



Unsupervised Machine Learning For Human Neural Electrophysiology Signals

Inaugural dissertation

For the attainment of the title of doctor
in the Faculty of Mathematics and Natural Sciences
at the Heinrich Heine University Düsseldorf

presented by

Abhinav Sharma

from India

Agra, Uttar Pradesh, January 2023

from the institute for Clinical Neuroscience and Medical Psychology
at the Heinrich Heine University Düsseldorf

Published by permission of the
Faculty of Mathematics and Natural Sciences at
Heinrich Heine University Düsseldorf

Supervisor: Prof. Dr. Esther Florin
Co-supervisor: Prof. Dr. Stefan Conrad

Date of the oral examination: 25 May 2023

“Tout le malheur des hommes vient de ne savoir pas demeurer en repos, dans une chambre“- Blaise Pascal

“All of humanity's problems stem from man's inability to sit quietly in a room alone” - Blaise Pascal

Acknowledgments

Before anything else, I want to express my gratitude to Prof. Dr. Esther Florin for giving me the chance to do research in such a fertile intellectual setting. Invaluable to my growth as a scientist, her encouragement and guidance spanned my whole PhD career. Without the help of many individuals at the Institute of Clinical Neuroscience and Medical Psychology at Uniklinik Düsseldorf, the doctorate program's research would not have been feasible. I especially appreciate the cooperation of Professor Dr. Alfons Schnitzler and Dr. Jan Hirschmann.

Both in my professional and personal life, I have been fortunate to count on an amazing group of friends and co-workers. For their outstanding professional and personal assistance, I am eternally indebted to Matthias Sure, Levent Kandemir, and Dr. Bahne Bahners. In addition, I'd want to express my gratitude to Marius Krösche and Fayed Rassoulou, whose patient and insightful philosophical and scientific talks have been crucial in shaping the direction of my study. Lindsay Bahureksa and Tanvi Nagpal, two of my closest friends who live on opposite sides of the world, were there for me when the COVID epidemic hit. Their presence was a tremendous asset and their contribution was vital in my professional success.

My parents, Drs. Rupali Gupta and Anurag Sharma, and my brother, Additya Sharma, have been my greatest sources of inspiration and encouragement throughout my research and nothing would have been possible without their support.

CONTENTS

ACKNOWLEDGMENTS	3
GLOSSARY	6
ABSTRACT	7
1 INTRODUCTION TO THE THESIS	8
1.1 ELECTROPHYSIOLOGICAL NEURAL SIGNALS	9
1.2 HUMAN NEURAL SIGNALS AND CONNECTIVITY	11
1.3 MACHINE LEARNING FOR HUMAN NEURAL SIGNALS	12
2 STUDY 1 EVALUATING A DATA DRIVEN PIPELINE FOR EXTRACTING EEG RESTING STATE NETWORKS IN HUMANS	17
CONTRIBUTIONS	17
ABSTRACT	17
2.1 INTRODUCTION	18
2.2 METHODS	19
2.2.1 <i>Structural MRI acquisition and pre-processing</i>	19
2.2.2 <i>EEG data acquisition and pre-processing</i>	19
2.2.3 <i>Extracting EEG resting state networks using ICA</i>	20
2.2.4 <i>Extracting EEG resting state networks using megPAC</i>	21
2.2.5 <i>fMRI data acquisition and extraction of resting state networks</i>	21
2.2.6 <i>Comparison of the EEG with fMRI resting state networks</i>	22
2.2.7 <i>Statistics for comparison of networks</i>	22
2.3 RESULTS	22
2.3.1 <i>Bootstrapped EEG-fMRI comparison</i>	23
2.3.2 <i>EEG megPAC results</i>	26
2.4 DISCUSSION	26
2.4.1 <i>Optimal filter settings</i>	27
2.4.2 <i>Frequency specificity of results</i>	27
2.4.3 <i>Alpha band D-value results for the visual RSN</i>	28
2.4.4 <i>High frequency contamination</i>	28
2.5 CONCLUSION	28
3 STUDY 2 DIFFERENTIAL DOPAMINERGIC MODULATION OF SPONTANEOUS CORTICO–SUBTHALAMIC ACTIVITY IN PARKINSON’S DISEASE	30
CONTRIBUTIONS	30
ABSTRACT	30
3.1 INTRODUCTION	31
3.2 METHODS	32
3.2.1 <i>Subjects</i>	32
3.2.2 <i>Pre-processing</i>	33
3.2.3 <i>HMM analysis</i>	35
3.2.4 <i>Estimation of the HMM</i>	35
3.2.5 <i>Statistical analysis of the states</i>	36
3.2.6 <i>Intra-medication analysis</i>	36
3.2.7 <i>Inter-medication analysis</i>	38
3.2.8 <i>Temporal properties of HMM states</i>	38
3.3 RESULTS	39
3.3.1 <i>Spontaneous brain activity in PD can be resolved into distinct states</i>	39
3.3.2 <i>Ctx–Ctx state is characterised by increased frontal coherence due to elevated dopamine levels</i>	45
3.3.3 <i>Dopaminergic medication selectively reduced connectivity in the Ctx–STN state</i>	46
3.3.4 <i>Dopamine selectively modifies delta/theta oscillations within the STN–STN state</i>	46

3.3.5	<i>States with a generic coherence decrease have longer lifetimes</i>	47
3.4	DISCUSSION	50
3.4.1	<i>Increased tonic dopamine causes excessive frontal cortical activity</i>	50
3.4.2	<i>Selective spectral connectivity remains preserved with increased dopamine levels</i>	51
3.4.3	<i>Tonic dopamine has a limited effect on local STN–STN interactions</i>	52
3.5	LIMITATIONS OF THE STUDY	52
3.6	CONCLUSION	54
3.7	SUPPLEMENTARY FIGURES	54
4	STUDY 3 LATENT COGNITIVE NETWORK TRANSITIONS TRACK SOMATOSENSORY PERCEPTUAL VARIABILITY	56
	CONTRIBUTIONS	56
	ABSTRACT	56
4.1	INTRODUCTION	57
4.2	METHODS	58
4.2.1	<i>Behavior</i>	58
4.2.2	<i>Electrophysiological recording</i>	58
4.2.3	<i>Pre-processing</i>	58
4.2.4	<i>TDE-HMM pipeline</i>	60
4.2.5	<i>HMM model fitting</i>	60
4.2.6	<i>Data-driven frequency bands</i>	61
4.2.7	<i>Group-level coherence ring figures</i>	62
4.2.8	<i>Power spectral density test</i>	62
4.2.9	<i>Diffusion map analysis</i>	62
4.2.10	<i>Uniqueness of the state-specific embedding</i>	63
4.2.11	<i>Statistical testing on the connectivity manifold</i>	63
4.2.12	<i>Transition matrix test</i>	63
4.2.13	<i>Temporal properties</i>	64
4.3	RESULTS	64
4.3.1	<i>Spontaneous pre-stimulus activity organizes into different spectral networks</i>	64
4.3.2	<i>Latent HMM states track multiple cognitive networks</i>	67
4.3.3	<i>Differential network transitions relate to varying percept</i>	69
4.3.4	<i>Faster transitions lead to correct tactile perception</i>	71
4.4	DISCUSSION	72
4.4.1	<i>Connectivity manifolds for network descriptions</i>	73
4.4.2	<i>Transitions between higher-order cognitive networks for perception</i>	74
4.4.3	<i>Limitations of the study</i>	75
4.5	CONCLUSION	76
4.6	SUPPLEMENTARY FIGURES	77
5	THESIS DISCUSSION	83
5.1	SUMMARY OF PUBLICATIONS	83
5.1.1	<i>Comparison of EEG and fMRI based resting brain networks in humans</i>	83
5.1.2	<i>Neuromodulation of spontaneous cortico–subthalamic activity in Parkinson’s disease</i>	85
5.1.3	<i>Transitions between latent cognitive networks track somatosensory perceptual variability</i>	86
5.1.4	<i>A general note on Hidden Markov Model fitting</i>	87
5.1.5	<i>Contributions of our research</i>	88
	REFERENCES	91
	STATUTORY DECLARATION	108
	APPENDIX	109
	CONTRIBUTIONS	109

Glossary

EEG Electroencephalogram

e-phys Electrophysiology

fMRI Functional magnetic resonance imaging

FPs Field potentials

HMM Hidden Markov model

ICA Independent component analysis

LCMV Linear constraint minimum variance

LFP Local field potential

MEG Magnetoencephalography

ML Machine learning

MLPA Machine learning and pattern analysis

MRI Magnetic resonance imaging

ROI Region of interest

SOA Stimulus onset asynchrony

SVI Stochastic variational inference

TDE Time delay embedding

UML Unsupervised machine learning

Abstract

Electrophysiological (e-phys) brain signals recorded in humans pose an arduous challenge both from an interpretability and analysis standpoint. Issues arise due to the very nature of the signal recorded. Known as "field potentials (FPs)," these signals arise from complex multiscale physiological properties ranging from a single neuron to a group of neurons extending all the way to the whole brain level. Hence, analysis-wise, differentiating the origin of e-phys signals is difficult. From an interpretability standpoint, ascribing neural functions to e-phys signals recovered under wildly different circumstances becomes even more convoluted. Additionally, the multiscale physiological properties endow the e-phys signals with varied statistical properties. In this thesis, we argue that various neurophysiological aspects of e-phys signals recorded in humans can be extracted by leveraging unsupervised machine learning (UML) designed to exploit the underlying statistical properties of the data.

To test our hypothesis, we used different unsupervised algorithms on varying types of human e-phys datasets. We used independent component analysis to extract spontaneous resting state networks in healthy adults, achieving results comparable to those of imaging-based methods in humans. Furthermore, we were able to retrieve information regarding the spectral characteristics of the networks that are closely related to inter-region communication in the brain. Next, we delineated the different effects a complex neuromodulator like dopamine could have on cortico-subthalamic nucleus connectivity using a hidden Markov model (HMM). thus paving the way to retrieve the spatial, spectral, and temporal properties of brain networks. In addition, with our HMM pipeline, we addressed the issue of nonstationarity of brain signals, which is often ignored during analysis. Finally, we could retrieve the effects of inter-network dynamics on behavioral variability using a task dataset. Analytically, we used a combination of HMM and diffusion maps to gain novel insights into the relationship between spectral activity and whole brain networks responsible for different types of cognitive control.

The complexity of e-phys FPs presents a unique opportunity to use modern machine learning (ML) algorithms, which have become increasingly sophisticated over the years, backed by computational hardware progress. However, we should emphasize that the current success of machine learning relies on years of human-labelled and carefully crafted datasets. This has led to high signal-to noise ratio datasets for problems that are the majority targets of ML as a field. Thus, this thesis also provides a note of caution against blindly using ML algorithms on increasingly complex low signal-to-noise ratio biological datasets, especially neuroscientific ones, without comprehending biological theory. Appreciation and assimilation of ML and neuroscientific theory promise to reveal insights into brain structure and function.

1 Introduction to the thesis

In humans, different brain regions are endowed with rich neurochemical and cytoarchitectural properties. But it is the interaction between these brain regions that is critical to giving rise to the speed, flexibility, and heterogeneity of the tasks that we as humans can achieve. We record extracellular FPs in humans using either invasive or non-invasive methods to study how different parts of the brain work together. But the brain is complicated on many different levels, from non-linear computations in a single neuron to long inter-regional white matter connectivity. This, along with the need for computation, makes the FPs recorded noisy, non-linear, and not always the same. So, it's still hard to study how different parts of the brain interact across the whole brain and then connect those interactions to cognitive, sensory, and motor processes.

In our research, we wanted to use the statistical properties of e-phys time series to figure out how different parts of the brain work together to form what neuroscientists call "brain networks." We also argue for the use of data-driven methods, which are often called unsupervised or self-supervised ML algorithms, to deal with the problems of nonlinear interactions and nonstationarity in the brain's e-phys time series. In addition, we also develop pipelines for pre-processing noisy datasets. We want to use UML to get back to well-known neuroscientific results, find new results, and figure out what the new results mean in terms of well-known principles of human behavior and neurophysiology.

As such, we begin Chapter 1 section 1(1.1): Electrophysiological neural signals by providing a background on electric signals that can be recorded at different scales in the brain. We discuss how properties like nonstationarity and non-linearity are built into the nervous system, starting from a single neuron to a group of neurons, and finally how these affect the FPs recorded in humans across the whole brain. In Chapter 1 section 2(1.2) : Human neural signals and Connectivity, we introduce the concept of "brain networks," as studied in the field of human neuroscience, and how they can be extracted from the FP recording previously introduced. We also back up our claim that the cytoarchitectural, physiological, computational, and communicational needs of the brain affect the recorded e-phys signals, giving them statistical properties. Chapter 1 section 3(1.3): Machine learning for human neural signals, we briefly revisit the concept of non-linearity and nonstationarity of e-phys signals. We talk about the idea behind how machine learning algorithms work and warn against blindly using unsupervised methods in neuroscience or medicine. Lastly, we talk about the goals of the thesis and go over three key publications, in chapters 2,3 and 4, that helped us reach those goals. In the discussion (Chapter 5), we give a critical analysis of our publications and explain what our research has added to neuroscience.

1.1 Electrophysiological neural signals

The human brain is never at rest. Even when we sleep, the brain produces a constant stream of neural activity that controls our sleep cycles and regulates the sleep-wake mechanisms. During rest, the brain continues to generate spontaneous activity associated with cognition even when we are not actively performing an overt task. Ongoing brain activity is sculpted in a complex manner. From a single neuron to the whole brain, different neural entities (neurons, groups of neurons, large brain regions, or the entire brain) receive multiple inputs, process the inputs, and send out outputs. The outputs are either sent to the neural entity downstream or cause a perceptual or behavioral output that lets an organism interact with its environment. Input-output relationships are non-linear at every level in the brain (Cifre et al., 2021; Ecker et al., 2011; Shamir and Sompolinsky, 2004; Sheth et al., 2004; Yang et al., 2021). The relationship is controlled by what goes into a neural entity and how that neural entity is structurally wired. This is evident in the brain at multiple scales. In the case of a single neuron, for example, inputs from other neurons are transmitted in the form of electrical impulses. The neuron receives many inputs via dendrites. However, single neuron dendrites have a diverse range of structural arborization patterns and a complex physiology that can influence how neuron inputs are processed at the dendritic level (Gidon et al., 2020; London and Hausser, 2005). The input-output processing then changes the action potential of the single neuron, which affects other neurons further downstream. A single neuron does not have sufficient computational power to give rise to the diverse behavior exhibited by, for example, the higher vertebrates. But the complexity of dendritic single-neuron computation is needed for a group of neurons that work together to create a more complex input-output relationship that influences the behavior of an organism (Poirazi and Papoutsis, 2020).

Over the past decade, it has become increasingly apparent that a group of neurons, often referred to as a "neural population," produces neural activity responsible for computation and inputs to downstream brain regions (Panzeri et al., 2022). Coupled neural populations enable cognitive and motor tasks (Churchland et al., 2012, 2010; Raposo et al., 2014). Different structural connectivity motifs exist at the population level as well. Typically, we see feed-forward, feedback, and recurrent connectivity within and across neural populations (Kreiman and Serre, 2020; Layton et al., 2014; Miska et al., 2018; Roelfsema and Holtmaat, 2018; Semedo et al., 2022; van Bergen and Kriegeskorte, 2020; Wyatte et al., 2012). These connections constrain the types of computations a neural population can perform. Also, connectivity patterns enable selective communication across different neural populations (Semedo et al., 2019). Furthermore, the dendritic patterns differ in a region-specific manner in the brain, adding to the functional specialization of different brain regions (Fleischhauer, 1974; Helmstaedter et al., 2009; Kawaguchi et al., 2006; Leguey et al., 2018). To probe the relationship between different neural populations, modern neuroscience experiments, typically in non-human primates (NHPs) and rodents, record large neural populations from a single or multiple brain regions using microelectrode arrays (Urai et al., 2022). The data from such large-scale neural recordings is collected as neural spiking activity over time (Humphries,

2021). Given the multitude of functions the neural population(s) undertake, it becomes an analytical challenge to extract interpretable representations from the spiking activity recorded. But it has been shown that physiological and structural constraints on a neural population can produce specific activity patterns that have discernible statistical properties (Goris et al., 2014; Litwin-Kumar and Doiron, 2012; Pillow et al., 2008; Shadlen and Newsome, 1998; Wilson and Cowan, 1972). As a result, recent breakthroughs in ML have been utilized to split large-scale spiking brain activity into physiologically understandable representations. (Glaser et al., 2020; Smith et al., 2021; Williams and Linderman, 2021). ML methods have basic and translational capabilities in understanding the brain (Kemere et al., 2008; Santhanam et al., 2006). ML methods have been successful at parsing neural population activity into behaviorally relevant and irrelevant clusters (Hennig et al., 2018; Kaufman et al., 2014). According to Semedo et al. (2019), ML techniques have demonstrated how neural population activity from a particular region can promote communication between various neural populations without affecting its computational resources. The methods have been applied to explain the variable timing and initiation of tasks in rodents (Ashwood et al., 2022; Bagi et al., 2022; Recanatesi et al., 2022; Wiltschko et al., 2015). To create an effective method to assess the collected data for each of these applications, the authors relied on a combination of physiological and biological insight and the statistical features of brain signals.

Neural population activity provides data that resembles the fundamental electric signal via which neurons perform their functions. But current recording methods involve invasive procedures that are ethically challenging to perform on healthy humans or patient participants. Some patients have had successful brain-computer interface (BCI) systems implanted, but no country has approved a standard care procedure for BCI-based therapies (Moses et al., 2021; Santhanam et al., 2006; Willett et al., 2021). To understand neural activity in the human brain, a number of non-invasive methods have existed for a long time. They are namely functional magnetic resonance imaging (fMRI), electroencephalography (EEG), and magnetoencephalography (MEG). Other techniques, such as functional near-infrared spectroscopy (fNIRS), also exist and are becoming increasingly popular. None of these techniques has the spatial resolution to record neural data at the level of a single neuron or neural population. But with the help of sophisticated physics-based modeling, we can obtain data from different cubic millimeter-sized locations on the brain commonly termed "voxels" (for fMRI and fNIRS) and/or sources (EEG and MEG) (Baillet, 2017; Henson et al., 2009; Munck et al., 2012). About a hundred thousand neurons and a million synapses are packed into a cubic millimeter of brain tissue. So, the neural data collected by the above methods is a signal from all of the thousands of neurons in a millimeter-sized cube of brain tissue. We don't have a complete understanding of how single neural action potentials transform into aggregate noninvasively recorded human brain signals (Baillet, 2017; Pesaran et al., 2018). In addition to non-invasive methods, we occasionally have the opportunity to collect invasive data from human participants. One such opportunity arises when Parkinson's disease patients are undergoing surgery for the implantation of deep brain stimulation (DBS) devices to manage the neurological symptoms of the disease (Okun, 2012). A DBS electrode delivers electrical

stimulation to a specific brain region, but it can also be used to record electrical signals (Litvak et al., 2021; Oswal et al., 2016). As in the case of non-invasive human brain signals, invasive DBS electrodes do not have the spatial resolution of multi-electrode arrays. But given the surgical precision during the implantation of the DBS electrodes, the electrical signals can be better localized to specific brain regions as compared to EEG and MEG signals. Hence, using DBS electrodes, we are recording the aggregate signals of a few thousand neurons. For this thesis, we will only talk about brain signals from humans that were recorded with MEG, EEG, and DBS electrodes. Signals recorded using these methods are collectively called extracellular FPs (Buzsáki et al., 2012). Due to the better localization afforded by precisely implanted DBS electrodes, signals recorded using DBS electrodes are termed "local field potentials," or LFPs.

1.2 Human neural signals and connectivity

As mentioned previously, FPs capture the electrical activity of thousands of neurons (Pesaran et al., 2018). Hence, we cannot study the properties of a neural population. Instead, we focus on the properties of macro-level regions termed "sources" or "voxels" in the brain. There are multiple contributors to the FP signals recorded in the brain. Some of these contributions occur from single-neuron electrical characteristics, such as synaptic activity, calcium, and sodium spikes (Golding and Spruston, 1998; Kim et al., 2015; Tsubokawa et al., 2000), and neuron-glia interactions (Fields and Stevens-Graham, 2002). Two major factors seem to affect the extracellular FPs strongly: local neural architecture and the temporal synchrony of large neural populations (Finn et al., 2019; Helmstaedter et al., 2009; Huntenburg et al., 2017; Khambhati et al., 2018; Paquola et al., 2019; Seidlitz et al., 2018; van den Heuvel and Yeo, 2017). The spatial architecture of neurons is most prominent in the cerebral cortex. Cortical neurons are arranged in layers and have distinct projections to downstream structures (Baker et al., 2018; Gerfen, Economo and Chandrashekar, 2018). In the visual cortex, for example, each layer can have different contributions to the FPs (Xing, Yeh and Shapley, 2009). Local architectural heterogeneity also exists in the connectivity patterns of single neurons, which can lead to functionally distinct cortical microcircuits (Gerfen, Economo and Chandrashekar, 2018). For our research, we did not have access to cytoarchitectural details of the brain; hence, we focused on the contribution of temporal synchronicity to FPs.

Temporal synchronicity arises in FPs when large and spatially separated neural populations coordinate their spiking patterns to accomplish a specific task (Uhlhaas et al., 2009). Synchronization between vast neuronal populations adds significantly to FPs (Buzsaki, 2006; Olejniczak, 2006; Uhlhaas et al., 2009). The contribution of synchrony to FPs is clearer in the cortex than in subcortical regions. The clarity could be because the cerebral cortex recruits spatially distributed regions to accomplish computations. In addition, communication between distinct cortical regions involves oscillatory activity, which is reflected in the power spectra of the FPs (Fries, 2015). Synchronicity, neural oscillations, and their relationship to neural spiking and computation bring us to the concept of functional networks in the brain. Even though stable canonical networks have been found in spontaneous "resting-state" brain activity (Biswal et al., 1995; Cordes et al., 2001; Greicius et al., 2009), the definition of brain

networks, especially for human brain data, is continually evolving because of new ways to study human brain datasets. The basic premise of functional networks involves the temporal coordination of spatially distinct brain regions that work together to accomplish a specific behavioral or cognitive task. Temporal coordination is assessed at the neural population level through sample covariance between neural signals from the set of recorded neurons. This means that the information in the multidimensional covariance matrix aids in determining which neurons in the recorded population collaborate and which function independently. Covariance at the population level of neurons has a physiological meaning because neurons connected by a single synaptic connection tend to fire together more when they are working on a task. Synchronous firing patterns in turn result in a larger shared covariance between such neurons (Diesmann et al., 1999; Gill et al., 2020; Montani et al., 2009; Reyes, 2003). Hence, the covariance between a set of neurons serves as a proxy for connectivity between neurons.

Covariance has been used as a proxy for connectivity in FP recordings of human data. The covariance principle works because large brain regions are physically connected via white matter tracts (Basile et al., 2022; Warling et al., 2021). The white matter tracts are a collection of myelinated axons that give rise to a structural architecture of the brain along which signals can travel at high speeds between brain regions. This structural architecture is referred to as the "structural connectivity map of the brain." Different brain regions exhibit different structural connectivity architectures; hence, the structural connectivity map helps produce distinct co-activation patterns between spatially distant brain regions. The co-activation patterns lead to different synchronous dynamics that can be extracted from the FPs recorded in humans. Lastly, sample covariance can be used on the regional FP time series to figure out which parts of the brain are working in sync with each other. The covariance patterns obtained can be interpreted as functional connectivity (FC) maps of the human brain.

1.3 Machine learning for human neural signals

We have discussed how neuroanatomical and neurophysiological properties affect FPs recorded at different scales. When neuroscientists look at FPs, it is important to remember that they are the result of computations in the brain. Numerous brain networks work together to carry out the computations. A single neuron is a non-linear computing unit for several large-scale brain networks. This is because it has complex dendritic patterns and receives input from nearby neurons at different times. So, networks of these non-linear units that produce a wide range of cognitive activities are likely to act as non-linear computational entities. So, there will always be a non-linear relationship between the FPs collected from neurons, groups of neurons, major brain regions, or the entire brain when trying to figure out how different neural entities relate to each other at any scale. But in practice, first, due to a lack of concrete theory in human neuroscience and second, due to issues with the interpretability of current non-linear methods, regional relationships are usually studied using linear methods. Also, the current linear methods are good at explaining behavioral and neural differences in neuroscience (Kriegeskorte and Douglas, 2019).

