philosophy of Science*. Oxford University Press.
Putnam, H. (1975). The meaning of meaning. *Minnesota Studies in the Philosophy of Science*, *7*, 131–193.
Putnam, H. (1981). *Reason, Truth, and History.* Cambridge University Press.
Putnam, H. (1987). *The many faces of realism. The Paul Carus Lectures 16. LaSalle*. Open Court.
Quine, W. O. (1969). *Natural Kinds. Ontological Relativity and Other Essays*. Columbia University Press. 114–138.
Salmon, N. U. (1981). *Reference and essence*. Princeton University Press.
Sankey, H. (2024). The objectivity of science. *Journal of Philosophical Investigations*, *17*, 45. https://doi.org/10.22034/jpiut.2023.17473
Scholz, S. (2024). Conceptual spaces: A solution to Goodman's new riddle of induction? *Philosophia*.
Schrenk, M. (2016). *Metaphysics of Science: A systematic and historical introduction*. Routledge. https://doi.org/10.4324/9781315639116
Schrenk, M. (2017). The emergence of better best system laws. *Journal for General Philosophy of Science*, 48(3), 469–83. https://doi.org/10.1007/s10838-017-9374-z
Schwartz, S. P. (1978). Putnam on artifacts. *Philosophical Review*, *87*(4), 566–574. https://doi.org/10.2307/2184460
Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.

Slater, M. H., & Borghini, A. (2013). Introduction: Lessons From the Scientific Butchery. In J. K. Campbell, M. O'Rourke, & M. H. Slater (Eds.), *Carving Nature at Its Joints: Natural Kinds in Metaphysics and Science*. MIT Press.

Vosgerau, G. (2024). Ecological empiricism. *Philosophia*.

Warglien, M., & Gärdenfors, P. (2013). Semantics, conceptual spaces, and the meeting of minds. *Synthese*, *190*(12), 2165–2193.

Zenker, F., & Gärdenfors, P. (2014). Modeling diachronic changes in structuralism and in conceptual spaces. *Erkenntnis*, 79(81), 1547–61. https://doi.org/10.1007/s10670-013-9582-9

Zenker, F., & Gärdenfors, P. (2016). Continuity of theory structure: A conceptual spaces approach. *International Studies in the Philosophy of Science*, 30(41), 343–60. https://doi.org/10.1080/02698595.2017.1331983