Nonstationarity is a hallmark of most neuroscience time series data (Guan et al., 2020; Jones et al., 2012; Kaplan et al., 2005). Once again, a signal is simply an observation of an underlying process. Hence, it is not the signal but the generating process that is inherently nonstationary. Essentially, if the properties of the generative process change over time, then the process is said to be nonstationary. FP nonstationarity can be explained by anatomical connectivity, cellular architecture, temporal synchronicity, and, finally, the purpose of computation and communication in the brain. Given the speed and flexibility of the computations in the brain, even in the absence of external influences, the generative process of the observed FPs undergoes rapid transitions. Furthermore, in complex systems such as the brain, there is a large range of time scales at which numerous processes evolve to give rise to the final dynamics. Hence the parameters governing the dynamics of the FPs become sensitive to the various time scales (Indic et al., 1999), thus giving rise to nonstationarities in the FPs. When it comes to the analysis of neural data, nonstationarity becomes crucial since the linear or nonlinear methods measuring relationships between regional time series make implicit assumptions about the probability distributions of the data examined. Hence, if nonstationarity is left unaddressed, it could lead to misleading conclusions regarding brain functions. Therefore, a possible direction of analysis is to harness the nonstationary and non-linear properties of the time series and link them to interpretable physiological motifs to gain insight into brain functioning.

Since its inception, the objective of machine learning and pattern analysis (MLPA) has been to exploit the statistical properties of the underlying datasets. The natural world, on multiple scales, is endowed with rich statistical properties. Hence the signals recorded from the natural world are already endowed with those properties. A simple example is the identification of different trees based on the patterns of their leaves. We can call the leaf patterns features of the signal that are helpful in identification. Machine learning algorithms have long used handcrafted rules (Viola and Jones, 2001) or different layers of neural networks (Krizhevsky et al., 2012) to take advantage of these features. Audio signals (Hershey et al., 2017), human and computer languages, films, and images all contain such examples (Brown et al., 2020; Szegedy et al., 2015; Vaswani et al., 2017). As a result, the ability of any MLPA algorithm to extract these statistical regularities or features is critical to its effectiveness. It is crucial to remember that the final goal of MLPA algorithms might be whatever the user desires, such as categorization, auto-generation of code, or the creation of original graphics depending on user input.

An MLPA algorithm finds features suitable to meet the end objective during what is known as the "training" or "learning" phase of the algorithm. Primarily, there are two types of learning strategies followed in machine learning: supervised and unsupervised/self-supervised learning. The main difference lies in the kind of dataset used in supervised vs. unsupervised learning. In supervised learning, the algorithm "learns" from the training data by iteratively making predictions on the data and adjusting for the correct answer. Humans supply the correct response as part of the dataset utilized by the algorithm. Unsupervised algorithms, on

the other hand, operate on their own to identify the intrinsic structure of unlabeled data. Supervised methods are time-consuming to train and require labels for input and output data, which relies on expertise. Meanwhile, unsupervised algorithms can produce inaccurate results unless we have human intervention to validate the outputs produced by the algorithm.

An implicit and often ignored fact in the current state of the art of machine learning is that the datasets on which algorithms or "models" are trained belong to categories of problems for which some ground truth is universally agreed upon. We can call this the implicit ground truth dataset assumption. This assumption is present irrespective of the learning strategy used to train an algorithm. For example, we can take the problem of image classification. In the supervised scenario, we can have labeled datasets. Then we train an algorithm to classify images, and if it gets it wrong, the expert labels guide the algorithm. In the unsupervised case, an algorithm learns to group images with similar features together and label them. For example, modern state-of-the-art algorithms can easily distinguish dog images from cat images, among other things. But here is where the implicit ground truth dataset assumption we talked about plays an immense role. The unsupervised algorithm's clustering result is human-verifiable since humans agree on categories like dog and cat. This is true for learning a language, playing computer games, creating art, etc.

The implicit ground-truth assumption cannot be made about neuroscientific data. We do not have a concrete theory in neuroscience that supports an exclusive hypothesis. Multiple e-phys patterns that emerge in datasets are open to different interpretations. Therefore, e-phys data must be supported by either experimental results or current neurophysiological understanding. Even in medicine, there is ample noise in the classification of diseases and prognoses based on patient data. Hence, creating an incorruptible, labeled dataset is extremely challenging. Furthermore, unsupervised outcomes and/or features discovered are subject to highly variable interpretations. Furthermore, when machine learning algorithms are designed to deal with medical patient data, the assumption of an absolute ground truth cannot and should not be made.

We have discussed that there are hidden statistical properties in e-phys data that can be linked to meaningful neurophysiological underpinnings of brain function. Furthermore, machine learning methods are intended to exploit the inherent statistical structure of data in order to achieve user-defined end goals. We also mentioned how, when it comes to neuroscientific data, relying on supervised labeled datasets may not be the best option. As a result, the goal of this thesis was to use UML methods to find intrinsic structure in human e-phys datasets:

- by developing principled pipelines for pre-processing noisy e-phys datasets.
- by recovering neurophysiologically established results in the field while using unsupervised algorithms.
- by discovering new neuroscientific findings by exploiting the statistical properties of human datasets.

- by interpreting the findings in the context of established principles of human behavior or neurophysiology.

We were able to reach the above goals thanks to the three scientific studies that the doctoral candidate wrote as the first author during the doctoral research program. The first study dealt with the application of independent component analysis (ICA) on resting state (spontaneous) EEG datasets acquired from healthy young adults. ICA is a blind source separation (BSS) algorithm that decomposes a signal into statistically independent components. Hence, it qualifies as an unsupervised learning approach. The first BSS application in fMRI (Mckeown et al., 1998) was motivated by principles of localization and connectionism in the brain (Phillips et al., 1984). Localization meant that different parts of the brain were used for different tasks. For example, the precentral and postcentral gyri are used for somatosensory and motor tasks, respectively. On the other hand, connectionism states that different parts of the brain collaborate to produce brain functions. For example, the visual and auditory cortical regions might work in tandem to help localize the source of a sound. Given our previous discussion regarding neural architecture, white matter tracts, and functional connectivity, we can expect that there is a neurophysiological basis for the underlying e-phys signals to decompose into statistically independent components. This would be a reasonable assumption since the statistical properties of the e-phys signals would depend on how many local or distributed processes are present in the neural activity underneath. Hence, spatial ICA would be able to show independent spatiotemporal patterns of signal covariation, which could be caused by processes that are local or spread out in the brain and drive neural activity. Another data-driven approach rooted in neurophysiological principles of phase-amplitude coupling (PAC), called MEG-PAC (Florin and Baillet, 2015), was applied to the same healthy EEG dataset.

In the second study, we looked at the statistical nonstationarity of brain data and the effects of neuromodulators from the point of view of neurophysiology. So far, we have discussed how nonstationarity is a feature of e-phys data due to the signal being affected by multiple factors. What was missing was the impact of neuromodulators on the statistical properties of the signal over time. Neuronal firing rate is altered by ascending projections from the brainstem and subcortical nuclei, making neurons more or less responsive to incoming signals (Bell and Shine, 2016; Shine, 2019). The resulting adjustments in neural gain or suppression will result in a nonstationary brain signal. Moreover, different neuromodulators can produce either targeted or diffuse modulation, which is exerted over large temporal scales (Gershman and Uchida, 2019). For the second publication, we relied on our ability to acquire simultaneous cortical MEG and subthalamic nucleus (STN) LFP recordings in Parkinson's disease (PD) patients undergoing surgery for deep brain stimulation (DBS) therapy. We recorded PD patients with and without the influence of dopaminergic medication (fast-acting L-DOPA). To get around the statistical problem, we used a hidden Markov model (HMM) that could split the time series data into separate time intervals. This made it easier to evaluate how dopaminergic modulation changed spatial (cortico-STN) and temporal properties.

For the third and final study, we aimed to examine the relationship between spontaneous whole-brain networks and behavioral variability in healthy humans. According to research, the neural activity recorded a few milliseconds before the task begins contains a variety of information that can predict variability and task performance (Iemi et al., 2019; Li et al., 2020; Linkenkaer-Hansen et al., 2004; Podvalny et al., 2019; Rajagovindan and Ding, 2011; Rassi et al., 2019). Usually, the inter-trial interval is kept short for task-based studies. But it prevents us from delineating the impact of long-duration spontaneous brain activity on task-related variability. Furthermore, we are unable to probe the flexibility with which the brain might work. To address these issues, we used an established tactile discrimination task and set the inter-trial interval randomly between ten and fifteen seconds, which was significantly longer than what has been used in previous perceptual studies. The perceptual task we employ does not warrant any explicit cue-based attention modulation or suppression of extrinsic distractors. It was not a reward-based paradigm where the subject had to spend cognitive or memory resources to reach a decision. As a result, to properly respond to an incoming tactile stimulus, the internal dynamics of spontaneous activity preceding the stimulus had to be configured flexibly to make the best use of the limited external information that arrived in the form of stimuli. To characterize the temporal flexibility, spatiality, and spectral characteristics of spontaneous neural activity, we expanded on our previous HMM framework. We also used diffusion maps to describe the relationship between the spectral signatures of different HMM states and whole-brain networks. Diffusion maps are a class of non-linear dimensionality reduction techniques (Coifman et al., 2005). Intuitively, they convert connectivity between two brain regions into a distance in Euclidean space. Using a mix of two UML algorithms, we revealed a unique association between various whole-brain networks, their spectral signatures, temporal features, and their relationship with tactile discriminating performance. Our findings show how effective UML can be at comprehending human e-phys data.

Overall, we discussed how the electrical properties of neural entities, including individual neurons and neural populations, affect the functional behavior of large-scale human brain networks. Furthermore, we stressed how multiscale structural connectivity constraints might play a critical role in governing functional connectivity. The statistical properties of brain activity arise from functional synchronization and anatomical constraints, which may be exploited by UML techniques. In our publications, we elucidate how UML can be used to reveal the spatiotemporal properties and spectral signatures of e-phys time series recorded under different conditions. The different conditions included pure resting state activity, resting state under neuromodulatory conditions, and pre-stimulus resting state activity subjected to a tactile discrimination task. In the discussion of the thesis, we summarize our findings and point out the limitations of our results.

In the following sections, we present the complete manuscripts for the three publications.

2 Study 1 Evaluating a data driven pipeline for extracting EEG resting state networks in humans

Contributions

Abhinav Sharma: Conceptualization, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (55% contribution)

Dr. Esther Pelzer: Analysis of fMRI data

Dr. Esther Florin: Conceptualization, Software, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing

Abstract

There exist multiple ways of extracting electroencephalography (EEG) based resting state networks (RSNs). Yet there is no consensus on an approach to derive these networks in a way that is comparable to the canonical RSNs observed in the fMRI literature. Ideally we would extract EEG-RSNs that are spatially comparable to their fMRI counterparts while leveraging on the EEG ability to reveal the frequency characteristics of the RSNs. In our opinion such an approach should include the following essential factors a) Independent component analysis (ICA)-fMRI networks, b) whole brain data driven EEG networks without any atlas constraints, c) networks extracted independently in different frequency bands, d) statistical tests between source level EEG and fMRI RSNs across all frequencies. Furthermore, we need a principled data driven way of extracting networks where cross frequency interaction could play a role. In this paper we test Hilbert envelope based ICA and phase-amplitude coupling based approaches to extract EEG-RSNs. We found that the envelope based approach provides the best spatial correspondence with fMRI-RSNs but none of the EEG networks were dominated by a specific frequency.

2.1 Introduction

Resting state networks are defined as temporally correlated networks of brain activity which emerge without the presence of any explicit stimuli or tasks (Biswal, 2012). Though these networks emerge spontaneously they are robustly determined from fMRI recordings across centres and studies (Buckner et al., 2011; Doucet et al., 2011; Yeo et al., 2011) and they have been associated with multiple cognitive processes (Fox and Raichle, 2007; Cole et al., 2010; Deco and Corbetta, 2011; Laird et al., 2011). The two most popular analysis approaches include the independent component analysis (ICA) (Biswal and Ulmer, 1999) and a seed based correlation approach (van den Heuvel and Hulshoff Pol, 2010; Chen and Glover, 2015). Both approaches yield canonical resting state networks (Calhoun et al., 2001). Despite these well-established canonical fMRI RSNs a link between electrophysiology and RSNs has been difficult to underpin and is an area of intense research.

On the electrophysiology front EEG studies have used diverse pipelines for analysis of resting state data. However, most of the studies did not aim at identifying the canonical RSN. Mantini et al. (2007) were the first to directly address the relationship between EEG based electrophysiological rhythms in different frequency bands and ICA derived fMRI networks. But no EEG networks were derived and hence a direct spatial relationship between EEG and fMRI networks could not be assessed on either the sensor or source level. Jann et al. (2010) produced topographic sensor level results that demonstrated the spatial and spectral properties of EEG data and fMRI-ICA results. Moving to source-reconstructed EEG data Congedo et al. (2010) used source level ICA to extract RSN but did not directly compare to fMRI RSNs. Two further studies used ICA on EEG data with the objective to examine EEG networks formed by frequency interaction. Specifically, Aoki et al. (2015) applied eLORETA-ICA but they did not compare the EEG networks directly to fMRI results. Given the way ICA is applied in Aoki et al. (2015) each independent component comprises information from multiple spatial locations and frequency bands. For example, independent component number 5 is called the ventral visual pathway. The ventral visual pathway in Aoki et al. (2015) is comprised of right occipitotemporal cortex in the alpha band and the right ventral prefrontal cortex (vPFC) in the beta band. Hence Aoki et al. (2015) report cross frequency interaction based on a mixture of frequencies recovered from their ICs. On the other hand, Sockeel et al. (2016) perform spatial ICA on source level data concatenated across all frequency bands. The resultant ICs contain spatially distinct information with contributions from all frequency bands in each and every IC. Furthermore, Sockeel et al. (2016) also compared their results to the fMRI networks. Overall these studies, due to the nature of their objectives, did not delineate the frequency specificity of the EEG networks. Deligianni et al. (2014) aimed at directly assessing the relationship between covariance matrices derived from source level EEG and fMRI data using partial correlation and canonical correlation analysis (CCA) in different frequency bands. However, for their analysis they used an atlas based approach which does not allow comparing their networks to the canonical fMRI networks. Finally, a recent study covering a variety of functional connectivity measures tried to assess

the relationship between EEG source space connectivity and resting state fMRI based connectivity (Rizkallah et al., 2020). In this study the process of deriving connectivity between brain regions was restricted by an atlas and connectivity was derived using correlation.

Therefore, to advance our understanding of the relationship between RSNs derived from EEG and fMRI a direct comparison of ICA-fMRI networks with whole brain data driven EEG networks without any atlas constraints of the same subjects is necessary. Furthermore, statistical tests should be performed to compare source level EEG and fMRI RSNs in addition to visual inspection of derived RSNs. Additionally, extraction of networks formed by interaction of different frequencies should be based upon physiological processes rather than assuming a linear mixture across frequencies. To tackle some of the missing parts in the previous literature we derived resting state networks from EEG recordings using data driven approaches at the source level without constraining the analysis to an atlas. The EEG ICA networks were extracted individually for each frequency band and statistically tested to investigate their frequency specificity. To investigate potential cross frequency interaction in the RSN we also used the megPAC approach (Florin and Baillet 2015). As a yardstick for performance of the EEG pipeline we compared the results to ICA networks that we derived from fMRI recordings of the same subjects. By providing a validated pipeline to extract RSNs and their frequency content, we hope to provide a tool which can also be used in clinical settings.

2.2 Methods

Our study comprised 26 healthy right-handed male subjects [age: 26.7+/-3.9 SD; laterality index 88.6+/-20.7; mini mental state test: 29.8+/-0.5]. We excluded 3 subjects due to movement artefacts during data acquisition. All subjects gave their informed consent in line with the ethical guidelines of the declaration of Helsinki (Ethics committee Cologne: 14-264, Ethics committee Düsseldorf: 5608R) and were then included to participate in the study. All participants were measured once with EEG, and once with fMRI while resting for 30 minutes. In both modalities they were asked to not move and fixate on a paper cross.

2.2.1 Structural MRI acquisition and pre-processing

All magnetic resonance data was obtained on a Siemens 3T PRISMA scanner using a 64-channel head coil. High-resolution T1-weighted images were acquired with a 3D MPRAGE sequence (TR = 2300ms, TE = 2.32ms, ES = 7.2 ms, FA= 8°, FOV = 230mm x 230mm, isotropic pixel resolution of 0.9 x 0.9 x 0.9 mm, slice thickness of 0.9 mm, 192 slices). Brain and cortex surface were extracted with FreeSurfer (5.3.0) from the T1-weighted images.

2.2.2 EEG data acquisition and pre-processing

EEG measurement was performed with a 128 EEG channel actiCAP with a sampling rate at 2.5Khz. The impedance was controlled to be less than 25 kOhm for each electrode. Electrode positions were digitized with a Zebris ultrasound device. All participants sat in the EEG chamber and fixated on a paper-cross.

We collected a total of 30 minutes of resting state activity for each subject. Subjects were seated for the duration of data collection and fixated on a cross printed on paper which was placed in front of the subjects. EEG data was acquired in blocks of ten minutes. After each block subjects were permitted to take a break. The ten-minute block duration was in line with previous research (Liuzzi et al., 2017). In addition to the EEG we also recorded an electrocardiogram (ECG) and electrooculogram (EOG). Moreover, we ensured that participants stayed awake by monitoring them through a camera system.

Post-acquisition data pre-processing was done using Brainstorm (Tadel et al., 2011). ECG and EOG signals were used to remove heartbeat and eye blink related artefacts respectively using signal space projections (SSPs). Line noise (50Hz) and its harmonics were removed by using notch filters. Power spectral density for all the channels was calculated using Welch's method and sensors with high noise levels were excluded from further analysis. Data was then visually examined for further artefacts which were then removed. The cleaned data was down sampled to 1000Hz.

The EEG forward model was created with the symmetric Boundary Element Method (BEM) in the open source software OpenMEEG. The inverse problem was solved on the individual cortical surface with a cortically constrained weighted minimum norm estimation (wMNE) using default settings in brainstorm where an identity matrix was used for the noise covariance. Individual cortical surfaces were obtained from the individual MRIs of each subjects using FreeSurfer software with a tessellation of 15002 vertices. Coregistration with the EEG data was based on the three landmarks (nasion, right and left pre-auricular point).

2.2.3 Extracting EEG resting state networks using ICA

EEG resting state networks were extracted using source level data. We used the Envelope ICA approach described in Brookes et al. (2011). For each subject the Hilbert envelope was extracted for the five most common electrophysiological bands for each source time series: delta (1-4 Hz), theta (4-8Hz), alpha (8-12 Hz), beta (12-30 Hz), and gamma (30-50 Hz). Following the approach by Brookes et al., (2011) we filtered the data across the whole frequency range of each frequency band. In order to account for the 1/f characteristics of the EEG signal (wide band) we also tested a filter option in which the bin size was adjusted to extract the envelope (small band). From 1-4 Hz we band-pass filtered the data in 1 Hz bins, from 4-30 Hz in 2 Hz bins, and above 30 Hz in 5 Hz bins. Afterwards each of these frequencies was Hilbert transformed. The signal was then averaged across bins in every band individually. After filtering the envelope time series were down sampled to 1Hz. Moreover, we tested for the influence of infinite and finite impulse response filters.

After determining the Hilbert envelopes at the individual subject level they were projected to the Colin 27 brain. The projected data was spatially smoothed with a 5mm Gaussian Kernel and then zscore-normalised in the time domain. For each frequency band, data from all subjects was concatenated along the temporal dimension as is commonly done in spatial group ICA approaches. Subsequently temporally independent components were calculated

using the FastICA (Hyvärinen, 1999) algorithm. Before ICA the data were pre-whitened with a principal component analysis of 30 components. As the results of an ICA are dependent on the starting seed they are variable with every new iteration of FastICA. In order to stabilise ICA output, we used the ICASSO algorithm (Himberg et al., 2004). FastICA was run 50 times in ICASSO. The resulting independent components from each ICASSO run were clustered based on the absolute value of the linear correlation coefficient between components from different repetitions. ICASSO then provides the centrotpe which is the best representative estimate of a cluster of independent components from the different runs of FastICA.

Finally, to obtain spatial resting state networks the top 20 representative ICs (centrotpe from ICASSO) were correlated with the envelope data. For visualization, the correlation networks were plotted on the Colin27 brain. These correlation networks of each frequency band were then compared to the resting state fMRI networks.

2.2.4 Extracting EEG resting state networks using megPAC

megPAC approach was applied to the source reconstructed EEG data. For the megPAC approach we used the parameters as described in the paper Florin and Baillet (2015). Chirplet transform was used to calculate phase and amplitude of the source signals (Mann and Haykin, 1995). The first step involved calculation of low-frequency phase that couples most strongly to the high gamma amplitude from 80-150 Hz based on phase-amplitude coupling (Özkurt and Schnitzler, 2011) for each source time series of each subject. Between the identified low frequency peaks and troughs, the gamma amplitude (80-150Hz) was interpolated. These new time series were down-sampled to 10 Hz and then projected to the Colin27 brain. On the Colin27 brain the cortical time-series from all subjects were spatially smoothed (5mm Gaussian Kernel) and then temporally concatenated. Subsequently the spatial correlation matrix between all time-series was calculated. Finally, the resting state networks were determined as the principal spatial modes based on a singular value decomposition.

2.2.5 fMRI data acquisition and extraction of resting state networks

Resting fMRI BOLD (blood oxygenation level-dependent) data were recorded with echo-planar-imaging (EPI) acquisition, (TR= 776 ms, TE = 37,4 ms, flip angle = 55°, resolution 2.0 x 2.0 x 2.0 mm, slice thickness of 2.0 mm, 72 slices). The resting fMRI scan lasted ~ 30 min. Subjects were also asked to fixate on a paper cross during the scan. MELODIC 3.0 (Multivariate Exploratory Linear Optimized Decomposition into Independent Component) was then used to extract independent components after affine registration to the structural space. FSL was used to process functional images. Affine registration to the MNI-152-2mm standard space was performed using FLIRT. To constrain fMRI results to the cortical surface we used a binarized cortical mask obtained from FreeSurfer. fMRI time series was variance normalised and then MELODIC was used to extract 20 independent components. After IC decomposition, we chose the standard threshold of 0.5 for the IC networks following the recommendations in FSL Melodic. This threshold enables rejecting noise from the estimated ICs using the MELODIC pipeline. In order to compare fMRI datasets with EEG resting state results we finally

registered individual IC components, located in the MNI-152-2mm standard space, to the Colin27 brain. For complete description of the fMRI dataset please refer to Pelzer et al. (2021).

2.2.6 Comparison of the EEG with fMRI resting state networks

We used two different procedures to measure the spatial correspondence of fMRI and EEG RSNs. First we calculated spatial correlation between all fMRI ICs and all EEG RSNs. We then selected the maximum value of this correlation for each fMRI IC to find the best correlated EEG IC frequency in each frequency band. As second measure we used binary overlap measure “D”. For this the EEG networks were thresholded at 0.3 correlation value. fMRI ICs were thresholded using probability networks produced by FSL. A threshold value of 0.85 was used. This threshold is used to select voxels within the fMRI ICs that have a high probability of belonging to a particular IC.

$$D(EEG_i, fMRI_j) = \frac{EEG_i \cap fMRI_j}{EEG_i \cup fMRI_j}$$

(Mesmoudi *et al.*, 2013)

The overlap is calculated for each cortical source. If there is a perfect spatial overlap between an EEG map i and an fMRI map j , D will be 1. For each fMRI-map the EEG-map with the highest D is determined.

2.2.7 Statistics for comparison of networks

To compare the EEG derived networks we bootstrapped across all subjects 100 times and extracted networks via ICA. The bootstrap was performed for each frequency band separately. We calculated the spatial correlation and the D value for each iteration which gave us a distribution for these values. We then performed a one-way ANOVA to determine the effect of frequency on the spatial correlation and D metrics for each Envelope-ICA based map. Furthermore, we tested post-hoc if a specific frequency band significantly stood out from all others for a resting state map.

2.3 Results

Out of the 20 independent components extracted from fMRI data we used the default mode, sensorimotor, visual, fronto-parietal left, fronto-parietal right and frontal resting state networks for further comparison with EEG networks. We show the fMRI and EEG networks in figure 2.1 for four networks. As compared to the fMRI networks EEG networks show less spatial selectivity.

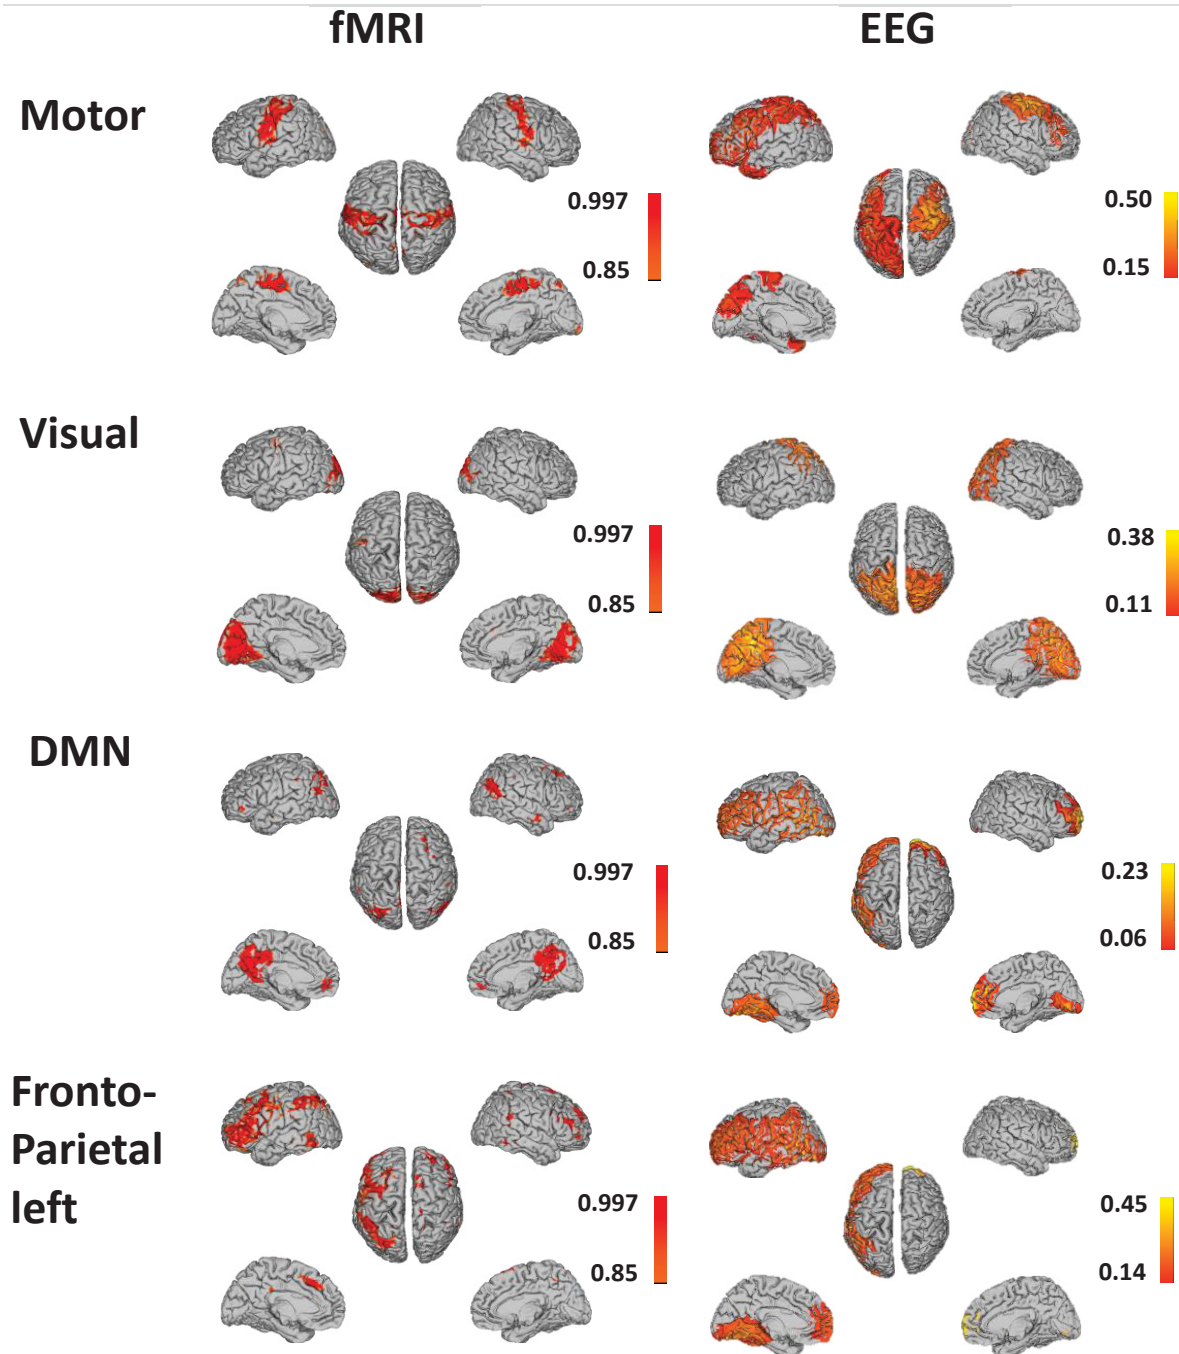


Figure 2.1: Resting state networks extracted using fMRI and best matched EEG RSNs. The EEG RSNs were derived using Hilbert envelopes using the small band iir filter.

2.3.1 Bootstrapped EEG-fMRI comparison

As explained in the methods we bootstrapped across frequency envelopes and extracted independent components for different filter settings. Bootstrapping helped us delineate both the effect of filter type and frequency specificity of RSNs extracted from EEG. For spatial correlation values the type of filter had a significant effect for all RSNs (Table 2.1) ($p < 0.01$). For the frontal and sensorimotor networks there was significant interaction between filter type and frequency of the envelope ($p < 0.01$). Post-hoc testing revealed that the small band iir filter led to significantly higher spatial correlation with fMRI RSNs for all networks except for the parietal network ($p < 0.01$).

For D-value based comparison between EEG and fMRI RSNs there was a significant effect of filter type ($p < 0.01$). We also observed a significant interaction between filter type and envelope frequency for all networks except the sensorimotor one ($p < 0.01$). Post-hoc testing revealed that small band fir filter produced the best D-value correspondence between EEG and fMRI RSNs for the fronto-parietal left and right, visual and the default mode network. No specific filter type-frequency pairs were dominant post-hoc for the D-value based comparison.


Overall the small band iir filter was the optimal filter to extract EEG networks based on spatial correlation comparison with fMRI networks whereas the small band fir filter was the best when using the D-value based comparison. Since spatial correlation is a widely reported metric of comparison and it captures spatial variability better than the binary D metric our further analysis and results were calculated on the envelopes derived using the small band iir filter.

Using the small band iir based envelopes there was an effect of envelope frequency for all networks except the left fronto-parietal map based on spatial correlation ($p\text{-val} < 0.05$; left fronto-parietal $p\text{-val} = 0.19$). However, post-hoc testing revealed that for a given map no specific frequency stood out from the others for a given networks. For the D-value based comparison the frequency of the envelope had a significant impact for all networks ($p < 0.01$). Post-hoc testing revealed that for the DMN, parietal, sensorimotor, fronto-parietal left and right networks no specific envelope frequency dominated. In case of the frontal and visual networks the alpha envelope derived EEG map had the highest D-values (frontal $p < 0.01$, visual $p < 0.01$).

Table 2.1: F-values for bootstrapped EEG-fMRI comparison for the small-band iir-filter. We tested for the frequency specificity within each network.

RSN(EEG-fMRI)	F(df1,df2)	F-value(spatial-corr)	p-value(spatial-corr)	F-value(D-value)	p-value(D-value)
DMN	4,495	3.74	$p < 0.05$	5.03	$p < 0.05$
sensorimotor	4,495	6.95	$p < 0.05$	14.84	$p < 0.05$
visual	4,495	2.43	0.05	69.07	$p < 0.05$
fp_left	4,495	1.53	0.19	6.51	$p < 0.05$
fp_right	4,495	5.09	$p < 0.05$	20.94	$p < 0.05$
pareital	4,495	3.82	$p < 0.05$	37.12	$p < 0.05$
frontal	4,495	5.90	$p < 0.05$	14.70	$p < 0.05$

A



B

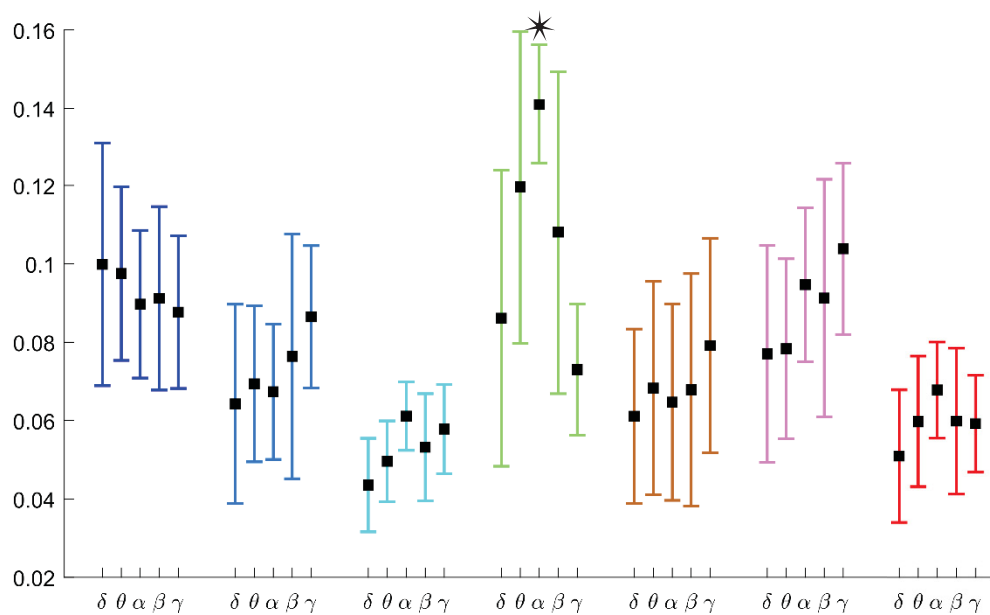


Figure 2.2: fMRI-EEG comparison based on the ICA-envelope approach with small band iir filter. A) shows the results for the spatial correlation based comparison. B) shows the results for D-value based comparison. The black square represents the mean calculated using the bootstrap method and the error bars represent the standard error. The asterisk over any specific error bar represents frequency that had significantly higher mean as compared to the other frequencies within a specific resting state network.

2.3.2 EEG megPAC results

To investigate cross-frequency interaction in RSNs we used the megPAC approach (Florin and Baillet 2015). The first step for the megPAC approach is to identify low frequency phase- high gamma amplitude coupling. Figure 2.3A shows the distribution of low frequencies on the cortex, which had the strongest phase-amplitude coupling. We found coupling with alpha as the low frequency within the ventral portions of the frontal, temporal and the visual cortex. Within previous MEG studies (Florin and Baillet, 2015; Pelzer et al., 2021) the average low frequencies were in the delta/theta range. The spatial distribution of those alpha frequencies as the low-frequency of phase-amplitude coupling points towards a contamination of the EEG from muscle activity – resulting in artificial phase-amplitude coupling. Therefore, we did not perform subsequent steps of the megPAC pipeline to calculate RSNs. Results in the following sections only focus on RSNs derived using the Hilbert envelope approach.




Figure 2.3: A) shows low frequency distribution on the cortex and B) shows high frequency distribution on the cortex.

2.4 Discussion

In our study we evaluated data driven methods for extracting resting state networks from EEG comparable to fMRI based ones. We compared the independent component analysis derived networks from different EEG frequency bands using multiple filter settings. Overall the small band iir filter performed the best for extracting EEG RSNs. In addition, the envelope frequency had a significant main effect on the EEG-fMRI spatial correlation and D-values based comparison for almost all resting state networks. However, post-hoc testing did not

reveal one frequency with a significantly higher mean than the other frequencies for any EEG network.

2.4.1 Optimal filter settings

An important step in finding EEG RSNs is using the correct filter settings in order to account for the $1/f$ characteristics of the EEG signal. By bootstrapping across multiple frequency bands using different types of filters we could discern the effect of filters on the spatial correlation and D-value based comparison between EEG and fMRI RSNs. Overall we found that small band iir filter produced the best spatial correlation based comparison whereas the small band fir filter produced best results for D-value comparison. As described in our methods section, the D-value is a binary measure of comparison between two networks, whereas spatial correlation not only measures the spatial extent of RSN overlap but also accounts for the spatial variability of the RSNs. Hence, an appropriate comparison to select an optimal filter for a time series should be based on spatial correlation.

2.4.2 Frequency specificity of results

Previous studies using envelope based approaches have attributed specific frequencies to different resting state networks (Laufs et al., 2003b; Mantini et al., 2007; Brookes et al., 2011). Contrary to these studies we did not reach such a conclusion in our work. Our results were much more in line with studies which report that resting state connectivity networks are linked across multiple frequencies though some connections might scale in strength in specific frequency bands (Scheeringa et al., 2011; Deligianni et al., 2014; Hipp and Siegel, 2015; Wirsich et al., 2017). Multiplexing using different frequency bands (Akam and Kullmann, 2014) is one plausible mechanism of communication between different networks due to which resting state networks might operate across multiple frequency bands. Another feature seen in EEG recordings is that of nested oscillations and cross frequency coupling (Florin and Baillet 2015, Penny et al., 2008). Another potential reason for resting state networks to operate across multiple frequency bands could be that frequency specific activity might emerge only under task conditions to engage specific cortical computations (Siegel et al., 2012). Methodologically the missing frequency specificity of the RSNs could be due to the fixed point iteration optimisation used in the Fast ICA algorithm. But our bootstrap results address this issue to a certain degree. In our bootstrap approach the stochastic seeding point for the fixed point iteration scheme was re-initialised at each bootstrap iteration. This re-initialisation produces 100 different initialisation conditions for the ICA algorithm. Furthermore, the resultant ICs were compared to the standard fMRI-RSNs at every bootstrap iteration using spatial correlation. In our results the resultant variability is visible for all EEG derived RSNs across all frequency bands. This entire bootstrap pipeline is similar to the ICASSO (Himberg et al. 2004) approach proposed to stabilise ICA outputs. In the ICASSO approach FastICA is run multiple times. The resulting independent components from each ICASSO run are then clustered based on the absolute value of the linear correlation coefficient between components from different repetitions. ICASSO then provides the centrotypic which is the best representative estimate of a cluster of independent components from the different runs

of FastICA. But in our approach given that spatial (linear) correlation between the standard fMRI RSN and different bootstrap outputs was variable, ICASSO stabilisation would not be effective. In the future in order to enforce strict independence in the frequency domain, researchers could also investigate networks found using spectrospatial decomposition using Fourier-ICA methods (Hyvärinen et al., 2010).

2.4.3 Alpha band D-value results for the visual RSN

For the visual RSN the alpha-band had the best correspondence to the fMRI RSN based on D-value (post hoc results). D-value is a binary measure of comparison which estimates the spatial overlap of RSNs. Hence it is possible that D-value is sensitive to detect EEG-fMRI RSN correspondence that is dictated by anatomical connectivity (Honey et al., 2009; Hermundstad et al., 2013). Furthermore, alpha based connectivity seems to be closely related to structural connectivity in the visual cortex (Toosy et al., 2004; Warbrick et al., 2017; Minami et al., 2020). Hence in line with the visual cortical findings we saw that EEG-fMRI comparison (based on D-value) showed a significantly higher overlap in the alpha range.

Interestingly for the same visual RSN no frequency had a better correspondence to the fMRI RSN based on spatial correlation. Spatial correlation takes into account spatial variability in RSNs in addition to spatial overlap (as measured by D-value). Previous studies have indicated a wide variety of relationships between alpha range electrophysiology signals and BOLD fMRI signals. EEG-fMRI research has indicated a negative correlation between alpha activity and fMRI signals in the visual, frontal and temporal regions whereas positive relationship with thalamic areas (Stern, 2002b, Laufs et al., 2003a, b; Moosmann et al., 2003; Feige et al., 2005; DiFrancesco et al., 2008). A study involving concurrent EEG-fMRI data suggested the presence of five different alpha sources under rest (Bridwell et al., 2013). Out of the five sources two had a negative relationship with fMRI activity. Given the complex relationship between alpha electrophysiological activity across the whole brain and fMRI BOLD signals it is possible that spatial correlation based EEG-fMRI comparison would not yield a significant result in the visual or as matter of fact in any RSNs, which is what we see in our results.

2.4.4 High frequency contamination

Muscle related activity manifests as oscillations within the gamma range (40-60 Hz) in EEG signals especially in the temporal, visual as well as frontal regions (Whitham et al., 2007; Muthukumaraswamy, 2013; Nottage and Horder, 2016). This was one of the reasons we decided not to extract EEG networks using the MEG-PAC approach (Florin and Baillet, 2015) because the PAC of the EEG data was shifted towards noisy high frequency coupling (see figure 2.3B).

2.5 Conclusion

We derived whole brain, non-atlas constrained and data driven EEG RSNs. Furthermore, we investigated the relationship between EEG RSNs and fMRI RSNs and found that the EEG RSNs had reasonable spatial profiles when compared to fMRI based RSNs, though they were less

spatially selective than their fMRI counterparts. We investigated the relationship between EEG-RSNs derived in different frequency bands and their fMRI counterparts and found no frequency specificity for EEG networks. One limitation of EEG based analysis is that of muscular artefacts contaminating high frequency data (gamma and higher) due to which one might be limited to the available methods to extract networks and thus we could not extract RSNs based on the megPAC approach. Given our results it seems plausible to extract resting state networks from EEG data using the small band iir filter coupled with Hilbert envelope based ICA. In essence it could bring down the cost of producing such networks and bring a more nuanced understanding of resting networks in clinical settings.

3 Study 2 Differential dopaminergic modulation of spontaneous cortico–subthalamic activity in Parkinson’s disease

Note: This study has been published in a peer reviewed journal (<https://elifesciences.org/articles/66057>)

Contributions

Abhinav Sharma: Conceptualization, Partial data acquisition and quality control, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (65% contribution)

Dr. Diego Vidaurre: Software, Validation, Methodology, Writing - review and editing

Dr. Jan Vesper: Resources, Investigation

Dr. Alfons Schnitzler: Resources, Investigation, Writing - review and editing

Dr. Esther Florin: Conceptualization, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing

Abstract

Pathological oscillations including elevated beta activity in the subthalamic nucleus (STN) and between STN and cortical areas are a hallmark of neural activity in Parkinson’s disease (PD). Oscillations also play an important role in normal physiological processes and serve distinct functional roles at different points in time. We characterised the effect of dopaminergic medication on oscillatory whole-brain networks in PD in a time-resolved manner by employing a hidden Markov model on combined STN local FPs and magnetoencephalography (MEG) recordings from 17 PD patients. Dopaminergic medication led to coherence within the medial and orbitofrontal cortex in the delta/theta frequency range. This is in line with known side effects of dopamine treatment such as deteriorated executive functions in PD. In addition, dopamine caused the beta band activity to switch from an STN-mediated motor network to a frontoparietal-mediated one. In contrast, dopamine did not modify local STN–STN coherence in PD. STN–STN synchrony emerged both on and off medication. By providing electrophysiological evidence for the differential effects of dopaminergic medication on the discovered networks, our findings open further avenues for electrical and pharmacological interventions in PD.

3.1 Introduction

Oscillatory activity serves crucial cognitive roles in the brain (Akam and Kullmann, 2010; Akam and Kullmann, 2014), and alterations of oscillatory activity have been linked to neurological and psychiatric diseases (Schnitzler and Gross, 2005). Different large-scale brain networks operate with their own oscillatory fingerprint and carry out specific functions (Keitel and Gross, 2016; Mellem et al., 2017; Vidaurre et al., 2018b). Given the dynamics of cognition, different brain networks need to be recruited and deployed flexibly. Hence, the duration for which a network is active, its overall temporal presence, and even the interval between the different activations of a specific network might provide a unique window to understanding brain functions. Crucially, alterations of these temporal properties or networks might be related to neurological disorders.

In Parkinson's disease (PD), beta oscillations within the subthalamic nucleus (STN) and motor cortex (13–30 Hz) correlate with the motor symptoms of PD (Marreiros et al., 2013; van Wijk et al., 2016; West et al., 2018). Beta oscillations also play a critical role in communication in a healthy brain (Engel and Fries, 2010). (For the purposes of our paper, we refer to oscillatory activity or oscillations as recurrent but transient frequency-specific patterns of network activity, even though the underlying patterns can be composed of either sustained rhythmic activity, neural bursting, or both [Quinn et al., 2019]. Disambiguating the exact nature of these patterns is, however, beyond the scope of this work.) At the cellular level, loss of nigral dopamine neurons in PD leads to widespread changes in brain networks, to varying degrees across different patients. Dopamine loss is managed in patients via dopaminergic medication. Dopamine is a widespread neuromodulator in the brain (Gershman and Uchida, 2019), raising the question of whether each medication-induced change restores physiological oscillatory networks. In particular, dopaminergic medication is known to produce cognitive side effects in PD patients (Voon et al., 2009). According to the dopamine overdose hypothesis, a reason for these effects is the presence of excess dopamine in brain regions not affected in PD (MacDonald et al., 2011; MacDonald and Monchi, 2011). Previous task-based and neuroimaging studies in PD demonstrated frontal cognitive impairment due to dopaminergic medication (Cools et al., 2002; Ray and Strafella, 2010; MacDonald et al., 2011).

Using resting-state whole-brain MEG analysis, network changes related to both motor and non-motor symptoms of PD have been described (Olde Dubbelink et al., 2013a; Olde Dubbelink et al., 2013b). However, these studies could not account for simultaneous STN–STN or cortico–STN activity affecting these networks, which would require combined MEG/electroencephalogram (EEG)–LFP recordings (Litvak et al., 2021). Such recordings are possible during the implantation of deep brain stimulation (DBS) electrodes, an accepted treatment in the later stages of PD (Volkman et al., 2004; Deuschl et al., 2006; Kleiner-Fisman et al., 2006). Combined MEG–LFP studies in PD involving dopaminergic intervention report changes in beta and alpha band connectivity between specific cortical regions and the STN (Litvak et al., 2011; Hirschmann et al., 2013; Oswal et al., 2016). Decreased cortico–STN coherence under dopaminergic medication (ON) correlates with improved motor functions in

PD (George et al., 2013). STN–STN intra-hemispheric oscillations positively correlate to motor symptom severity in PD without dopaminergic medication (OFF), whereas dopamine-dependent nonlinear phase relationships exist between inter-hemispheric STN–STN activity (West et al., 2016). Crucially, previous studies could not rule out the influence of cortico–STN connectivity on these inter-hemispheric STN–STN interactions.

To further characterise the differential effects of dopaminergic medication and delineate pathological versus physiological-relevant spectral connectivity in PD, we study PD brain activity via a hidden Markov model (HMM), a data-driven learning algorithm (Vidaurre et al., 2016; Vidaurre et al., 2018b). Due to the importance of cortico–subcortical interactions in PD, we investigated these interactions with combined spontaneous whole-brain magnetoencephalography (MEG) and STN local FPs (LFPs) recordings from PD patients. We study whole-brain connectivity including the STN using spectral coherence as a proxy for communication based on the communication through coherence hypothesis (Fries, 2005; Fries, 2015). This will allow us to delineate differences in communication OFF and ON medication. Furthermore, we extended previous work that was limited to investigating communication between specific pairs of brain areas (Litvak et al., 2011; George et al., 2013; Hirschmann et al., 2013). Moreover, we identified the temporal properties of the networks both ON and OFF medication. The temporal properties provide an encompassing view of network alterations in PD and the effect of dopamine on these networks.

We found that cortico–cortical, cortico–STN, and STN–STN networks were differentially modulated by dopaminergic medication. For the cortico–cortical network, medication led to additional connections that can be linked to the side effects of dopamine. At the same time, dopamine changed the cortico–STN network towards a pattern more closely resembling physiological connectivity as reported in the PD literature. Within the third network, dopamine only had an influence on local STN–STN coherence. These results provide novel information on the oscillatory network connectivity occurring in PD and the differential changes caused by dopaminergic intervention. These whole-brain networks, along with their electrophysiological signatures, open up new potential targets for both electric and pharmacological interventions in PD.

3.2 Methods

3.2.1 Subjects

In total, 17 (4 female) right-handed PD patients (age: 55.2 ± 9.3 years) undergoing surgery for therapeutic STN DBS were recruited for this study. Patients had been selected for DBS treatment according to the guidelines of the German Society for Neurology. The experimental procedure was explained to all participants and they gave written consent. The study was approved by the local ethics committee (study number 5608R) and conducted in accordance with the Declaration of Helsinki. Bilateral DBS electrodes were implanted in the dorsal part of the STN at the Department of Functional Neurosurgery and Stereotaxy in Düsseldorf. The implanted DBS electrodes used were the St. Jude Medical directional lead 6172 (Abbott

Laboratories, Lake Bluff, IL) and in one case the Boston Scientific Vercise segmented lead (Boston Scientific Corporation, Marlborough, MA). These electrodes have four contact heights and the two middle heights are segmented into three equally spaced contacts.

The DBS leads were externalised and we measured the patients after 1–3 days. To simultaneously acquire MEG and LFP signals, we connected the externalised leads to an EEG amplifier integrated with the MEG system. We used a whole-head MEG system with 306 channels (Elekta Vectorview, Elekta Neuromag, Finland) housed within a magnetically shielded chamber. All patients were requested to sit still and awake during data acquisition. To ensure that patients did not fall asleep, we tracked patients' pupil diameter with an eye tracker. To remove eye blink and cardiac artefacts, electrooculography and electrocardiography were recorded along with the LFP and MEG signals. In order to co-register the MEG recording with the individual MRI, four head position indicator coils were placed on the patient's head. Their position as well as additional head points were digitised using the Polhemus Isotrack system (Polhemus, Colchester, CT). The data were recorded with a sampling rate of 2400 Hz and a low-pass filter of 800 Hz was applied. An electrode was placed at the mastoid and all LFP signals were referenced to it.

For the clinical OFF medication state, oral PD medication was withdrawn overnight for at least 12 hr. If a patient had an apomorphine pump, this pump was stopped at least 1 hr before the measurement. First, we recorded resting-state activity in the medication OFF condition. The patients were then given their morning dose of L-DOPA in the form of fast-acting levodopa. Data were acquired in three runs of 10 min, for a total of 30 min for each medication condition. We started the ON medication measurement at least half an hour after the administration of the dose and after clinical improvement was seen. The same procedure as for the OFF medication state was followed for the ON medication measurement.

3.2.2 Pre-processing

All data processing and analyses were performed using Matlab (version R 2016b; Math Works, Natick, MA). Custom-written Matlab scripts (https://github.com/saltwater-tensor/HMM_pipeline) and the Brainstorm toolbox (<http://neuroimage.usc.edu/brainstorm/Introduction>) were used (Tadel et al., 2011). To ensure artefact-free data, two people independently inspected the data visually, cleaned artefacts, and compared the cleaning output. The final cleaned data included changes agreed upon by both the people involved in cleaning. The Neuromag system provides signal-space projection (SSP) vectors for the cleaning of external artefacts from the MEG channels, which were applied. The line noise was removed from all channels with a notch filter at 50, 100, 150, ..., 550, and 600 Hz with a 3 dB bandwidth of 1 Hz. The LFP recordings from the DBS electrode were re-referenced against the mean of all LFP channels. Very noisy and flat MEG/LFP channels were excluded from further analysis. Time segments containing artefacts were removed from the time series. However, if artefacts regularly occurred only in one single channel, this whole channel was removed instead. Frequently arising artefacts following the same basic pattern, such as eye blinks or cardiac artefacts, were removed via SSP. All data

were high-pass filtered with 1 Hz to remove movement-related low-frequency artefacts. Finally, the data were down-sampled to 1000 Hz.

Source estimation was performed on these recordings at an individual level using each individual's anatomy. Therefore, using Freesurfer (<https://surfer.nmr.mgh.harvard.edu/>, v.5.3.0), the individual cortical surfaces were extracted from the individual T1-weighted MRI scans (3T scanner and 1 mm³ voxel size). We used the overlapping spheres method with 306 spheres for the forward model. As the inverse model, we used a linearly constrained minimum variance (LCMV) beamformer. The data covariance matrix for the LCMV beamformer was computed directly from each 10 min recording. The data covariance was regularised using the median eigenvalue of the data covariance matrix. The noise covariance was obtained from an empty room recording on the same day as the actual measurement.

For each subject, the invasive entry point of the STN was identified based on intraoperative microelectrode recordings (Gross et al., 2006; Moran et al., 2006). Subsequently, the first recording height after the entry into the STN was selected to obtain the three directional LFP recordings from the respective hemisphere. In addition, we visualised the location of all electrodes using lead-DBS (Horn et al., 2019). All electrodes were properly placed within the STN – except for one (see figure 3.1 below). To exclude that our results were driven by outlier, we reanalysed our data without this patient. No qualitative change in the overall connectivity pattern was observed.




Figure 3.1: Deep brain stimulation (DBS) electrode location for all subjects. Lead-DBS reconstruction with all subjects. The red leads are the ones of a subject with one of the outside the STN. The red directional contacts are the ones from which the data was used for analysis.

The source-reconstructed MEG data were projected to the default cortical anatomy (MNI 152 with 15,002 vertices) and then down-sampled temporally to 250 Hz for each medication condition for every subject. We used the Mindboggle atlas to spatially reduce the data dimensions. For each of the 42 cortical regions in the atlas, a multidimensional time series consisting of the vertices within that anatomical region was extracted. To reduce the multivariate times series for each region to a single one, we employed the first principal

component explaining the highest variance share in each region. The first principal component row vectors from all 42 anatomical regions were stacked into a MEG cortical time series matrix. To correct for volume conduction in the signal, symmetric orthogonalisation (Colclough et al., 2015) was applied to each subject's resulting MEG cortical time series matrix. The row vectors of this orthogonalised matrix and the six LFPs (three each for left and right STN) were z-scored. Subsequently, they were stacked into one multidimensional time series (N by T) matrix. Here, N = 48 is the total number of nodes/regions (42 regions from the cortex and 6 LFP electrode contacts) and T denotes the length of the time dimension. This 48 by T data matrix obtained from each subject was concatenated along the temporal dimension across all subjects for each specific medication condition. Finally, to resolve sign ambiguity inherent in source-reconstructed MEG data as well as resolve polarity of LFP channels across subjects, a sign-flip correction (Vidaurre et al., 2016) procedure was applied to this final 48 by T (T by number of subjects) dataset within a medication condition. The pre-processing steps were performed for OFF and ON medication separately.

3.2.3 HMM analysis

The HMM is a data-driven probabilistic algorithm which finds recurrent network patterns in multivariate time series (Vidaurre et al., 2016; Vidaurre et al., 2018a). Each network pattern is referred to as a 'state' in the HMM framework, such that these networks can activate or deactivate at various points in time. Here onwards, 'state' or 'network' is used interchangeably. We used a specific variety of the HMM, the TDE-HMM, where whole-brain networks are defined in terms of both spectral power and phase coupling (Vidaurre et al., 2018b). Hence, for every time point, the HMM algorithm provided the probability that a network is active. Here onwards, a contiguous block of time for which the probability of a particular network being active remained higher than all the other networks is referred to as a 'state visit'. Hence, the HMM produced temporally resolved spatial networks for the underlying time series. In our approach, we also performed spectral analyses of these state visits, leading to a complete spatio-spectral connectivity profile across the cortex and the STN. By applying the HMM analysis to the combined MEG–LFP dataset, we were able to temporally, spatially, and spectrally separate cortico–cortical, cortico–STN, and STN–STN networks.

3.2.4 Estimation of the HMM

Since we were interested in recovering phase-related networks, the TDE-HMM was fit directly on the time series obtained after pre-processing steps described previously, as opposed to its power envelope. This preserved the cross-covariance within and across the underlying raw time series of the cortical regions and the STN. The model estimation finds recurrent patterns of covariance between regions (42 cortical regions and 6 STN contacts) and segregates them into 'states' or 'networks'. Based on these covariance patterns, for each state, the power spectra of each cortical region and the coherence amongst regions can be extracted.

We opted for six different states as a reasonable trade-off between the spectral quality of the results and their redundancy. The HMM-MAR toolbox (Vidaurre et al., 2016) was used for fitting the TDE-HMM. We employed the TDE version of the HMM where the embedding took place in a 60 ms window (i.e., a 15 time-point window for a sampling frequency of 250 Hz). Since time embedding would increase the number of rows of the data from 48 to 48 times the window length (also referred to as number of lags), an additional PCA (principal component analysis) (reduction across 48 by number of lags) step was performed after time embedding. The number of components retained was 96 (48×2). This approach follows Vidaurre et al., 2018b. To characterise each state, a full covariance matrix with an inverse Wishart prior was used. The diagonal of the prior for the transition probability matrix was set as 10. To ensure that the mean of the time series did not take part in driving the states, the ‘zero mean’ option in HMM toolbox was set to 1. To speed up the process of fitting, we used the stochastic version of variational inference for the HMM. In order to start the optimisation process, the ‘HMM-MAR’-type initialisation was used (for details, see Vidaurre et al., 2016). The HMM was fit separately OFF and ON medication.

3.2.5 Statistical analysis of the states

After the six states were obtained for HMM OFF and HMM ON medication, these states were statistically compared within each medication condition as well as between medication conditions. In addition, the temporal properties of these states were compared.

3.2.6 Intra-medication analysis

We investigated the spectral connectivity patterns across the different states within a medication condition (intra-medication or IntraMed). The objective was to uncover significant coherent connectivity standing out from the background within each frequency band (delta/theta [1–8 Hz], alpha [8–12 Hz], and beta [13–30 Hz]) in the respective states. The HMM output included the state time courses (i.e., when the states activated) for the entire concatenated data time series. The state time courses allowed the extraction of state- and subject-specific data for further state- and subject-level analysis. For each HMM state, we filtered the state-specific data for all the subjects between 1 and 45 Hz. (For state-wise data extraction, please refer the HMM toolbox wiki [<https://github.com/OHBA-analysis/HMM-MAR/wiki/User-Guide>].) Then we calculated the Fourier transform of the data using a multitaper approach to extract the frequency components from the short segments of each state visit. (See Vidaurre et al., 2018b for discussion on multitaper for short time data segments.) Seven Slepian tapers with a time–bandwidth product of 4 were used, resulting in a frequency resolution of 0.5 Hz and therefore binned frequency domain values. Subsequently, we calculated the coherence and power spectral density of this binned (frequency bins obtained during the multitaper step) data for every subject and every state. The coherence and the power spectral density obtained were three-dimensional matrices of size f (number of frequency bins) by N (42 cortical locations + 6 STN contacts) by N .

Based on the coherence matrices, we performed a frequency band-specific analysis. Canonical definitions of frequency bands assign equal weight to each frequency bin within a band for every subject. This might not be suitable when considering analyses of brain signals across a large dataset. For example, the beta peak varies between individual subjects. Assigning the same weight to each bin in the beta range might reduce the beta effect at the group level. To allow for inter-subject variability in each frequency bin's contribution to a frequency band, we determined the frequency modes in a data-driven manner (Vidaurre et al., 2018b). Because we focused on interactions that are important to establish the STN–cortex communication, the identification of the relevant frequency modes was restricted to the cross-coherence between the STN–LFPs and cortical signals; in other words, the block matrix consisting of rows 1–6 (STN) and columns 7–48 (cortex). For each subject, this extracted submatrix was then vectorised across columns. This gave us a (number of frequency bins by 252 [6 STN locations by 42 cortical locations]) matrix for each state. For every subject, this matrix was concatenated along the spatial dimension across all states producing a (number of frequency bins by [252 by 6 (number of states)]) matrix. We called this the subject-level coherence matrix. We averaged these matrices across all subjects along the spectral dimension (number of frequency bins) to yield a (number of frequency bins by [252 by 6]) group-level coherence matrix. We factorised the group-level coherence matrix into four frequency modes using a non-negative matrix factorisation (NNMF) (Lee and Seung, 2001). Each of the resulting four frequency modes obtained was of size (one by number of frequency bins). The values of frequency modes are the actual NNMF weights obtained from the NNMF estimation (which, just like a regression coefficient, are unit-less, because coherence is unit-less). Three of them resembled the canonical delta/theta (delta and theta frequencies were combined into one band), alpha, and beta bands whereas the last one represented noise. Since NNMF does not guarantee a unique solution, we performed multiple instances of the factorisation. In practice we could obtain frequency modes, which showed correspondence to the classical frequency bands, within four iterations of the algorithm. At each instance, we visualised the output to ensure frequency specificity of the frequency modes. The stability of the output was ensured by using 'robust NNMF', which is a variant of the NNMF algorithm (Vidaurre et al., 2018b). While these frequency modes were derived in fact from coherence measures (as detailed in Vidaurre et al., 2018a), they can be applied to power measures or any other frequency-specific measure. We then computed the inner product between the subject- and group-level coherence matrix and the frequency modes obtained above. We called these the subject-level and group-level projection results, respectively.

To separate background noise from the strongest coherent connections, a Gaussian mixture model (GMM) approach was used (Vidaurre et al., 2018b). For the group-level projection results, we normalised the activity in each state for each spectral band by subtracting the mean coherence within each frequency mode across all states. As a prior for the mixture model, we used two single-dimensional Gaussian distributions with unit variance: one mixture component to capture noise and the other to capture the significant connections. This GMM with two mixtures was applied to the coherence values (absolute value) of each

state. Connections were considered significant if their p-value after correction for multiple comparisons was smaller than 0.05.

3.2.7 Inter-medication analysis

To test for differences in coherence across medication conditions (inter-medication or InterMed), the first step was to objectively establish a comparison between the states found in the two HMMs fit separately for each condition. There is no a priori reason for the states detected in each condition to resemble each other. To find OFF and ON medication states that may resemble each other, we calculated the Riemannian distance (Förstner and Moonen, 2003) between the state covariance matrices of the OFF and ON HMM. This yielded an OFF states by ON states (6×6) distance matrix. Subsequently, finding the appropriately matched OFF and ON states reduced to a standard linear assignment problem. We found an ON state counterpart to each OFF state by minimising the total sum of distances using the Munkres linear assignment algorithm (Vidaurre et al., 2018a). This approach yielded a one-to-one pairing of OFF and ON medication states, and all further analysis was conducted on these pairs. For ease of reading, we gave each pair its own label. For example, when we refer to a 'Ctx–STN' state in the following sections, then such a state was discovered OFF medication and its corresponding state ON medication is its distance-matched partner. In subsequent sections, all mentions of ON or OFF medication refer to these state pairs unless mentioned otherwise.

We used the subject-level projection results obtained during IntraMed analysis to perform InterMed analyses. We performed two-sided independent sample t-tests between the matched states to compare the coherence, which was calculated between different regions of interest (see Dataset preparation). We grouped individual atlas regions into canonical cortical regions like frontal, sensorimotor, parietal, visual, medial PFC (prefrontal cortex), and STN contacts. For example, in the beta band, STN (contacts)–sensorimotor coherence in the OFF condition was compared to the STN (contacts)–sensorimotor coherence in the ON condition. The p-values obtained were corrected for multiple comparisons for a total number of possible combinations.

3.2.8 Temporal properties of HMM states

To test for changes in the temporal properties OFF versus ON medication, we compared the lifetimes, interval between visits, and FO for each state both within and across HMMs using two-way repeated measures ANOVA followed by post hoc tests. Lifetime/dwell time of a state refers to the time spent by the neural activity in that state. Interval of visit was defined as the time between successive visits of the same state. Finally, the FO of a state was defined as the fraction of time spent in each state. Extremely short state visits might not reflect neural processes, hence we only used values that were greater than 100 ms for lifetime comparisons.

3.3 Results

Under resting-state conditions in PD patients, we simultaneously recorded whole-brain MEG activity with LFPs from the STN using directional electrodes implanted for DBS. Using an HMM, we identified recurrent patterns of transient network connectivity between the cortex and the STN, which we henceforth refer to as an ‘HMM state’. In comparison to classic sliding window analysis, an HMM solution can be thought of as a data-driven estimation of time windows of variable length (within which a particular HMM state was active): once we know the time windows when a particular state is active, we compute coherence between different pairs of regions for each of these recurrent states. Each HMM state itself is a multidimensional, time-delay embedded (TDE) covariance matrix across the whole brain, containing information about cross-regional coherence and power in the frequency domain. Additionally, the temporal evolution of the HMM states was determined. The PD data were acquired under medication (L-DOPA) OFF and ON conditions, which allowed us to delineate the physiological versus pathological spatio-spectral and temporal changes observed in PD. To allow the system to dynamically evolve, we use time delay embedding. Theoretically, delay embedding can reveal the state space of the underlying dynamical system (Packard et al., 1980). Thus, by delay-embedding PD time series OFF and ON medication, we uncover the differential effects of a neurotransmitter such as dopamine on underlying whole-brain connectivity. OFF medication, patients had on average a Unified Parkinson’s Disease Rating Scale (UPDRS) part III of 29.24 ± 10.74 . This was reduced by L-DOPA (176.5 ± 56.2 mg) to 19.47 ± 8.52 , indicating an improvement in motor symptoms.

3.3.1 Spontaneous brain activity in PD can be resolved into distinct states

Using an HMM, we delineated cortico–subthalamic spectral changes from both global source-level cortical interactions as well as local STN–STN interactions. Three of the six HMM states could be attributed to physiologically interpretable connectivity patterns. We could not interpret the other three states within the current physiological frameworks both OFF and ON medication and they are therefore not considered in the following (see Figure 3.3—supplementary figure 3.1). The connectivity between different brain regions for each state was visualised for the frequency modes shown in Figure 3.2. Figures 3.3–3.5 show the connectivity patterns for the three physiologically meaningful states in both the OFF (top row) and ON medication condition (bottom row). We refer to the state obtained in Figure 3.3 as the cortico–cortical state (Ctx–Ctx). This state was characterised mostly by local coherence within segregated networks OFF medication in the alpha and beta band. In contrast, there was a widespread increase in coherence across the brain from OFF to ON medication. Therefore, ON medication, the connectivity strength in the alpha and beta band was not significantly different from the mean noise level. Figure 3.4 displays the second state. A large proportion of spectral connections in this state enable cortico–STN communication via spectral coherence (Lalo et al., 2008; Litvak et al., 2011; Hirschmann et al., 2013; Oswal et al., 2013; van Wijk et al., 2016) and thus we labelled this as the cortico–STN state (Ctx–STN). This

state was characterised by connectivity between multiple cortical regions and the STN OFF medication, but increased specificity of cortical–STN connectivity ON medication. Finally, Figure 3.5 shows the third state. Within this state, highly synchronous STN–STN spectral connectivity emerged, both OFF and ON medication and therefore we named it the STN–STN state (STN–STN). The spectral characteristics of this state largely remain unaffected under the influence of dopaminergic medication. In the following sections, we describe these three states in detail.

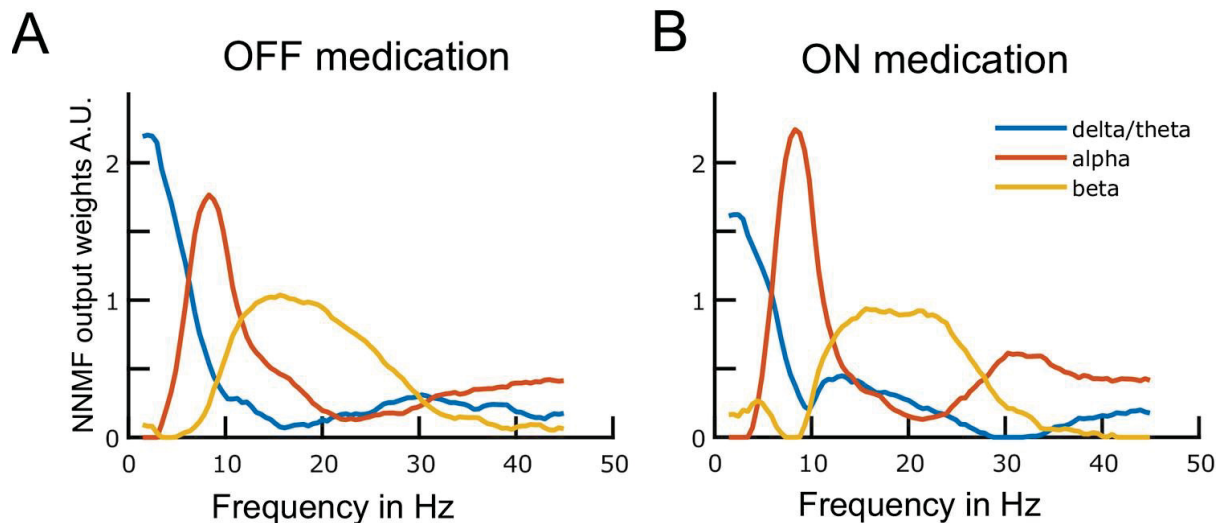


Figure 3.2: Data-driven frequency modes. Each plotted curve shows a different spectral band. The x-axis represents frequency in Hz and the y-axis represents the weights obtained from the non-negative matrix factorisation (NNMF) in arbitrary units. The NNMF weights are like regression coefficients. The frequency resolution of the modes is 0.5 Hz. Panels A and B show the OFF and ON medication frequency modes, respectively.

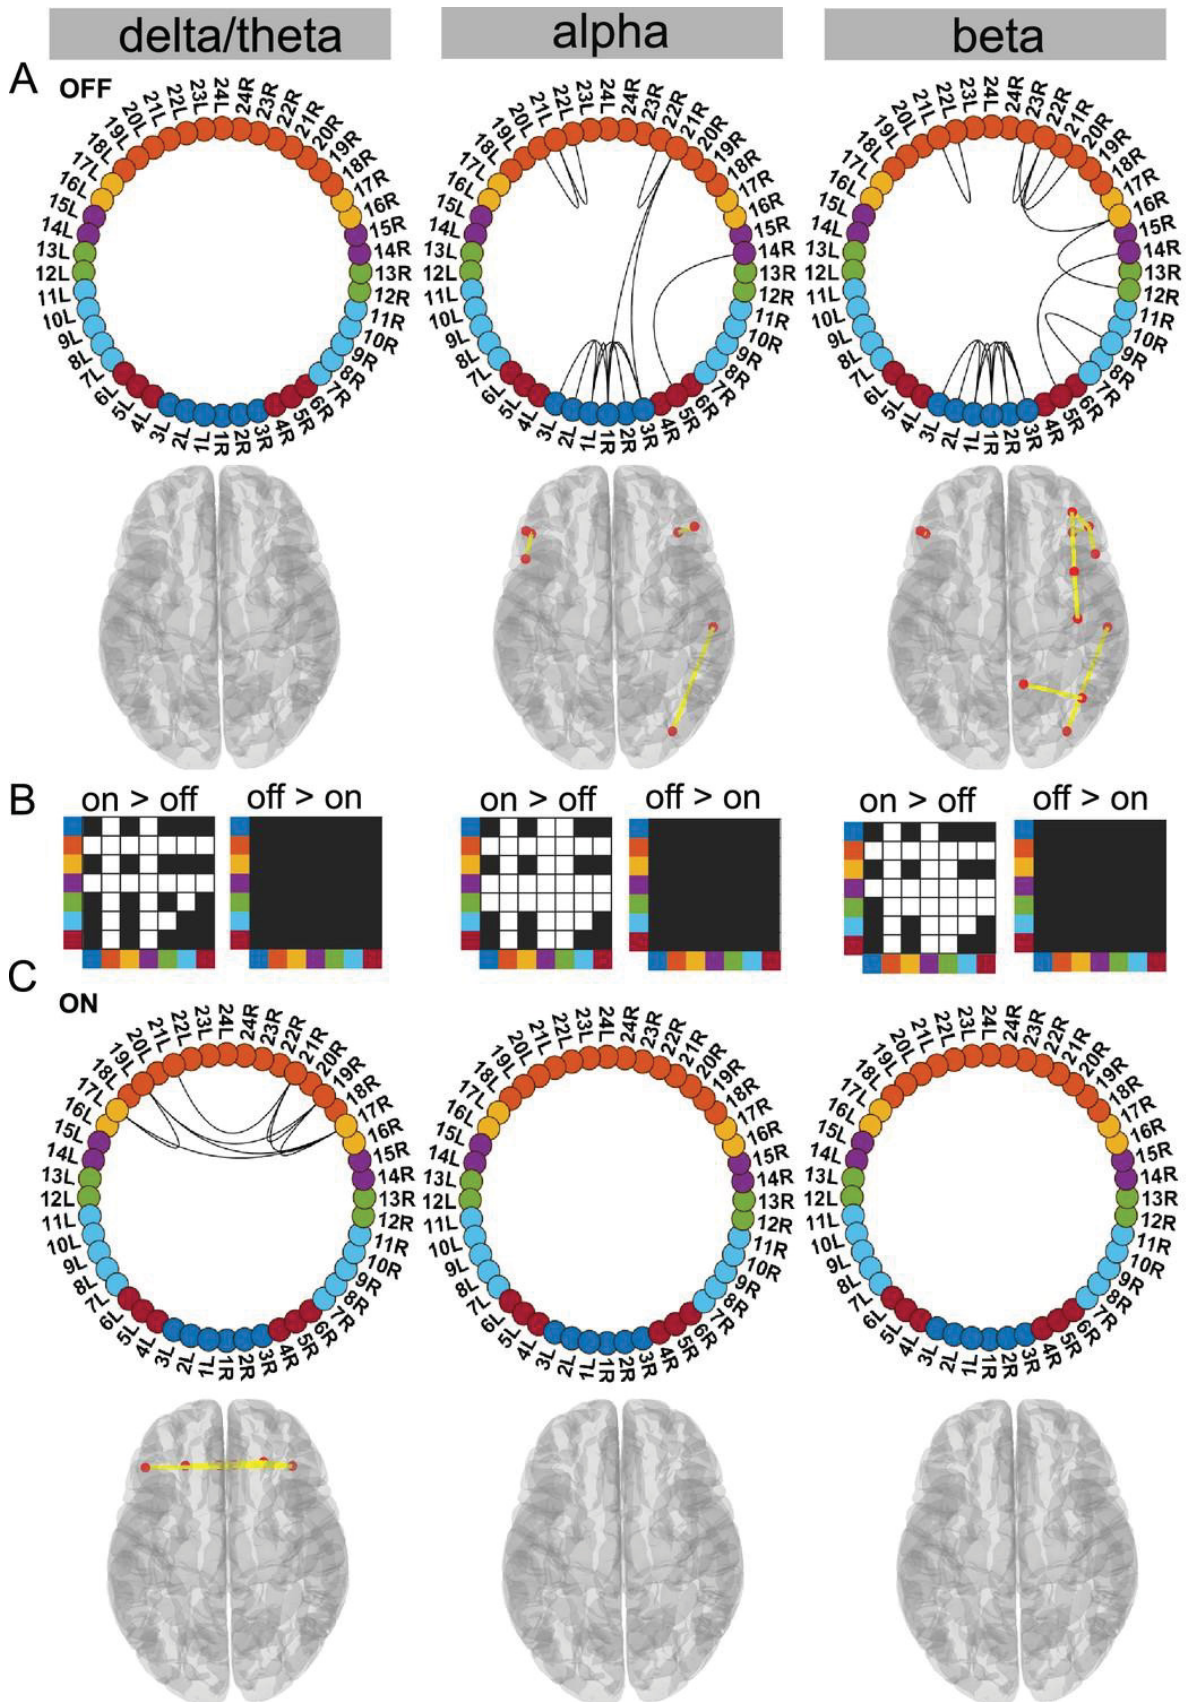


Figure 3.3: Cortico-Cortical State. The cortico–cortical state was characterised by a significant increase in coherence ON compared to OFF medication (see panel B). Due to this, no connections within the alpha and beta band ON medication were significantly higher than the mean (panel C). However, in the delta band, ON medication medial prefrontal–orbitofrontal connectivity emerged. (A and C) Each node in the circular graph represents a brain region based on the Mindboggle atlas. The regions from the atlas are listed in [Table 1](#) along with their corresponding numbers that are used in the circular graph. The colour code in the circular graph represents a group of regions clustered according to the atlas (starting from node number 1) STN contacts (contacts 1, 2, 3 = right STN and contacts 4, 5, 6 = left STN), frontal, medial frontal, temporal, sensorimotor, parietal, and visual cortices. In the circular graph, only the significant connections ($p < 0.05$; corrected for multiple comparisons, IntraMed analysis) are displayed as black curves connecting the nodes. The circles from left to right represent the delta/theta, alpha, and beta bands. Panel A shows results for OFF medication data and panel C for the ON medication condition. For every circular graph, we also show a corresponding top view of the brain with the connectivity represented by yellow lines and the red dot represents the anatomical seed vertex of the brain region. Only the cortical connections are shown. Panel B shows the result for inter-medication analysis (InterMed) for the cortico–cortical state. In each symmetric matrix, every row and column corresponds to a specific atlas cluster denoted by the dot colour on the side of the matrix. Each matrix entry is the result of the InterMed analysis where OFF medication connectivity between i th row and j th column was compared to the ON medication connectivity between the same connections. A cell is white if the comparison mentioned on top of the matrix (either ON >OFF or OFF >ON) was significant at a threshold of $p < 0.05$.

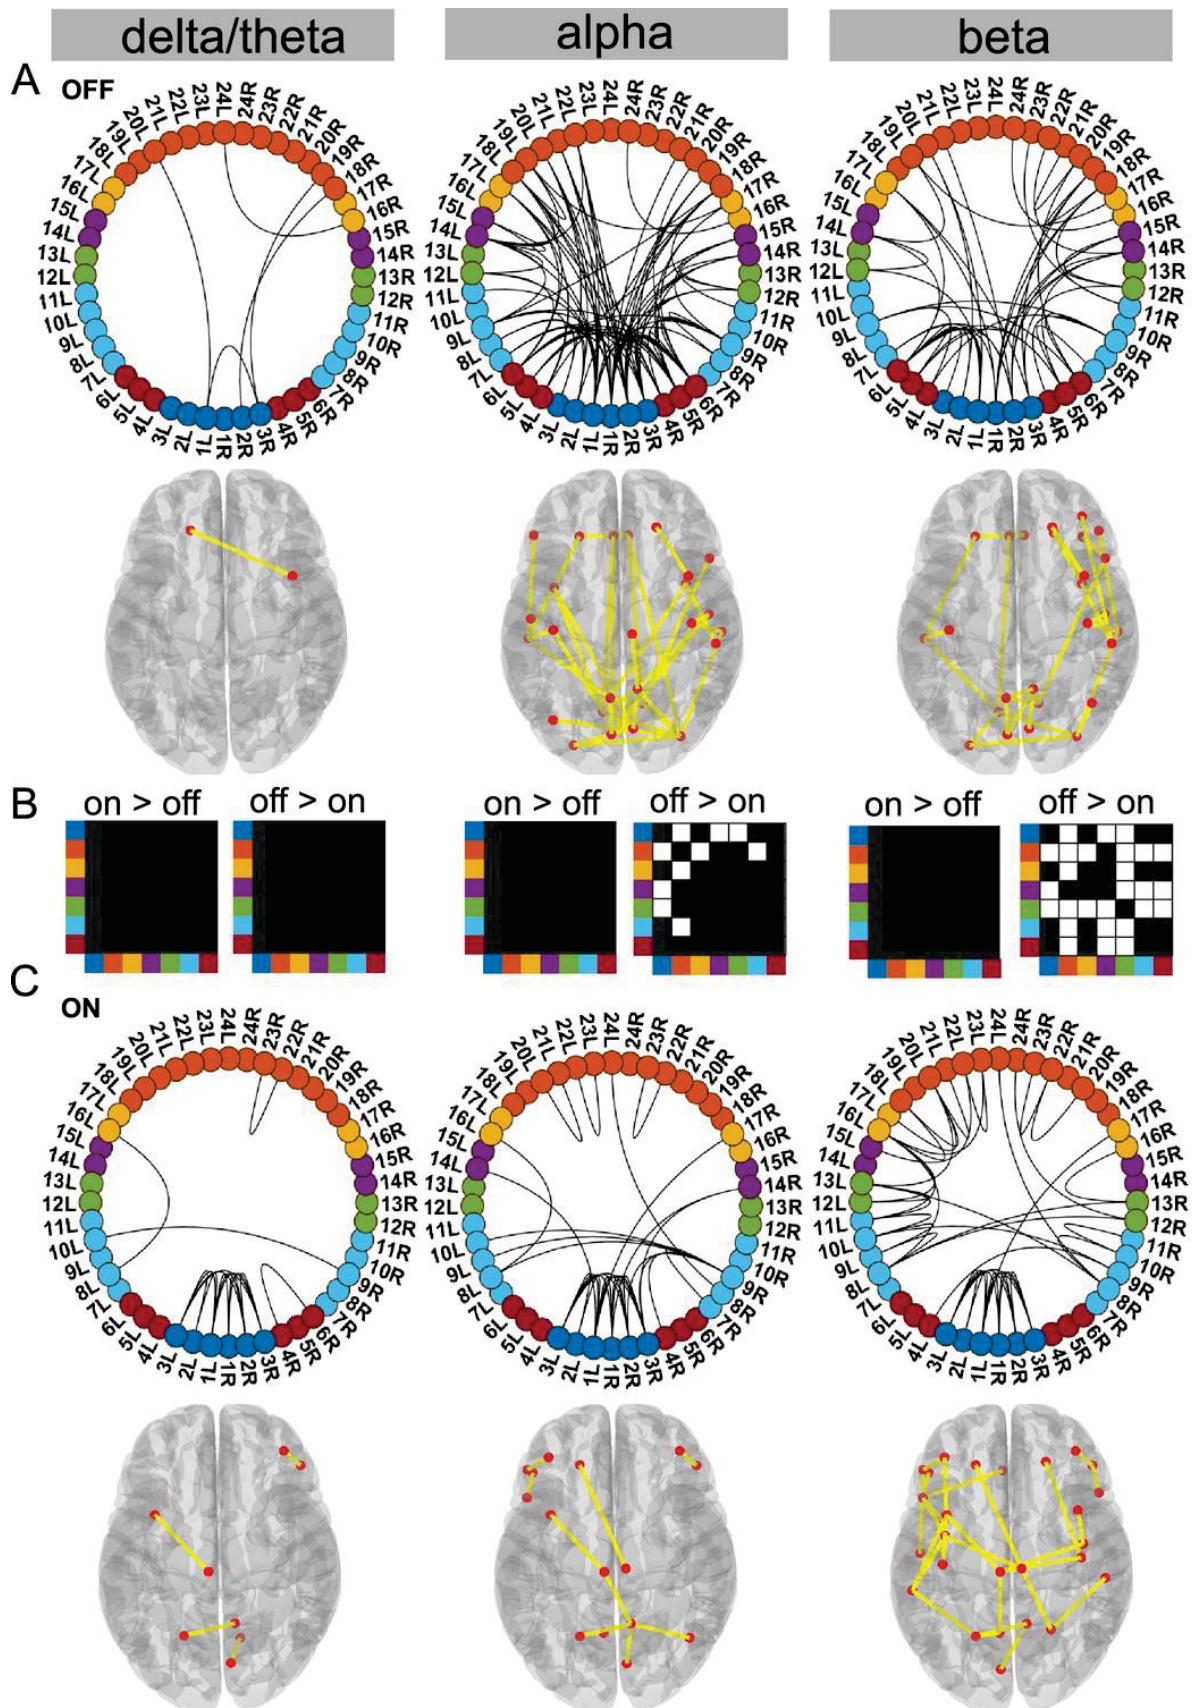


Figure 3.4: Cortico–STN state. For the general description, see the note to figure 3.2. The cortico–STN state was characterised by preservation of spectrally selective cortico–STN connectivity ON medication. Also, ON medication, a sensorimotor–frontoparietal network emerged. $p < 0.05$.

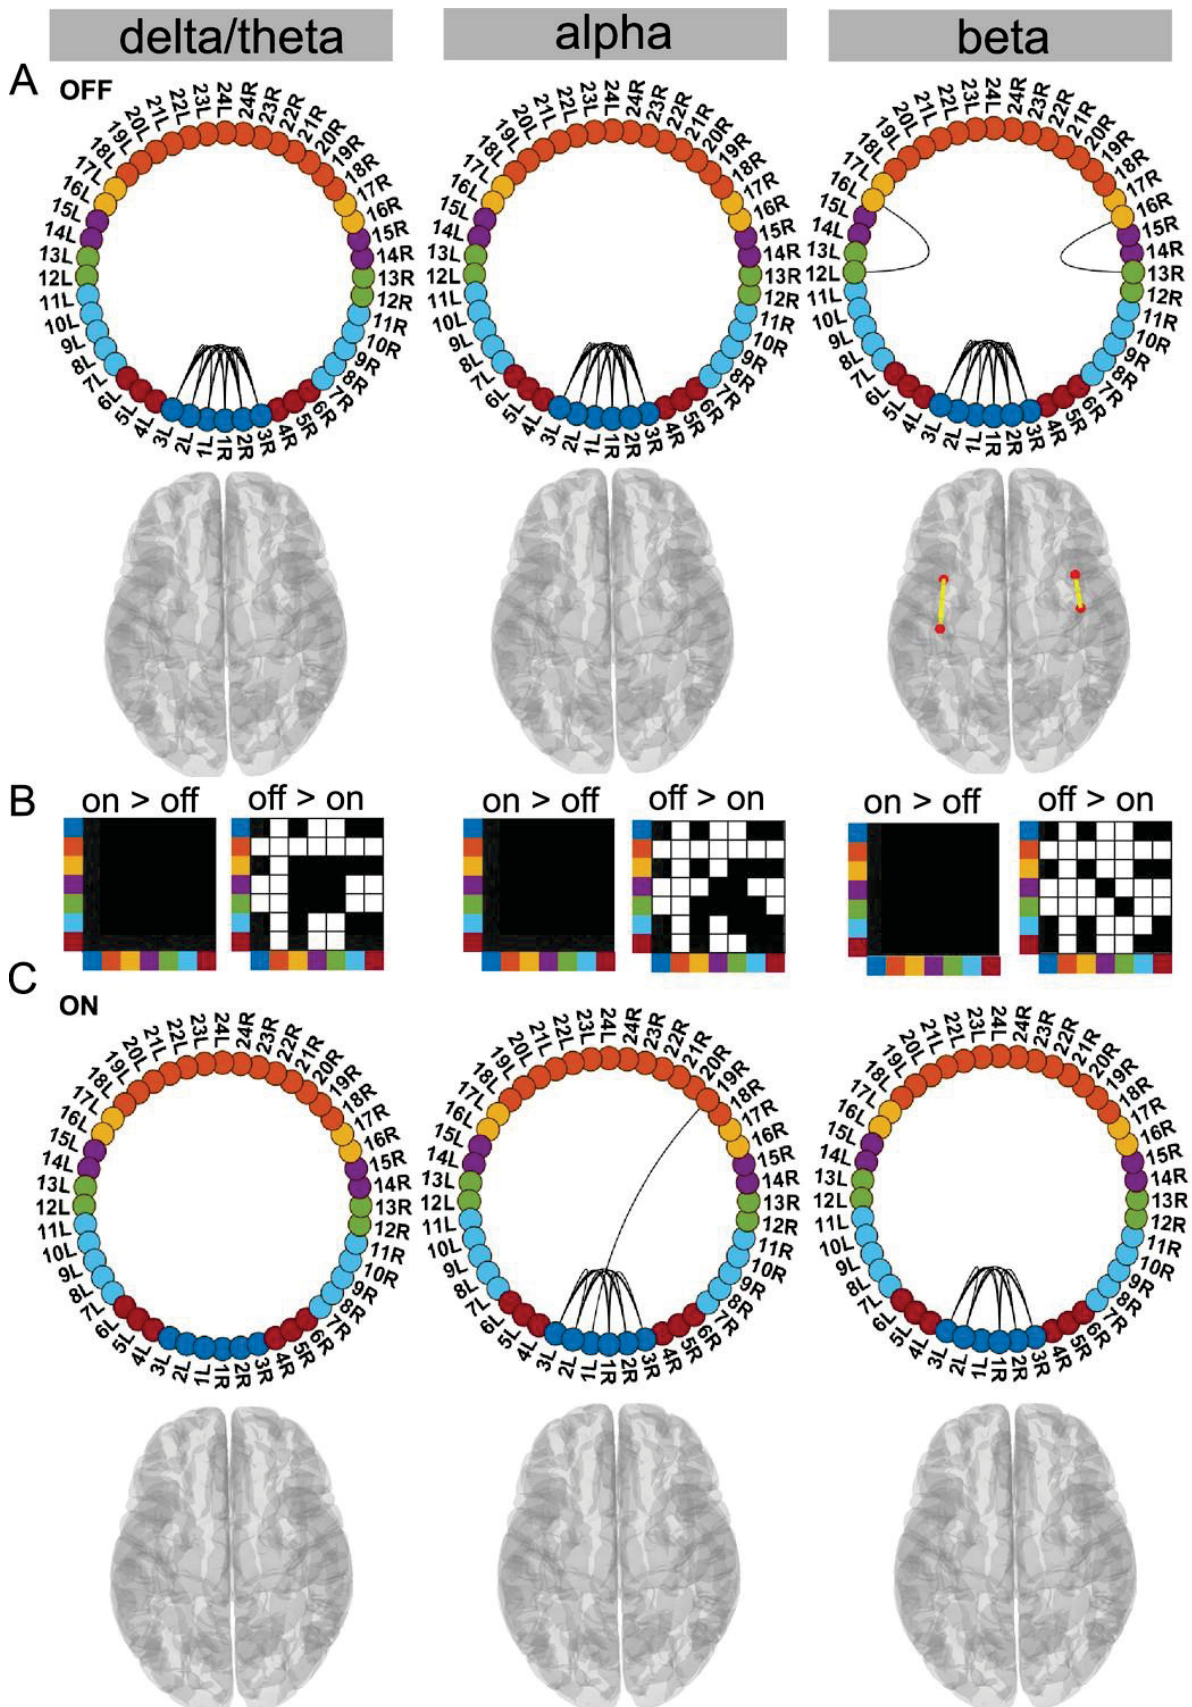


Figure 3.5: STN–STN state. For the general description, see the note to figure 3.2. The STN–STN state was characterised by preservation of STN–STN coherence in the alpha and beta band OFF versus ON medication. STN–STN theta/delta coherence was no longer significant ON medication.

Table 3.1: Regions of the Mindboggle atlas used.

List of regions in the connectivity ring figures

STN	1	contact1 right	12	postcentral	Smtr Ctx
	2	contact2 right	13	precentral	
	3	contact3 right	14	middletemporal	
Vis Ctx	1	contact4 left	15	superiortemporal	Tmp Ctx
	2	contact5 left	16	caudalmiddlefrontal	
	3	contact6 left	17	medialorbitofrontal	Mpf Ctx
	4	cuneus	18	insula	
	5	lateraloccipital	19	lateralorbitofrontal	Frnt Ctx
	6	lingual	20	parsopercularis	
Par Ctx	7	inferiorparietal	21	parsorbitalis	
	8	paracentral	22	parstriangularis	
	9	precuneus	23	rostralmiddlefrontal	
	10	superiorparietal	24	superiorfrontal	
	11	supramarginal			

STN- Subthalamic nucleus, Vis- Visual, Par-Parietal, Smtr-Sensory motor, Tmp-Temporal, Mpf-Medial prefrontal, Frnt-Frontal, Ctx- Cortex

3.3.2 Ctx–Ctx state is characterised by increased frontal coherence due to elevated dopamine levels

Supporting the dopamine overdose hypothesis in PD (Kelly et al., 2009; MacDonald and Monchi, 2011), we identified a delta/theta oscillatory network involving intra-hemispheric connections between the lateral and medial orbitofrontal cortex as well as the pars orbitalis. The delta/theta network emerged between the lateral and medial orbitofrontal as well as left and right pars orbitalis cortex ON medication ($p < 0.05$, Figure 3.3C delta). On the contrary, OFF medication no significant connectivity was detected in the delta/theta band. In the alpha and beta band OFF medication there was significant connectivity within the frontal regions, STN, and to a limited extent in the posterior parietal regions ($p < 0.05$, Figure 3.3A).

Another effect of excess dopamine was significantly increased connectivity of frontal cortex and temporal cortex both with the STN and multiple cortical regions across all frequency modes ($p < 0.01$, Figure 3.3 delta, alpha, and beta). The change in sensorimotor–STN connectivity primarily took place in the alpha band with an increased ON medication. Sensorimotor–cortical connectivity was increased ON medication across multiple cortical regions in both the alpha and beta band ($p < 0.01$, Figure 3.3 alpha and beta). However, STN–STN coherence remained unchanged OFF versus ON medication across all frequency modes.

Viewed together, the Ctx–Ctx state captured increased coherence across the cortex ON medication within the alpha and beta band. This, however, implies that ON medication, no connectivity strength was significantly higher than the mean noise level within the alpha or beta band. ON medication, significant coherence emerged in the delta/theta band primarily between different regions of the orbitofrontal cortex.

3.3.3 Dopaminergic medication selectively reduced connectivity in the Ctx–STN state

Our analysis revealed that the Ctx–STN state ON medication was characterised by selective cortico–STN spectral connectivity and an overall shift in cortex-wide activity towards physiologically relevant network connectivity. In particular, ON medication, connectivity between STN and cortex became more selective in the alpha and beta band. OFF medication, STN–pre-motor (sensory), STN–frontal, and STN–parietal connectivity was present ($p < 0.05$, Figure 3.4A alpha and beta). Importantly, coherence OFF medication was significantly larger than ON medication between STN and sensorimotor, STN and temporal, and STN and frontal cortices ($p < 0.05$ for all connections, Figure 3.4B alpha and beta). Furthermore, ON medication, in the alpha band only the connectivity between temporal, parietal, and medial orbitofrontal cortical regions and the STN was preserved ($p < 0.05$, Figure 3.4C alpha). Finally, ON medication, a sensorimotor–frontoparietal network emerged ($p < 0.05$, Figure 3.4C beta), where sensorimotor, medial prefrontal, frontal, and parietal regions were no longer connected to the STN, but instead directly communicated with each other in the beta band. Hence, there was a transition from STN-mediated sensorimotor connectivity to the cortex OFF medication to a more direct cortico–cortical connectivity ON medication.

Simultaneously to STN–cortico and cortico–cortical, STN–STN connectivity changed. In the ON condition, STN–STN connectivity was significantly different from the mean noise level across all three frequency modes ($p < 0.05$, Figure 3.4C). But on the other hand, there was no significant change in the STN–STN connectivity OFF versus ON medication ($p = 0.21$ delta/theta; $p = 0.25$ alpha; $p = 0.10$ beta; Figure 3.4B).

To summarise, coherence decreased ON medication across a wide range of cortical regions both at the cortico–cortical and cortico–STN level. Still, significant connectivity was selectively preserved in a spectrally specific manner ON medication both at the cortico–cortical (sensorimotor–frontoparietal network) and the cortico–STN levels. The most surprising aspect of this state was the emergence of bilateral STN–STN coherence ON medication across all frequency modes.

3.3.4 Dopamine selectively modifies delta/theta oscillations within the STN–STN state

In this STN–STN state, dopaminergic intervention had only a limited effect on STN–STN connectivity. OFF medication, STN–STN coherence was present across all three frequency modes ($p < 0.05$, Figure 3.5A), while ON medication, significant STN–STN coherence emerged

only in the alpha and beta band ($p < 0.05$, Figure 3.5C alpha and beta). ON medication, STN–STN delta/theta connectivity strength was not significantly different from the mean noise level ($p < 0.05$, Figure 3.5C delta).

OFF compared to ON medication, coherence was reduced across the entire cortex both at the inter-cortical and the STN–cortex level across all frequency modes. The most affected areas were similar to the ones in the Ctx–STN state, in other words, the sensorimotor, frontal, and temporal regions. Their coherence with the STN was also significantly reduced, ON compared to OFF medication (STN–sensorimotor, $p < 0.01$ delta/theta, beta; $p < 0.05$ alpha; STN–temporal, $p < 0.01$ delta/theta, alpha, beta; and STN–frontal, $p < 0.01$ delta/theta, alpha and beta; Figure 3.5B).

In summary, STN–STN connectivity was not significantly altered OFF to ON medication. At the same time, coherence decreased from OFF to ON medication at both the cortico–cortical and the cortico–STN level. Therefore, only significant STN–STN connectivity existed both OFF and ON medication, while cortico–STN or cortico–cortical connectivity changes remained at the mean noise level.

3.3.5 States with a generic coherence decrease have longer lifetimes

Using the temporal properties of the identified networks, we investigated whether states showing a shift towards physiological connectivity patterns lasted longer ON medication. A state that is physiological should exhibit increased lifetime and/or should occur more often ON medication. An example of the state time courses is shown in figure 3.6.

Figure 3.7A-C shows the temporal properties for the three states for both the OFF and ON medication conditions. Two-way repeated measures ANOVA on the temporal properties of the HMM states revealed an effect of HMM states on the fractional occupancy (FO) ($F(2,96)=10.49$, $p < 0.01$), interval of visits ($F(2,221513)=9783.13$, $p < 0.01$), and lifetime ($F(2,214818)=50.36$, $p < 0.01$). There was no effect of medication (L-DOPA) on FO ($F(1,96)=2.00$, $p=0.16$) and lifetime ($F(1,214818)=0.15$, $p=0.7026$). Medication had a significant effect on the interval of visits ($F(1,221513)=4202.96$, $p < 0.01$). Finally, we found an interaction between the HMM states and medication on the interval of visits ($F(2,221513)=1949.98$, $p < 0.01$) and lifetime ($F(2,214818)=172.25$, $p < 0.01$). But there was no interaction between HMM states and medication on FO ($F(2,96)=0.54$, $p=0.5855$).

We performed post hoc testing on the ANOVA results. OFF medication, the STN–STN state was the one with the longest lifetime (STN–STN $>$ Ctx Ctx, $p < 0.01$; STN–STN $>$ Ctx–STN, $p < 0.01$). The Ctx–STN state OFF medication had the shortest lifetime among all three states (Ctx–STN $<$ Ctx–Ctx, $p < 0.01$; Ctx–STN $<$ STN–STN, $p < 0.01$) and the shortest interval between visits (interval of visit Ctx–STN $<$ Ctx–Ctx, $p < 0.01$; Ctx–STN $<$ STN–STN, $p < 0.01$). The largest interval between visits was for the Ctx–Ctx state OFF medication (Ctx–Ctx $>$ STN–STN, $p < 0.01$; Ctx–Ctx $>$ Ctx–STN, $p < 0.01$). The FO for the STN–STN and Ctx–STN states was similar, but significantly higher than for the Ctx–Ctx state (STN–STN $>$ Ctx–Ctx, $p < 0.01$; STN–STN \approx Ctx–STN,

$p=0.82$; Ctx–STN >Ctx–Ctx, $p<0.01$). ON medication, the comparison between temporal properties of all three states retained the same significance levels as OFF medication, except for the lifetime of the Ctx–STN state, which was no longer significantly different from that of the Ctx–Ctx state ($p=0.98$). Within each medication condition, the states retained their temporal characteristics relative to each other.

Across medication conditions, significant changes were present in the temporal properties of the states. The lifetimes for both the STN–STN and Ctx–STN state were significantly increased by medication (ON >OFF: STN–STN, $p<0.01$; Ctx–STN, $p\leq 0.01$) but the lifetime for the Ctx–Ctx state was not significantly influenced by medication. The Ctx–Ctx state was visited even less often ON medication (interval: ON >OFF Ctx–Ctx, $p<0.01$). The interval between visits remained unchanged for the STN–STN and Ctx–STN states. The FO for all three states was not significantly changed from OFF to ON medication. In summary, the cortico–cortical state was visited least often compared to the other two states both OFF and ON medication. The cortico–STN and STN–STN states showing physiologically relevant spectral connectivity lasted significantly longer ON medication.




Figure 3.6: Example of a probability time course for the six hidden Markov model (HMM) states OFF medication. Note that within the main text of the paper, we are only discussing the first three states.

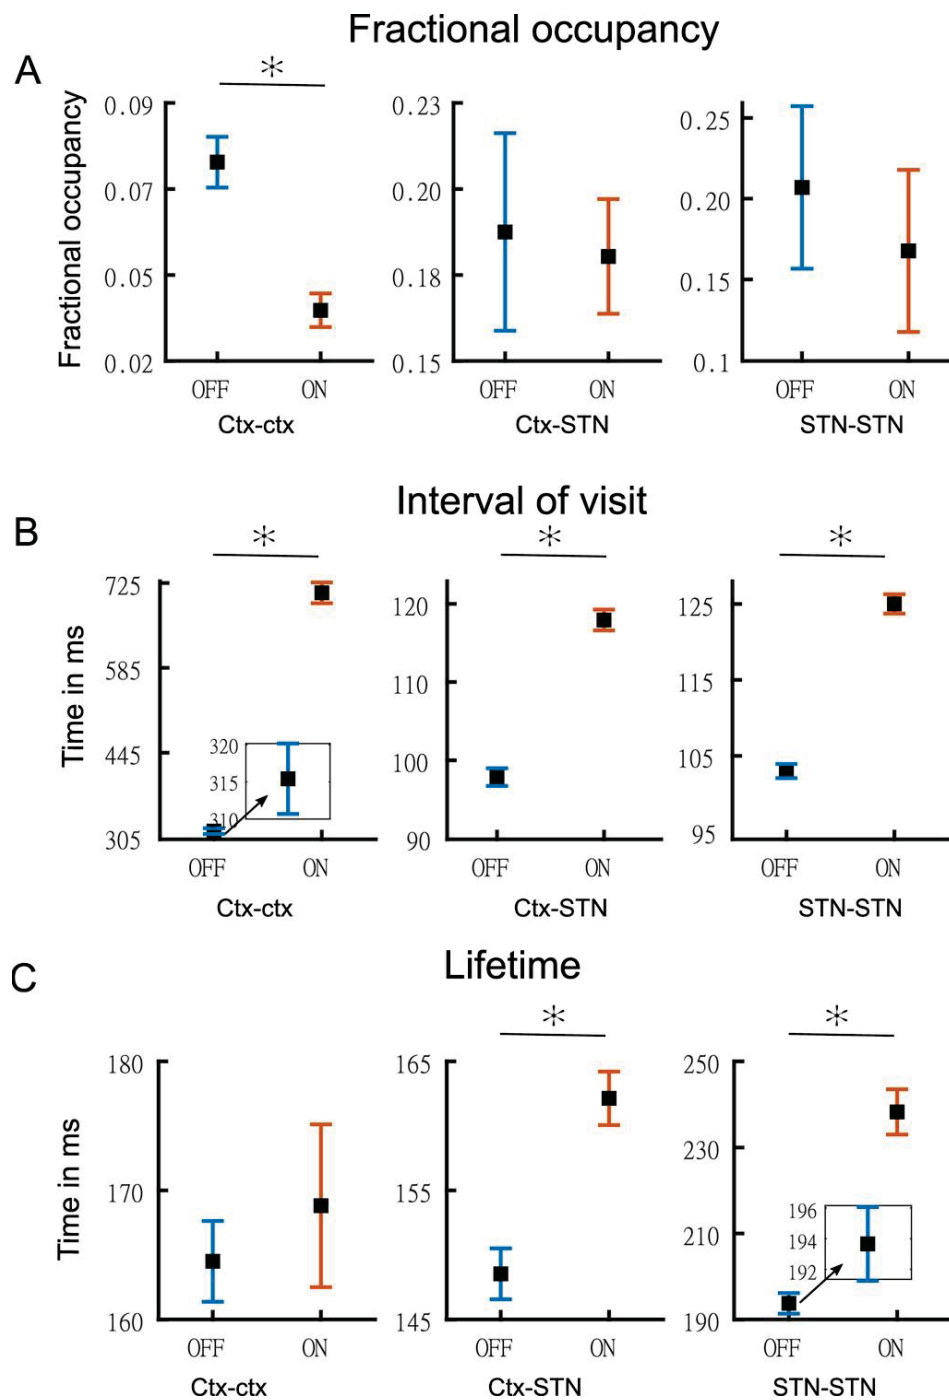


Figure 3.7: Temporal properties of states. Panel A shows the fractional occupancy for the three states for the cortico–cortical (Ctx–Ctx), cortico–STN (Ctx–STN), and the STN–STN (STN–STN). Each point represents the mean for a state and the error bar represents standard error. Orange denotes ON medication data and blue OFF medication data. Panel B shows the mean interval of visits (in milliseconds) of the three states ON and OFF medication. Panel C shows the lifetime (in milliseconds) for the three states. Figure insets are used for clarity in case error bars are not clearly visible. The y-axis of each figure inset has the same units as the main figure.

3.4 Discussion

In this study, we parsed simultaneously recorded MEG–STN LFP signals into discrete time-resolved states to reveal distinct spectral communication patterns. We identified three states exhibiting distinct coherence patterns ON and OFF medication: a cortico–cortical, a cortico–STN, and a STN–STN state. Our results indicate a tendency of neural activity to engage in connectivity patterns in which coherence decreases under the effect of dopaminergic medication and which maintain selective cortico–STN connectivity (Ctx–STN and STN–STN states). Only within the Ctx–Ctx state did coherence increase under dopaminergic medication. These results are in line with the multiple effects of dopaminergic medication reported in resting and task-based PD studies (Jubault et al., 2009; West et al., 2016; Tinkhauser et al., 2017).

The differential effect of dopamine allowed us to delineate pathological and physiological spectral connectivity. The Ctx–Ctx state provided electrophysiological evidence in the delta/theta band for the overdose effect of dopaminergic medication in PD. Prior to the electrophysiological evidence in our study, there was only evidence through task-based or functional magnetic resonance imaging (fMRI) studies (Cools et al., 2002; Ray and Strafella, 2010; MacDonald et al., 2011). The Ctx–STN state revealed that simultaneous cortico–cortical and STN–STN interactions emerge ON medication, with spectrally and spatially specific cortico–STN interactions. In addition, ON medication, a frontoparietal motor network was present, indicating a shift from STN-mediated motor connectivity to a cortical one. These findings have not been reported in previous studies. The STN–STN state exhibited the limited ability of dopaminergic medication to modify local STN–STN delta oscillations. Our analysis also revealed significant changes in the temporal properties of the connectivity profiles, including lifetime and FO, under the effect of dopaminergic medication. This insight might in the future prove important for modifying medication as well as DBS-based strategies for therapeutic purposes.

3.4.1 Increased tonic dopamine causes excessive frontal cortical activity

The Ctx–Ctx state showed significant coherent connectivity between the orbitofrontal cortical regions in the delta/theta band ON medication. According to the dopamine overdose hypothesis in PD (Cools, 2001; Kelly et al., 2009; MacDonald and Monchi, 2011; Vaillancourt et al., 2013), the commonly used doses of dopaminergic medication to mitigate the motor symptoms cause the ventral frontostriatal cortical circuits to experience an excessive increase in tonic dopamine levels. This medication-induced increase is due to excessive compensation of dopamine in the ventral striatal circuitry, which experiences a lower loss of dopamine than its dorsal counterpart. The reason is that in PD dopaminergic neurons in the substantia nigra are primarily lost and therefore the dopamine depletion within the dorsal circuitry is higher than within the ventral one (Kelly et al., 2009; MacDonald and Monchi, 2011). Frontal regions involved with the ventral striatal circuitry include the orbitofrontal cortex, anterior cingulate,

and the inferior temporal cortex (Cools, 2006; MacDonald and Monchi, 2011). Increased frontal cortex connectivity potentially explains the cognitive deficits observed in PD (Shohamy et al., 2005; George et al., 2013). Our detected emergence of frontal cortico–cortical coherence (between orbitofrontal and medial orbitofrontal regions) specifically in the delta/theta band could explain the cognitive deficits observed in PD due to dopaminergic medication, given the role of frontal delta/theta oscillations in cognition (Harmony, 2013; Zavala et al., 2014).

A comparison of temporal properties of the Ctx–Ctx state OFF versus ON medication revealed that the interval between visits was significantly increased ON medication, while the FO of this state was significantly reduced. In fact, the FO of the Ctx–Ctx state was the lowest among the three states. The temporal results indicate that the Ctx–Ctx state is least visited. Neural activity ON medication is not likely to visit this state, but whenever it does, its visits are of the same duration as OFF medication. Hence, the Ctx–Ctx state’s presence could explain the cognitive side effects observed ON medication in PD.

3.4.2 Selective spectral connectivity remains preserved with increased dopamine levels

An interesting feature of the Ctx–STN state was the emergence of local STN–STN coherence in all three frequency modes. Bilateral STN–STN coherence in the alpha and beta band did not change in the Ctx–STN state ON versus OFF medication (InterMed analysis). However, STN–STN coherence was significantly higher than the mean level ON medication (IntraMed analysis). Since synchrony limits information transfer (Cruz et al., 2009; Cagnan et al., 2015; Holt et al., 2019), the high coherence within the STN ON medication could prevent communication with the cortex. A different explanation would be that a loss of cortical afferents leads to increased local STN coherence. The causal nature of the cortico–basal ganglia interaction is an endeavour for future research.

Previous studies have reported STN–sensorimotor (Hirschmann et al., 2011; Litvak et al., 2011), STN–parietal, and STN–frontal (Litvak et al., 2011) coherence in the beta band OFF medication. Consistent with previous studies STN–sensorimotor, STN–parietal (inferior parietal), STN–frontal (insular cortex, pars orbitalis, pars opercularis, and lateral orbitofrontal), and STN–medial prefrontal (medial orbitofrontal) coherence emerged in the Ctx–STN state. In contrast, ON medication sensorimotor regions were coherent with parietal (para central) and frontal (superior frontal)/medial prefrontal (caudal middle frontal) regions in the beta frequency range. Previous research has not reported the emergence of such a coherent frontoparietal motor network ON medication. But consistent with previous research (Hirschmann et al., 2013), sensorimotor–STN coherence was reduced ON compared to OFF medication.

In addition, critical processing regarding sensorimotor decision-making involves frontoparietal regions (Gertz and Fiehler, 2015; Siegel et al., 2015; Gallivan et al., 2018; Martínez-Vázquez and Gail, 2018). Hence, the emergence of frontoparietal connectivity with

motor regions points towards the physiological relevance of the Ctx–STN state. Moreover, neural activity ON medication remained longer in the Ctx–STN state as the lifetime of this state significantly increased compared to OFF medication. The finding is in line with our hypothesis that a state showing physiologically relevant spectral connectivity lasts longer ON medication.

3.4.3 Tonic dopamine has a limited effect on local STN–STN interactions

In the Ctx–STN state, STN–STN coherence accompanied network changes affecting cortico–STN communication ON medication, thereby likely having a functional role. In contrast, in the STN–STN state, STN–STN coherence emerged without the presence of any significant cortico–STN coherence either OFF or ON medication. This may indicate that the observed STN–STN activity in the STN–STN state emerged due to local basal ganglia circuitry. No changes were observed in the alpha and beta band in the STN–STN state ON versus OFF medication, which may indicate the inability of tonic dopamine to modify basal ganglia circuit activity. These results provide more evidence that the changes in STN–STN coherence observed in previous studies (Little et al., 2013; Oswal et al., 2013; Shimamoto et al., 2013) reflect cortical interaction affecting STN activity. Future studies should analyse changes occurring within the STN. To the best of our knowledge, we are the first to uncover modulation of STN–STN delta/theta oscillations by dopaminergic medication. Studies have shown that local basal ganglia delta oscillations, which do not require input from the motor cortex, are robust biomarkers of dopamine depletion (Whalen et al., 2020). Hence, selective elimination of delta/theta oscillations under dopaminergic medication in the STN–STN state points towards restoration of physiologically relevant network activity.

3.5 Limitations of the study

In the present study, we employed a data-driven approach based on an HMM. In order to find the appropriate model, we had to specify the number of states a priori. We selected the number of HMM states based on a compromise between spectral quality of results and their redundancy. The number of states could also be determined by selecting the one with the highest negative free energy. However, model selection based on free energy often does not yield concrete results (Baker et al., 2014). Another limitation is the use of multivariate Gaussian distributions to characterise the state covariance matrices. Although it improves the tractability of the HMM estimation process, it is by construction unable to capture higher-order statistics of the data beyond the first two moments. For example, burst activity might also be a relevant property of brain networks (Florin et al., 2015). Lastly, we would like to note that the HMM was used as a data-driven, descriptive approach without explicitly assuming any a priori relationship with pathological or physiological states. The relation between biology and the HMM states, thus, purely emerged from the data; that is, is empirical. What we claim in this work is simply that the features captured by the HMM hold some relation

with the physiology even though the estimation of the HMM was completely unsupervised (i.e., blind to the studied conditions).

Besides these limitations inherent in the analysis approach, there are also some related to the experimental design. As this is a study containing invasive LFP recordings, we can never have a healthy control group. In addition, we only recorded four female patients because during the study period fewer female patients underwent a DBS surgery at our centre. To the best of our knowledge, there is no previous literature reporting a sex difference in MEG markers or the prescribed dopaminergic medication (Umeh et al., 2014). The medication led to a marked motor improvement in these patients based on the UPDRS, but the patients still have impairments. Both motor impairment and motor improvement can cause movement during the resting state in PD. While such movement is a deviation from a resting state in healthy subjects, such movements are part of the disease and occur unwillingly. Therefore, such movements can arguably be considered part of the resting state of PD. None of the patients in our cohort experienced hyperkinesia during the recording. All patients except for two were of the akinetic-rigid subtype. We verified that tremor movement is not driving our results. Recalculating the HMM states without these two subjects, even though it slightly changed some particular aspects of the HMM solution, did not materially affect the conclusions. A further potential influencing factor might be the disease duration and the amount of dopamine patients are receiving. Both factors were not significantly related to the temporal properties of the states.

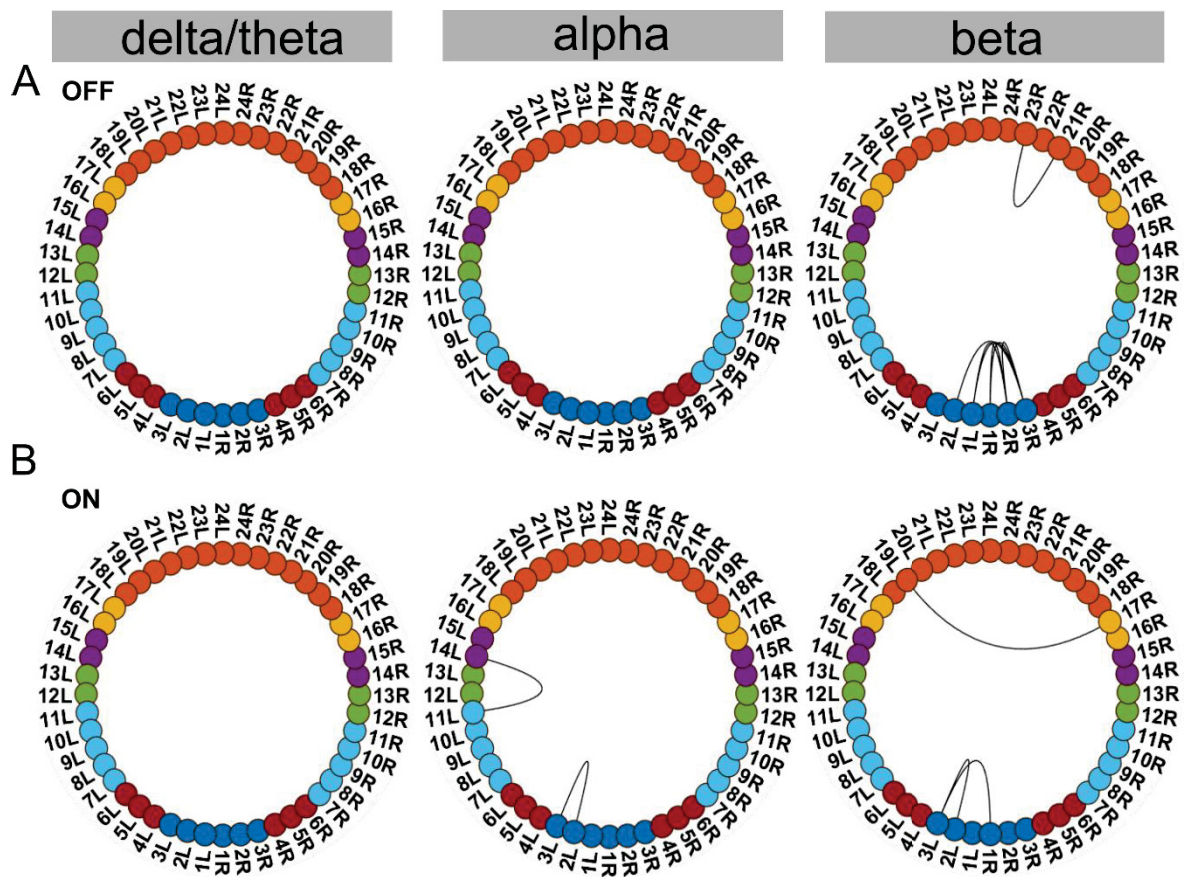
To differentiate pathological and physiological network activity, we had to rely on the temporal properties of the networks. A further limitation was that all our recordings were made under resting conditions, preventing us from discerning the functional role of oscillations within the discovered networks. We opted for the current design because resting-state data allows the study of networks independent of a task and because using a specific task bears the risk that the patients are not able to properly perform it. Nevertheless, future studies should analyse the behaviour of specific networks using tasks to probe them.

Lastly, we recorded LFPs from within the STN – an established recording procedure during the implantation of DBS electrodes in various neurological and psychiatric diseases. Although for Parkinson patients results on beta and tremor activity within the STN have been reproduced by different groups (Reck et al., 2010; Litvak et al., 2011; Florin et al., 2013; Hirschmann et al., 2013; Neumann et al., 2016), it is still not fully clear whether these LFP signals are contaminated by volume-conducted cortical activity. However, while volume conduction seems to be a larger problem in rodents even after re-referencing the LFP signal (Lalla et al., 2017), the same was not found in humans (Marmor et al., 2017). Moreover, we used directional contacts, which have a smaller surface area than the classical ring contacts. Based on the available literature, our sampling rate was high enough to resolve oscillatory activity in the STN (Telkes et al., 2020; Nguyen et al., 2020).

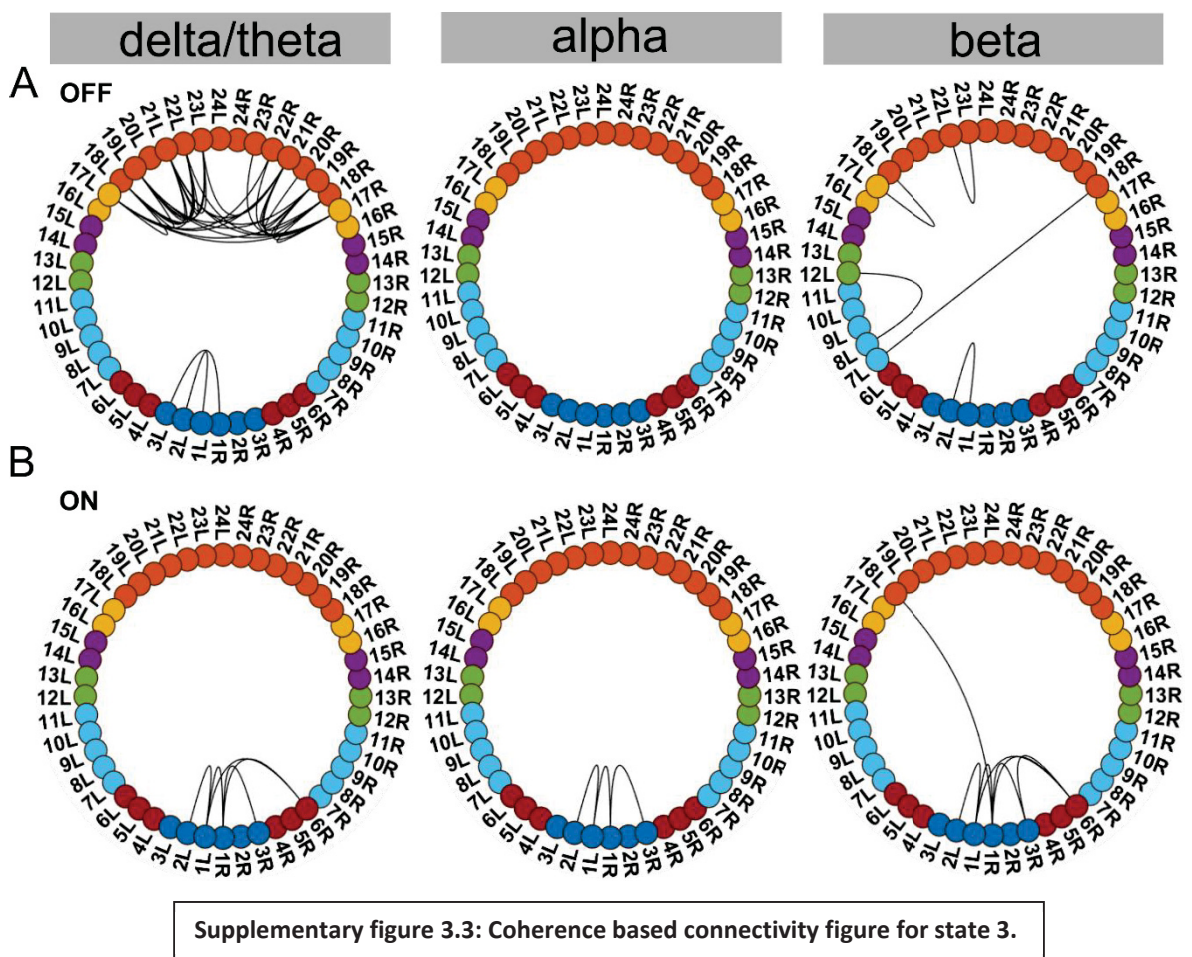
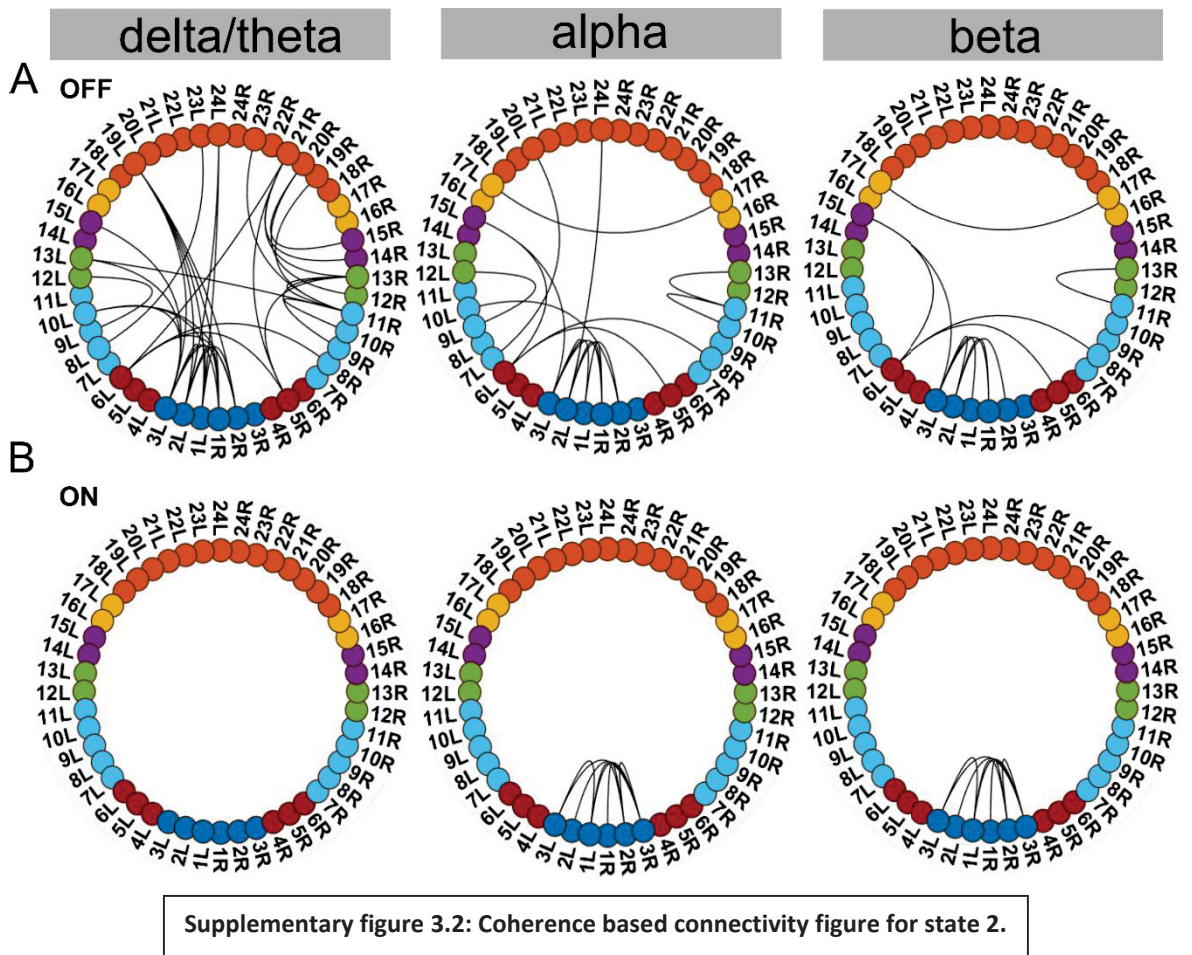
3.6 Conclusion

Using a data-driven machine learning approach, we identified three distinct networks (states) that captured differential effects of dopaminergic medication on spectral connectivity in PD. Our findings uncovered a Ctx–Ctx state that captured the potentially adverse effects of increased dopamine levels due to dopaminergic medication. Furthermore, a Ctx–STN state was identified that maintained spatio-spectrally selective cortico–STN connectivity ON medication. We also found an STN–STN coherent state, pointing towards the limited effect of dopaminergic medication to modify local basal ganglia activity. Our findings bring forth a dynamical systems perspective for differentiating pathological versus physiologically relevant spectral connectivity in PD. Furthermore, we were able to uncover differential changes induced by altered levels of a neuromodulator such as dopamine in a completely data-driven manner without providing detailed information about large-scale dopaminergic networks to the HMM. This shows another advantage of our dynamical systems-level approach. Furthermore, our whole-brain STN approach provides novel electrophysiological evidence of distributed changes due to dopaminergic medication in brain connectivity, extending previous pairwise connectivity results reported in PD.

3.7 Supplementary figures



Supplementary figure 3.1: Coherence based connectivity figure for state 1.



4 Study 3 Latent cognitive network transitions track somatosensory perceptual variability

Contributions

Abhinav Sharma: Conceptualization, Complete data acquisition and quality control, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (65% contribution)

Dr. Joachim Lange: Behavioral paradigm, Methodology

Dr. Diego Vidaurre: Software

Dr. Esther Florin: Conceptualization, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing

Abstract

Sensory perception is essential for transforming incoming information in the brain into a targeted behavior. Thus, perceptual variability is ubiquitously associated with human behavioral performance. Since our brains are active even while resting, the question arises whether spontaneous brain activity has any connection to perceptual variability observed during subsequent task performance. We intended to answer this question by analyzing whole-brain neural activity recorded in humans during a somatosensory discrimination task. By using a data-driven approach we were able to find time-resolved brain networks that were associated with different features of top-down cognition. We discovered that the transitions to and from these different cognitive networks before the stimulus occurred were correlated to the trial-by-trial perceptual variability. Furthermore, post-stimulus perception depended on the speed of network switching before the stimulus. Overall, these findings indicate that the interaction between whole-brain networks in the spontaneous pre-stimulus period contributes to perceptual variability during the task.

4.1 Introduction

Humans are not constantly engaged in tasks, but our brains are still buzzing with activity not pertaining to any specific cognitive demand. Imagine your daily morning routine for two days. Even if everything remains the same you are bound to produce variability (good or bad) on one or multiple tasks on the two days. Neural activity milliseconds before a task can be predictive of such variability (Linkenkaer-Hansen et al., 2004; Rajagovindan and Ding, 2011; Iemi et al., 2019; Podvalny et al., 2019; Rassi et al., 2019; Li et al., 2020). However, these findings do not consider the whole-brain network interactions, which provide us with deeper insight into the relevance of spontaneous, resting brain activity and its connection to perception (Thiebaut de Schotten and Forkel, 2022).

In our current work, we optimized an established tactile discrimination task (Baumgarten et al., 2015). During whole-brain MEG recordings, the participants receive two short consecutive electric pulses and indicate with a button press the percept of either one or two pulses. We determined the inter-pulse interval termed as stimulus onset asynchrony (SOA) for each individual with a staircase procedure so that the hit rate is 50% for the critical SOA. To enable an analysis of whole-brain networks the inter-trial interval was extended to 10 to 15 seconds. Importantly the critical SOA and stimulus remain the same during multiple trials yet the percept and therefore the response of the participant varies. We did not use cues to modulate attention and the tactile discrimination does not require cognitive or memory resources. Rather the percept and response purely depend on the internal dynamics of the whole-brain networks preceding the stimulus. To characterize the temporal dynamics and spectral properties of these spontaneously forming whole-brain networks we employ a hidden Markov model (HMM) (Vidaurre et al., 2018a, b) and relate the obtained networks to the behavioral outcome. Through this analysis, we aimed at identifying network properties and interactions, which would be predictive of the stimulus percept.

Our findings indicate that the stimulus percept varies depending on the network transitions preceding stimulus presentation. Critically, a balance between functionally constrained spectrally specific and spectrally broadband brain networks led participants to perceive the stimulus as two, whereas excessive transitions to the constrained networks led to the perception of one electrical pulse. Moreover, slower inter-network transitions were detrimental to perceiving two stimuli. Thus, for optimal tactile perception, transitions between whole-brain networks are necessary in the pre-stimulus epoch. As a result, the traditional view that increased or decreased spectral activity is responsible for perceptual discrimination should be extended to also include transitions between whole-brain networks.

4.2 Methods

4.2.1 Behavior

Thirty healthy participants (15 females and 15 males, right-handed, 26.4 (mean), 25 (median), 13 (range) years in age) volunteered to perform the task. The experimental procedure was explained to all participants and they gave written consent. The study was approved by the local ethics committee (study number 2019-477) and conducted in accordance with the Declaration of Helsinki.

We adopted the tactile discrimination task described in detail in Baumgarten et. al 2015. Each trial started with the presentation of a start cue (500 ms). Next, the cue decreased in brightness, indicating the pre-stimulus period (randomized between 10-15 seconds), after which the subjects received either 1 or 2 short (0.3 ms) electrical pulses, applied by 2 electrodes placed between the 2 distal joints of the left index finger. The time between the two pulses (Stimulus Onset Asynchrony (SOA)) was determined with a staircase procedure for each participant so that they had a hit rate of 50%. This SOA was termed a critical SOA. Participants had to respond to whether they perceived one or two stimuli with their right hand by pressing a button with either their right index or right middle finger. The index and middle finger options were randomly swapped at each trial. SOA was varied as follows: 0ms, crit SOA, crit SOA + 10ms, crit SOA – 10ms, and 100ms. Participants were recorded on two days: Day 1 for each participant started with a pre-task eyes open resting state MEG recording for 10 minutes. This was followed by the staircase procedure to determine the critical SOA during which no MEG data was recorded. Then MEG data was recorded during the task. The task was divided into 4 blocks of about 45-50 trials and lasted about 10-15 minutes. Participants were allowed to take breaks between blocks. Finally, a post-task eyes-open resting state recording was performed for 10 minutes. On day 2 the staircase procedure was repeated. This was followed by 4 measurement blocks of task during which MEG data were recorded. In total 200 crit SOA, 50 crit SOA – 10ms, 50 crit SOA + 10ms, 50 0ms, and 50 100ms trials per participant were acquired. On the first day crit SOA across subjects was 71ms +- 34.3ms and on the second day crit SOA across subjects was 79.8ms +- 34ms. The grand average crit SOA was 75ms +- 34ms.

4.2.2 Electrophysiological recording

MEG data were recorded on both days of task performance. We used a whole-head MEG system with 306 channels (Elekta Vectorview, Elekta Neuromag, Finland) housed within a magnetically shielded chamber. To remove eye blink and cardiac artifacts, electrooculography and electrocardiography were recorded simultaneously with MEG data. Data were acquired at a sampling rate of 2500Hz with an 833Hz low pass filter.

4.2.3 Pre-processing

All data processing and analyses were performed using Matlab (version R 2021b; Mathworks, Natick, MA) and the Brainstorm toolbox (Tadel et al., 2011). To ensure artifact-free data, we

inspected the data visually. The Neuromag system provides signal-space projection (SSP) vectors for the cleaning of external artifacts from the MEG channels, which were applied. The line noise was removed from all channels with a notch filter at 50, 100, 150, ..., 550, and 600 Hz with a 3 dB bandwidth of 1 Hz. Very noisy and flat MEG channels were excluded from further analysis. Time segments containing artifacts were marked in the time series. However, if artifacts regularly occurred only in one single channel, this whole channel was removed instead. Frequently arising artifacts following the same basic pattern, such as eye blinks or cardiac artifacts, were removed via SSP. All data were high-pass filtered with 1 Hz to remove movement-related low-frequency artifacts. Finally, the data were down-sampled to 1000 Hz. All trials with artifacts were excluded from further analysis. For each trial, we extracted pre-stimulus activity that ranged from -8.0 seconds to -0.1 seconds with 0 indicating the onset of the first electric stimuli. Each trial was mean corrected on a single MEG channel basis, where the mean channel vector was calculated using activity from -8.0 to -0.1 seconds. We refer to this pre-stimulus extracted time series as the pre-stimulus rest activity.

Cortically constrained source estimation was performed on these recordings at the participant level using each participant's anatomy. Therefore, using *FreeSurfer* (<https://surfer.nmr.mgh.harvard.edu/>, v.5.3.0) the participant's cortical surfaces were extracted from the individual T1-weighted MRI scans (3T scanner and 1 mm³ voxel size). We used the overlapping spheres method (Huang et al., 1999) with 306 spheres for the forward model. As inverse model we used a linearly constrained minimum variance (LCMV) beamformer. The data covariance matrix for the LCMV beamformer was computed per measurement block of each participant. The data covariance was regularized by setting the values below the median to the median eigenvalue of the data covariance matrix. The noise covariance was obtained from a five-minute empty room recording on the same day as the actual measurement. A single shared source kernel was obtained for each measurement block of every participant.

The source-reconstructed MEG data were projected to the default cortical anatomy (MNI 152 with 15,002 vertices) using *FreeSurfer* registered spheres extracted using *FreeSurfer* and MRIs of individual participants (Reuter et al., 2012) and then down-sampled temporally to 250 Hz for every trial and each participant. We used the *Mindboggle* atlas to spatially reduce the data dimensions. For each of the 42 cortical regions in the atlas, we obtained the first principal component from the vertices' time series within that region. To correct for volume conduction in the signal, symmetric orthogonalization (Colclough et al., 2015) was applied to these PCA (principal component analysis) time series. This resulted in a 42 by T multivariate time series for each trial and each participant. The row vectors of this orthogonalized matrix were z-scored. The entire pre-processing pipeline was applied to all trials of every participant irrespective of the response type on a trial.

To avoid that trials of the same response type, occur one after the other we shuffled trials across each participant. This ensures that the HMM algorithm does not get stuck in local optimization minima solely based on a single response type. Finally, to resolve sign ambiguity

inherent in source-reconstructed MEG data channels across participants, a sign-flip correction (Vidaurre et al., 2016) procedure was applied to entire dataset. Subsequently, this dataset was then fed to the TDE-HMM pipeline described below.

4.2.4 TDE-HMM pipeline

The HMM is a data-driven probabilistic algorithm that finds recurrent network patterns in a multivariate time series (Vidaurre et al., 2016; Vidaurre et al., 2018a). Each network pattern is referred to as a 'state' in the HMM framework, such that these networks can activate or deactivate at various points in time. Here onwards, 'state' or 'network' is used interchangeably. We used a specific variety of the HMM, the time delay embedded (TDE) - HMM, where whole-brain networks are defined in terms of both spectral power and phase coupling (Vidaurre et al., 2018b). Hence, for every time point, the HMM algorithm provides the probability that a network is active. Here onwards, a contiguous block of time for which the probability of a particular network being active remained higher than all the other networks is referred to as a 'state visit'. Thereby the HMM produced temporally resolved spatial networks for the underlying time series. In our approach, we also performed spectral analyses of these state visits, leading to a complete spatio-spectral connectivity profile across the cortex.

4.2.5 HMM model fitting

Since we were interested in recovering phase-related networks, the TDE-HMM was fit directly on the time series obtained after pre-processing steps described previously, as opposed to its power envelope. This preserved the cross-covariance within and across the underlying raw time series of the cortical regions.

We opted for six different states as a reasonable trade-off between the spectral quality of the results and their redundancy. The HMM-MAR toolbox (Vidaurre et al., 2016) was used for fitting the TDE-HMM. The embedding took place in a 60 ms window (i.e., a 15-time point window for a sampling frequency of 250 Hz) for each trial separately. Since time embedding increases the number of rows of the data from 42 to 42 times the number of samples of the window, an additional PCA step was performed along the spatial dimension. The number of components retained were 84 (42 times 2) in accordance to our previous pipeline (Sharma et al., 2021) and as recommended in Vidaurre et al. 2018. To characterize each state, a full covariance matrix with an inverse Wishart prior was used. The diagonal of the prior for the transition probability matrix was set as 10. To ensure that the mean of the time series did not take part in driving the states, the 'zero mean' option in HMM toolbox was chosen. To speed up the process of fitting, we used the stochastic version of variational inference for the HMM. In order to start the optimization process, the 'HMM-MAR'-type initialization was used (for details, see Vidaurre et al., 2016). A single HMM was fit across all participants and trials irrespective of the response type and SOA.

4.2.6 Data-driven frequency bands

We investigated the spectral connectivity patterns across the different states. The objective was to uncover significant coherence within each frequency band (delta/theta [1–8 Hz], alpha [8–12 Hz], beta [13–30 Hz], and high-beta/gamma) in the respective states. The HMM output includes the state time courses (i.e., when the states activated) for each trial for every participant stored in the HMM model output variable called “gamma”. 0.85 was used as the probability threshold for gamma. The state time courses allowed the extraction of state-specific data for each trial for further analysis. For each HMM state, we filtered the state-specific data for all trials between 1 and 45 Hz. (For state-wise data extraction, please refer to the HMM toolbox wiki [<https://github.com/OHBA-analysis/HMM-MAR/wiki/User-Guide>].) Then we calculated the Fourier transform of the data using a multi-taper approach to extract the frequency components from the short segments of each state visit (See Vidaurre et al., 2018b). Seven Slepian tapers with a time-bandwidth product of 4 were used, resulting in a frequency resolution of 0.5 Hz. Subsequently, we calculated the coherence and power spectral density of these binned frequency domain data for every state within each trial separately. For each trial, the coherence and the power spectral density obtained were three-dimensional matrices of size f (number of frequency bins) by N (42 cortical locations) by N . We call these trial-level coherence matrices.

Based on the trial-level coherence matrices, we performed a frequency band-specific analysis. Canonical definitions of frequency bands assign equal weight to each frequency bin within a band for every participant. This might not be suitable when considering analyses of brain signals across a large dataset. Assigning for example the same weight to each bin in the alpha range might reduce the alpha effect at the group level. To allow for inter-participant variability in each frequency bin’s contribution to a frequency band, we determined the frequency modes in a data-driven manner (Vidaurre et al., 2018b). For each subject and all their trials, the lower triangular portion of the trial-level coherence matrix obtained above was vectorized across columns. This resulted in 903 (lower triangular entries of a 42 by 42 matrix including the diagonal) by f (number of frequency bins) matrices for each trial. Subsequently, we averaged all matrices across trials and participants along the spectral dimension (number of frequency bins) resulting in a single 903 by f matrix per state. This procedure was repeated for all six states. Finally, we concatenated matrices across all states resulting in a group-level coherence matrix (f by 903 by 6). We factorized the group-level coherence matrix into four frequency modes using a non-negative matrix factorization (NNMF) (Lee and Seung, 2001).

The values of frequency modes are the actual NNMF weights obtained from the NNMF estimation (which, just like a regression coefficient, are unitless, because coherence is unitless). Three of them resembled the canonical alpha, beta, and high beta/gamma bands whereas the fourth one represented the $1/f$ noise in neural signals. Since NNMF does not guarantee a unique solution, we performed multiple instances of factorization. In practice, we could obtain frequency modes, which showed correspondence to the classical frequency bands, within four iterations of the algorithm. At each instance, we visualized the output to

ensure the frequency specificity of the frequency modes. The stability of the output was ensured by using 'robust NMF', which is a variant of the NMF algorithm (Vidaurre et al., 2018). We then computed the inner product between the trial-and group level coherence matrix and the frequency modes obtained above. We called these the trial-level and group-level coherence projections, respectively. While these frequency modes were derived from coherence, they can be applied to power measures or any other frequency-specific measure. We applied these factors to the power spectra calculated for each of the 42 regions as well to obtain trial-level and group-level power projections.

4.2.7 Group-level coherence ring figures

To separate background noise from the strongest coherent connections, a Gaussian mixture model (GMM) was used (Vidaurre et al., 2018b). For the group-level coherence projections, we normalized the activity in each state for each spectral band by subtracting the mean coherence within each frequency mode across all states. As a prior for the mixture model, we used two single-dimensional Gaussian distributions with unit variance: one mixture component to capture noise and the other to capture the significant connections. This GMM with two mixtures was applied to the coherence values of each state. Connections were considered significant if their p-value after correction for multiple comparisons was smaller than 0.001.

4.2.8 Power spectral density test

We used permutation testing to assess power increases within a brain region relative to others within a state. For this, we used the trial-level power projections. We subtracted the mean power across all states and all trials from each of the 42 regions within a state for each trial. We then permuted the labels for the 42 brain regions within a state across trials to obtain a null distribution. The p-value was calculated as the number of permutations for which the power of a specific brain region was higher or lower than the shuffled set divided by the total number of permutations. We used 5000 permutations to compute our results.

4.2.9 Diffusion map analysis

To recover a low- dimensional network embedding from high-dimensional cortical data we used diffusion maps (Coifman et al., 2005; Lian et al., 2015). In order to calculate state-specific embedding we masked the original dataset, i.e. the one without any delay embedding, using the state probability time (HMM model output gamma). The probability threshold for gamma was the same as used to perform the multitaper spectral analysis. State probability time course masking resulted in state-specific 42-dimensional vectors. For clarity, we will call them diffusion vectors. For each state individually we computed a covariance matrix using all of the state's diffusion vectors across all trials and participants. For each state covariance matrix, an affinity matrix was computed using Spearman rank order correlation. This affinity matrix was then fed to the diffusion map algorithm to obtain a low-dimensional embedding. The dimensionality was reduced from 42 to 5. Only the top three dimensions were used for subsequent analysis. The diffusion parameter alpha was set to 0.5 for all states as

recommended for large-scale brain data (Margulies et al., 2016). The above procedure was implemented using custom-written scripts and the Brainspace toolbox (Vos de Wael et al., 2020).

4.2.10 Uniqueness of the state-specific embedding

In order to test if the state-specific embedding was unique for each state we bootstrapped with 100 iterations. For each bootstrap iteration for each state, we randomly sampled with repetition the diffusion vectors. The sample covariance matrix is a biased estimator and the bias is dependent upon the number of samples that are used to estimate covariance using the estimator (Smith, 2005). Hence, we kept the total number of diffusion vectors used to calculate the state-specific covariance matrix the same as the original number used to calculate state-specific embedding. Then diffusion map analysis was applied to each of these 100 bootstrapped covariance matrices. This step yielded six hundred low-dimensional embeddings i.e. 100 per state. Calculation of the diffusion map involves eigenvalue decomposition (Coifman et al., 2005). The resultant eigenvectors are unique up to a +/- sign. In order to align the bootstrapped low-dimensional embeddings we used the Procrustes analysis (Vos de Wael et al., 2020). Importantly, this does not affect the magnitude of the eigenvector basis but only rotates them. For each state, the bootstrapped embeddings were aligned to the original state-specific embedding. Pearson's correlation was calculated to measure the similarity between a state's original state-specific embedding and all other states (including its own) bootstrapped embeddings. Subsequently, the resulting correlation values were subjected to repeated measures ANOVA and posthoc testing. Since an embedding for the connectivity manifold comprises the first three eigenvectors of the diffusion map, the Pearson correlation and the subsequent ANOVA analysis were calculated separately for each eigenvector. See supplementary table 4.1 for the ANOVA results pertaining to the first three eigenvectors.

4.2.11 Statistical testing on the connectivity manifold

We refer to the state-specific embedding as a manifold. The key property of this manifold is the distance between the different brain regions on this manifold. Distance between different brain regions is directly related to the connectivity between regions (Margulies et al., 2016; Vos de Wael et al., 2020), i.e. the smaller the distance stronger the connectivity. To test if certain brain regions showed significant connectivity on the manifold we applied the same test as described previously in the section Methods: Group level coherence ring figures.

4.2.12 Transition matrix test

To obtain transition matrices for each trial we estimated the HMM on a single trial level using the group-level model. The critical SOA and the critical SOA +/- 10ms trials were included in this analysis. Non-parametric permutation testing (5000 permutations) was performed to compare transitions between response one and response two trials. Results were considered significant if $p < 0.01$ after correction for multiple comparisons using Benjamini Hochberg FDR correction.

4.2.13 Temporal properties

To test for changes in the temporal properties for percept one versus percept two trials, we compared the lifetimes, interval between visits, FO, and switching rate for each state both within and across the response types using two-way repeated measures ANOVA followed by post-hoc tests. The lifetime/dwell time of a state refers to the time spent by the neural activity in that state. The interval of visit was defined as the time between successive visits of the same state. FO is defined as the average state probability across time for a trial. Finally, the switching rate is defined as the number of times one state switches to another per second within a trial.

4.3 Results

Using an HMM we delineated spontaneous source-level cortical activity during the 8 seconds preceding stimulus presentation into distinct states or networks. In comparison to classic sliding window analysis, an HMM solution can be thought of as a data-driven estimation of time windows of variable length within which a particular HMM state was active. A single 6-state group level time-delay embedded (TDE)-HMM model was trained on spontaneous pre-stimulus activity across critical SOA trials, which included the individual critical SOA as well as the critical SOA \pm 10ms trials (see Methods: section Behavior) irrespective of the response type. As our stimuli were designed to achieve a 50% hit rate at the critical SOA we hypothesized that the same spontaneous networks would be relevant, but the interactions between these networks in the pre-stimulus period would be different depending on the percept. Furthermore, we used our group-level model to infer single-trial dynamics to explain the variability in the percept.

4.3.1 Spontaneous pre-stimulus activity organizes into different spectral networks

Out of the six states, four states had diverse spectral and spatial characteristics whereas the remaining two (states 5 and 6) had sparse spectral connections. Hence, we focus mainly on these four states. For each state, the coherence-based connectivity between different brain regions was visualized for the frequency modes presented in figure 4.1. In contrast to the traditional definition of frequency bands in electrophysiological data, the frequency modes are obtained from non-negative matrix factorization of the Fourier-transformed data.

As seen in Figure 4.2A state 1 consisted of an interacting anterior (frontal) and posterior (parietal) network in the alpha and beta bands. Due to the aforementioned characteristics, this state is termed the cross-brain state in the following. Posterior parietal connectivity in the alpha band dominated the state in Figure 4.2B. Parietal regions were coherent with different brain areas both within the parietal cortex and across the brain ($p < 0.001$). Interestingly, there was no connectivity within the beta and the high-beta/gamma bands. Hence, we term this state the alpha-parietal state. The state in Figure 4.2C is dominated by high-beta/gamma band and alpha-band coherence in the frontal and medial-frontal regions ($p < 0.001$) so it is termed the frontal state. The last state in Figure 4.2 had high-beta/gamma

band coherence between and within frontal and parietal areas ($p < 0.001$). Thus, we referred to it as the high-beta/gamma state.

Hence the HMM resolved the pre-stimulus spontaneous activity into distinct states with specific spatio-spectral fingerprints. Critically, the HMM states are the same irrespective of the response/percept of participants.

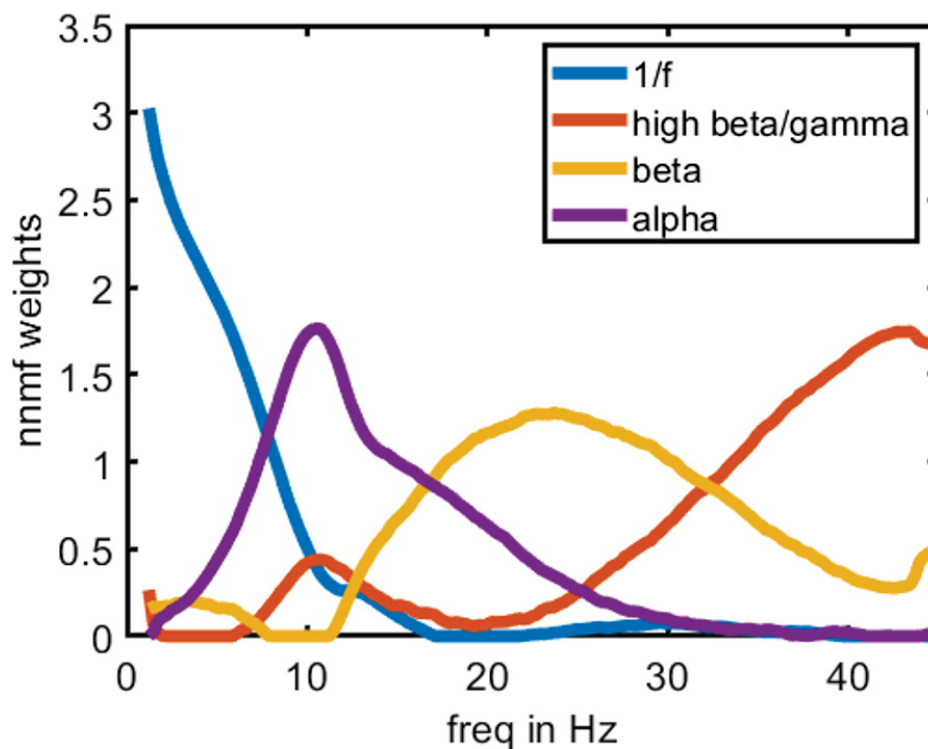


Figure 4.1: Data driven frequency bands, referred to as frequency modes in the text. See Methods: Data driven frequency bands

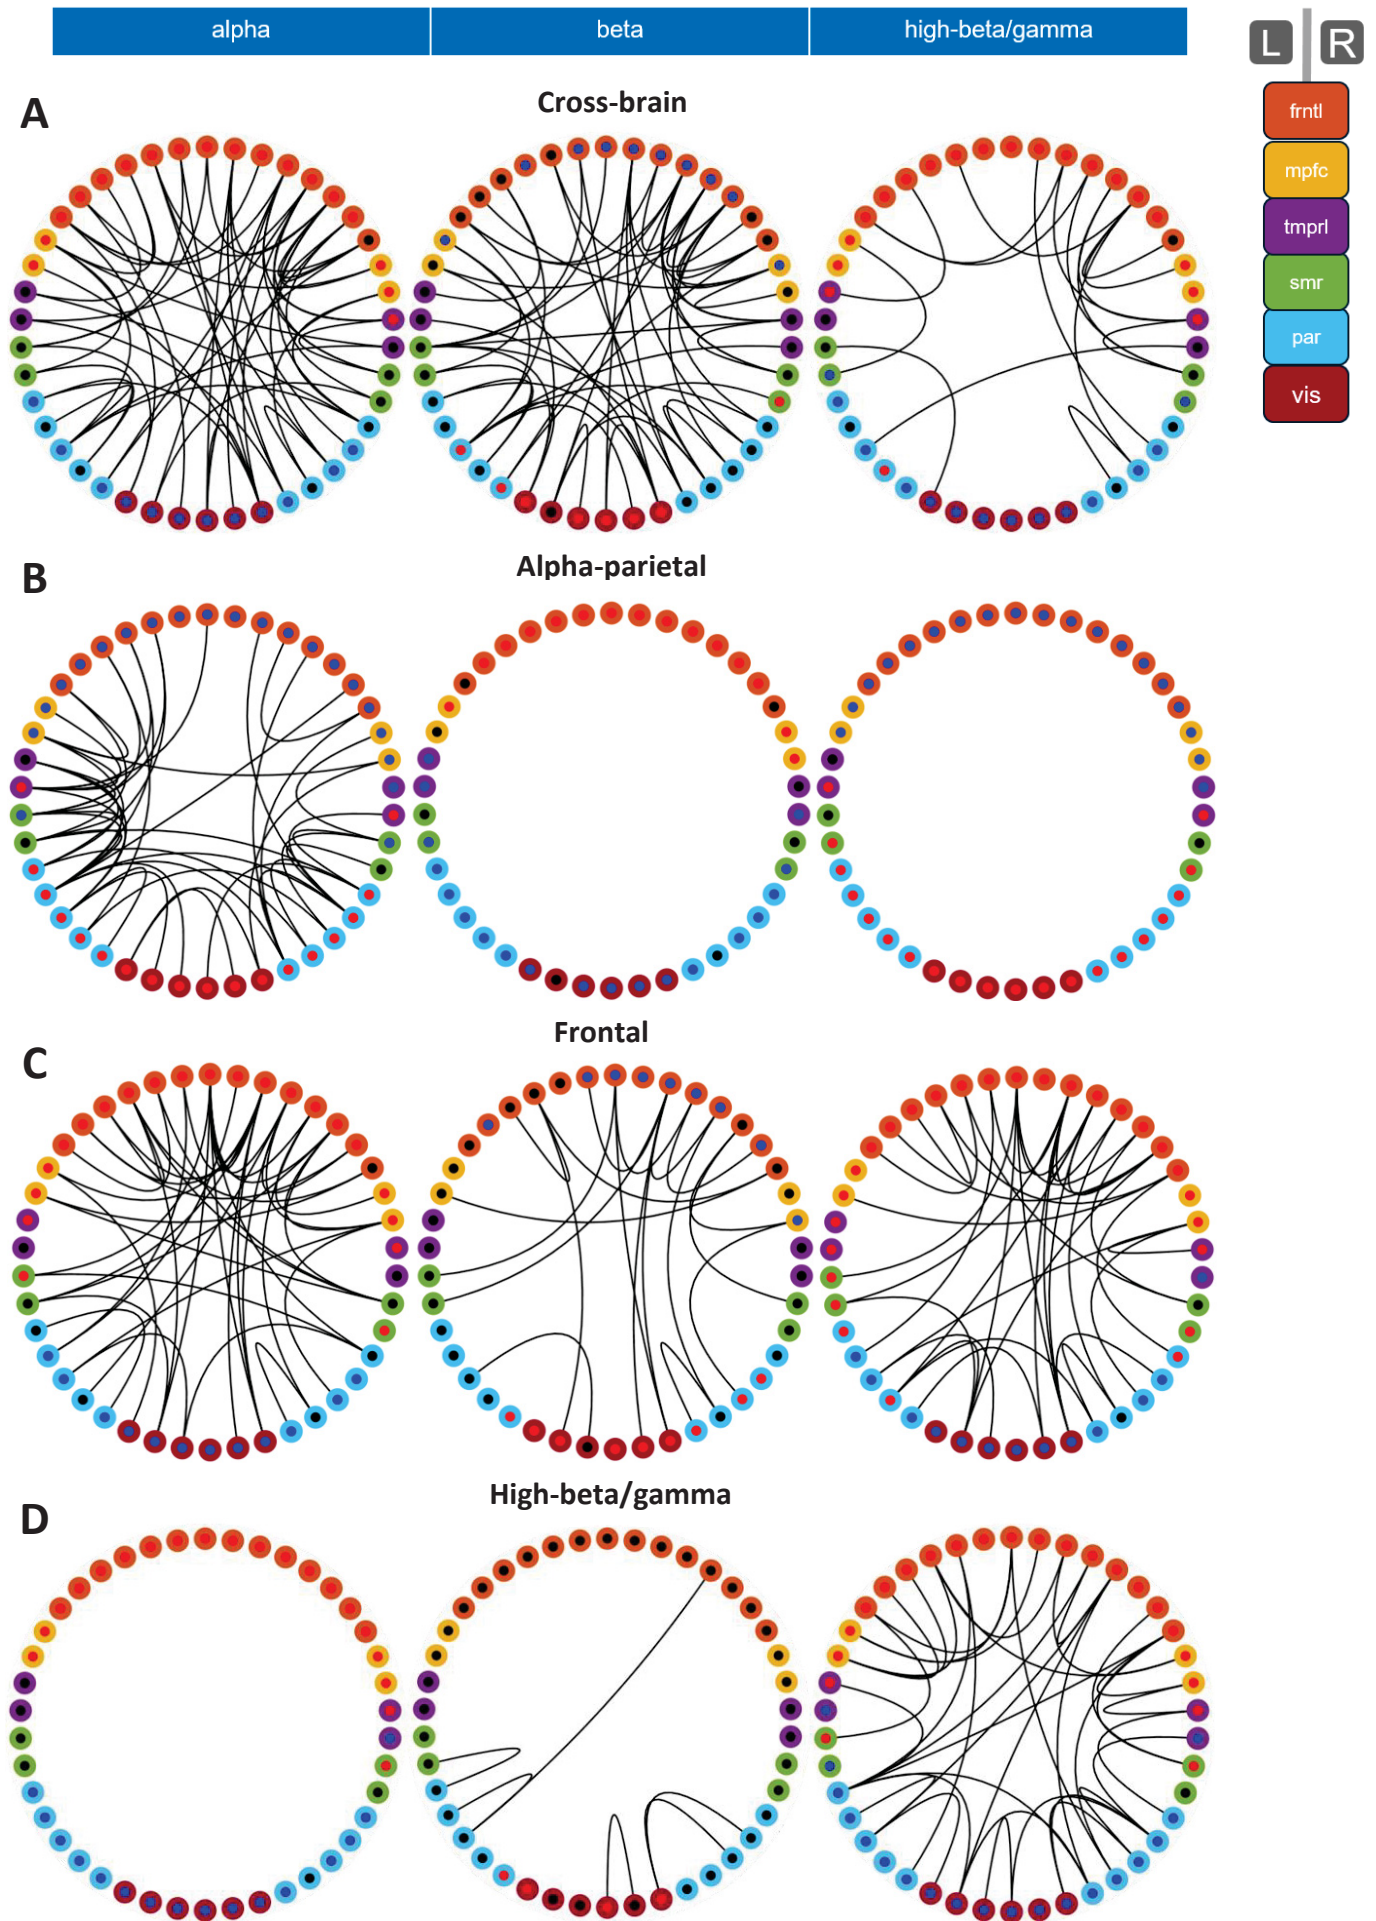


Figure 4.2: A) Coherence ring figure for the Cross Brain State. Each node of the ring contains two colours. The outer color represents a brain region (see Supplementary figure 4.1 for a list of brain regions). The inner color represents the outcome of the power spectral density test (see Methods: Power spectral density test). Inner color red represents significantly higher than mean power for that state and within a specific frequency band. Inner color blue represents significantly lower than mean power. Inner color black represents no significant difference than mean power for that state within a specific frequency band. The black coloured curve connecting a pair of two brain regions represents significant coherence between the pair, $p < 0.001$ (see Methods: Group level coherence ring figures). The coherence ring figures are shown for the alpha, beta and high-beta/gamma frequency bands shown in figure 4.1. B) Alpha-parietal state. This state was characterised posterior alpha band connectivity. Visuo-parietal regions showed significant increase in power in the alpha band along with significantly increase alpha band coherence. Beta and high-beta gamma band did not show any significant connectivity for this state. C) Frontal state. Frontal regions showed connectivity both within frontal cortex regions and to different brain regions. This connectivity pattern was observed across all three frequency bands. Hence this state can be characterised as a wideband frontal state. D) High-beta/gamma state. Fronto-parietal regions showed significant connectivity only within the high-beta/gamma band. Significant connections were present both within frontal and parietal regions as well between the regions.

4.3.2 Latent HMM states track multiple cognitive networks

To perform multiple functions a brain region might operate simultaneously in different frequencies which enables flexibility in the brain region's inputs and outputs (Akam and Kullmann, 2010, 2014). In line with this, our previous result indicates that multiple pairs of brain regions are coherent across several spectral bands within the same state. But, the same pairs were coherent in the same spectral band across different states. Such pairwise spectral coherence results lead to difficulties in interpreting HMM-based time-resolved states. To untangle the relevance of these pairwise connections within each state we employed the diffusion map framework (Lian et al., 2015; Margulies et al., 2016). With a diffusion map, we can embed the extracted non-linear relationship between different brain regions in Euclidean space. We refer to the resultant Euclidean space as the connectivity manifold. In this framework, two brain regions either directly connected to each other or connected via multiple brain regions will be embedded close to each other on the manifold. Hence, the extent of connectivity on the manifold between two brain regions is inversely proportional to the distance on it. Of note for the diffusion framework, the data has not been separated into the frequency modes; as a result, the diffusion maps provide an additional separate description of the HMM states.

Given the spectrally broadband and spatially non-specific characteristics of the cross-brain state, we expected its manifold to be densely connected. Surprisingly only a few brain regions were significantly connected with their neighbors. Thus, multiple pairwise spectral connections do not translate to dense connectivity on the manifold and this allows us to disentangle the spectral interactions between the frontal state and the cross-brain state. Most relevant for the cross-brain state the ipsi- and contralateral precentral gyri were connected to each other and embedded within a frontal network (Figure 4.3 A). But, the postcentral gyri of the two hemispheres were disconnected on the manifold. Of note, the right inferior parietal gyrus (IP-r) was connected via two frontal regions to the ipsi- and

contralateral precentral gyri (Figure 4.3 A, supplementary Figure 4.6A). Plausibly, the cross-brain state represents top-down controlled motor cortical connectivity.

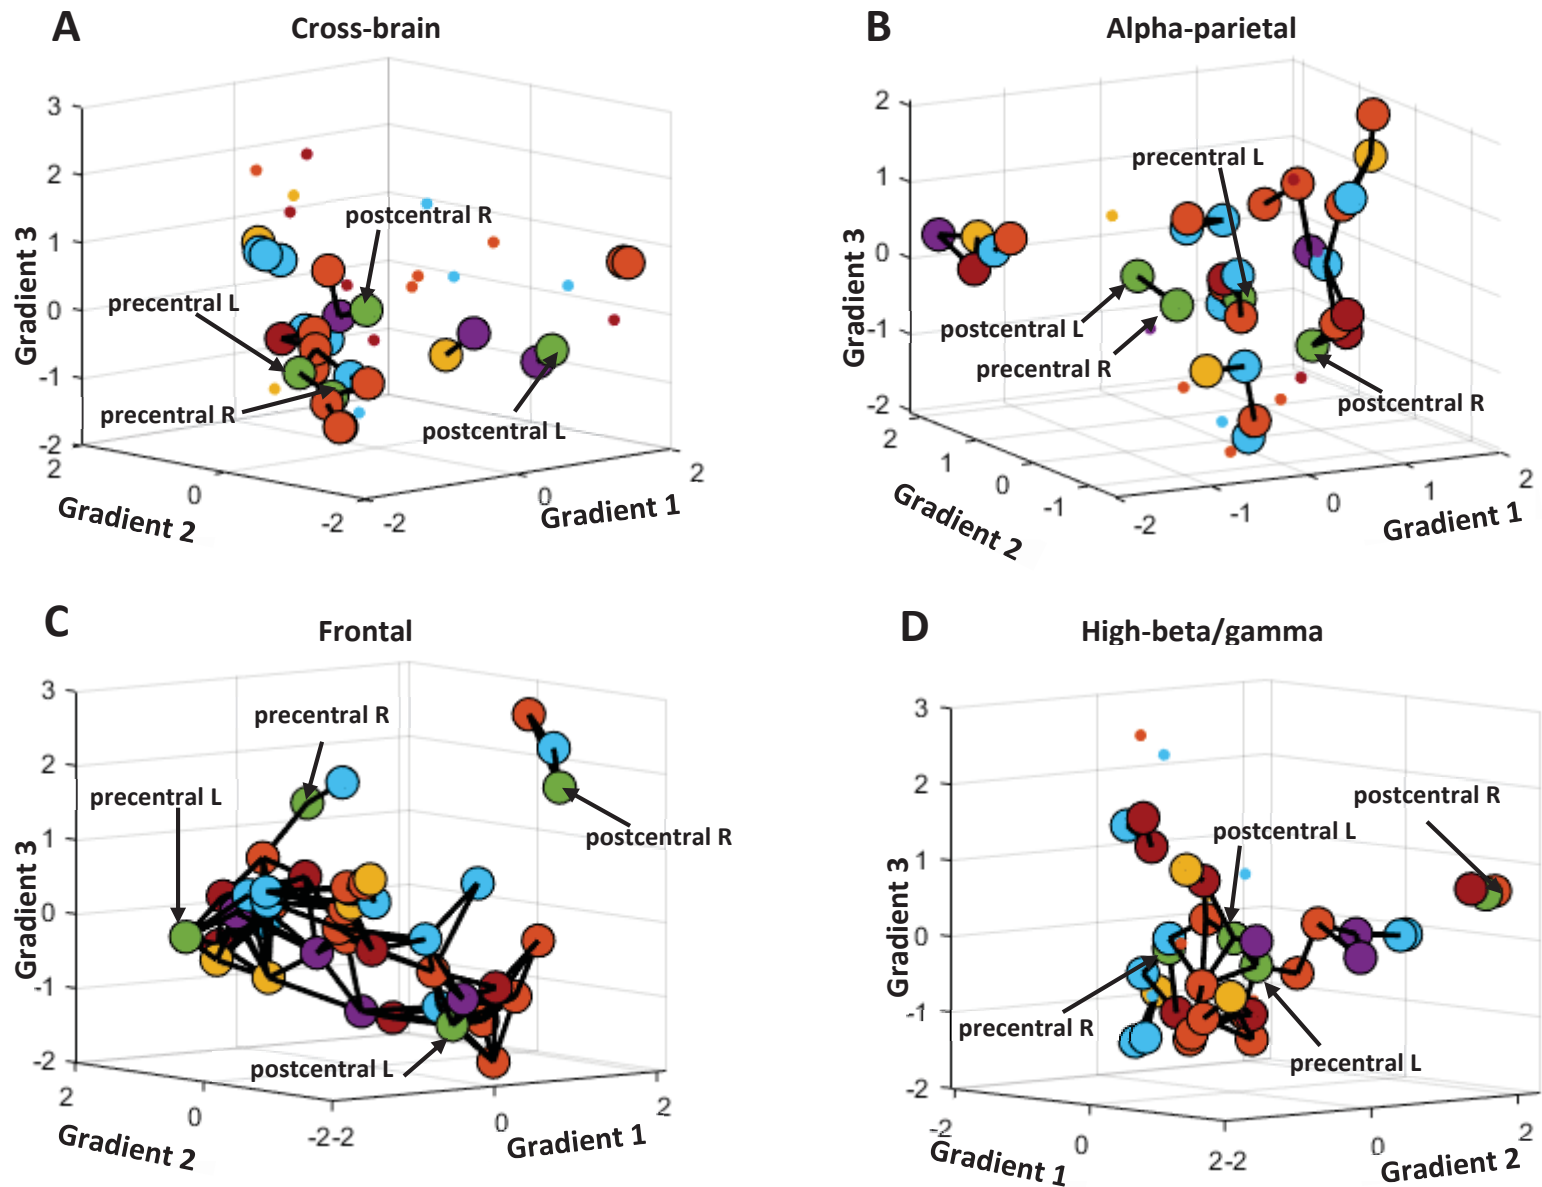


Figure 4.3: State specific manifolds for A) Cross-brain state B) Alpha-parietal state C) Frontal state D) High-beta/gamma state. In each figure the nodes represent brain regions from the Mindboggle atlas. The large dots are the ones for which a significant connection was found. The black lines are significant short distances that represent connectivity between brain regions (see Methods: Diffusion map analysis, statistical testing on the connectivity manifold).

In the alpha-parietal state, the ipsilateral postcentral gyrus and contralateral precentral gyrus were significantly connected on the manifold. Furthermore, fragmented connectivity occurred between parietal cortex regions and frontal cortical areas, ipsilateral motor cortex, superior temporal, and medial orbitofrontal regions (Figure 4.3 B). Taken together, the spectral characteristics and the connectivity manifold are consistent with the modulating role of parietal alpha activity for attention (Klimesch, 2012; Vollebregt et al., 2016; Helfrich et al., 2017). Importantly, dense pair-wise alpha-band coherence between brain regions in the alpha-parietal state resulted in a sparse pair-wise connectivity manifold dominated by parietal regions. This finding relates to the more generic role of posterior alpha activity in gating cortical processing (Jensen and Mazaheri, 2010; Foxe and Snyder, 2011; Zhou et al., 2021).

For the frontal state, as opposed to the spectrally constrained states, each and every brain region demonstrated statistically significant connectivity on the manifold (Figure 4.3 C). Still, it was not an all-to-all network but restricted to the significantly connected neighbors on the manifold. Frontal cortical regions perform their heterogeneous and complex functions via multiplexed oscillatory coding (Benchenane et al., 2011; Buschman et al., 2012; Akam and Kullmann, 2014; Voytek et al., 2015; Helfrich and Knight, 2016; Naud and Sprekeler, 2018). Hence, the frontal HMM state plausibly exerts top-down control via multiplexed oscillatory coding such that the frontal cortex selectively modulates coupling with other brain areas in specific frequency bands. This multiplexed information transmission in turn provides a substrate for information flow across the entire brain resulting in a densely connected manifold.

The manifold for the high-beta/gamma state displayed significant connectivity between precentral gyrus and postcentral gyrus ipsilateral to tactile stimulation which were embedded between multiple significantly connected frontal regions (Figure 4.3 D). We also found the IP-r connected via two frontal regions to the left pre and postcentral gyri (supplementary Figure 4.6 B). The spectral properties and the connectivity based on the manifolds are in line with previous findings (Tallon-Baudry et al., 1998, Spitzer et al., 2014b, a; Marco-Pallarés et al., 2015; Kornblith et al., 2016; Spitzer and Haegens, 2017). These studies highlight the relevance of high-beta and gamma activity within the fronto-parietal network for recruiting task-relevant regions.

4.3.3 Differential network transitions relate to varying percept

We will refer to the combined spectral and connectivity manifold-based description of an HMM state as a brain network or state interchangeably. Each of the brain networks described in the previous section seems to have a distinct functional role: The alpha-parietal state is relevant to gate cortical processing, the frontal state plausibly aids in whole-brain information transfer whereas the high-beta/gamma and the cross-brain state utilize top-down control to activate brain regions relevant for tactile perception. Hence, we hypothesized that interactions between these networks in the pre-stimulus period would relate to trial-by-trial perceptual variability. In order to test the relationship between trial-by-trial perceptual

variability and network transitions we inferred HMM dynamics on a single trial level and tested if transition probabilities both to and from a given network, including self-transitions, were significantly different between the two percepts at the critical SOA (see Figure 4.4).




Figure 4.4: A) Shows significantly increased probability to transition between states. Green arrows represent transitions that increased significantly when participants perceived one pulse. B) Orange arrows represent transitions that increased significantly when participants perceived two pulses. $P < 0.05$. The transitions that have not been shown exist in the model but are unchanged between the one or two pulse perception.

For the percept of two stimuli compared to one the transition probability increased from the alpha-parietal state to the cross-brain state, the frontal state, and the high-beta gamma state as well as the transition probability from the high-beta/gamma state to the cross-brain state, the frontal state, and the alpha-parietal state. At the same time transitions from the frontal state to both the alpha-parietal and the high-beta/gamma state increased for the percept as one stimulus. Interestingly, the state transition between the high-beta/gamma and alpha-parietal state was the only pair where the transition probability between the two increased in both directions for the percept of two stimuli. Overall, for the percept of two stimuli states transitioned away from the high-beta/gamma state as well as the alpha-parietal state and more towards the frontal and cross-brain states.

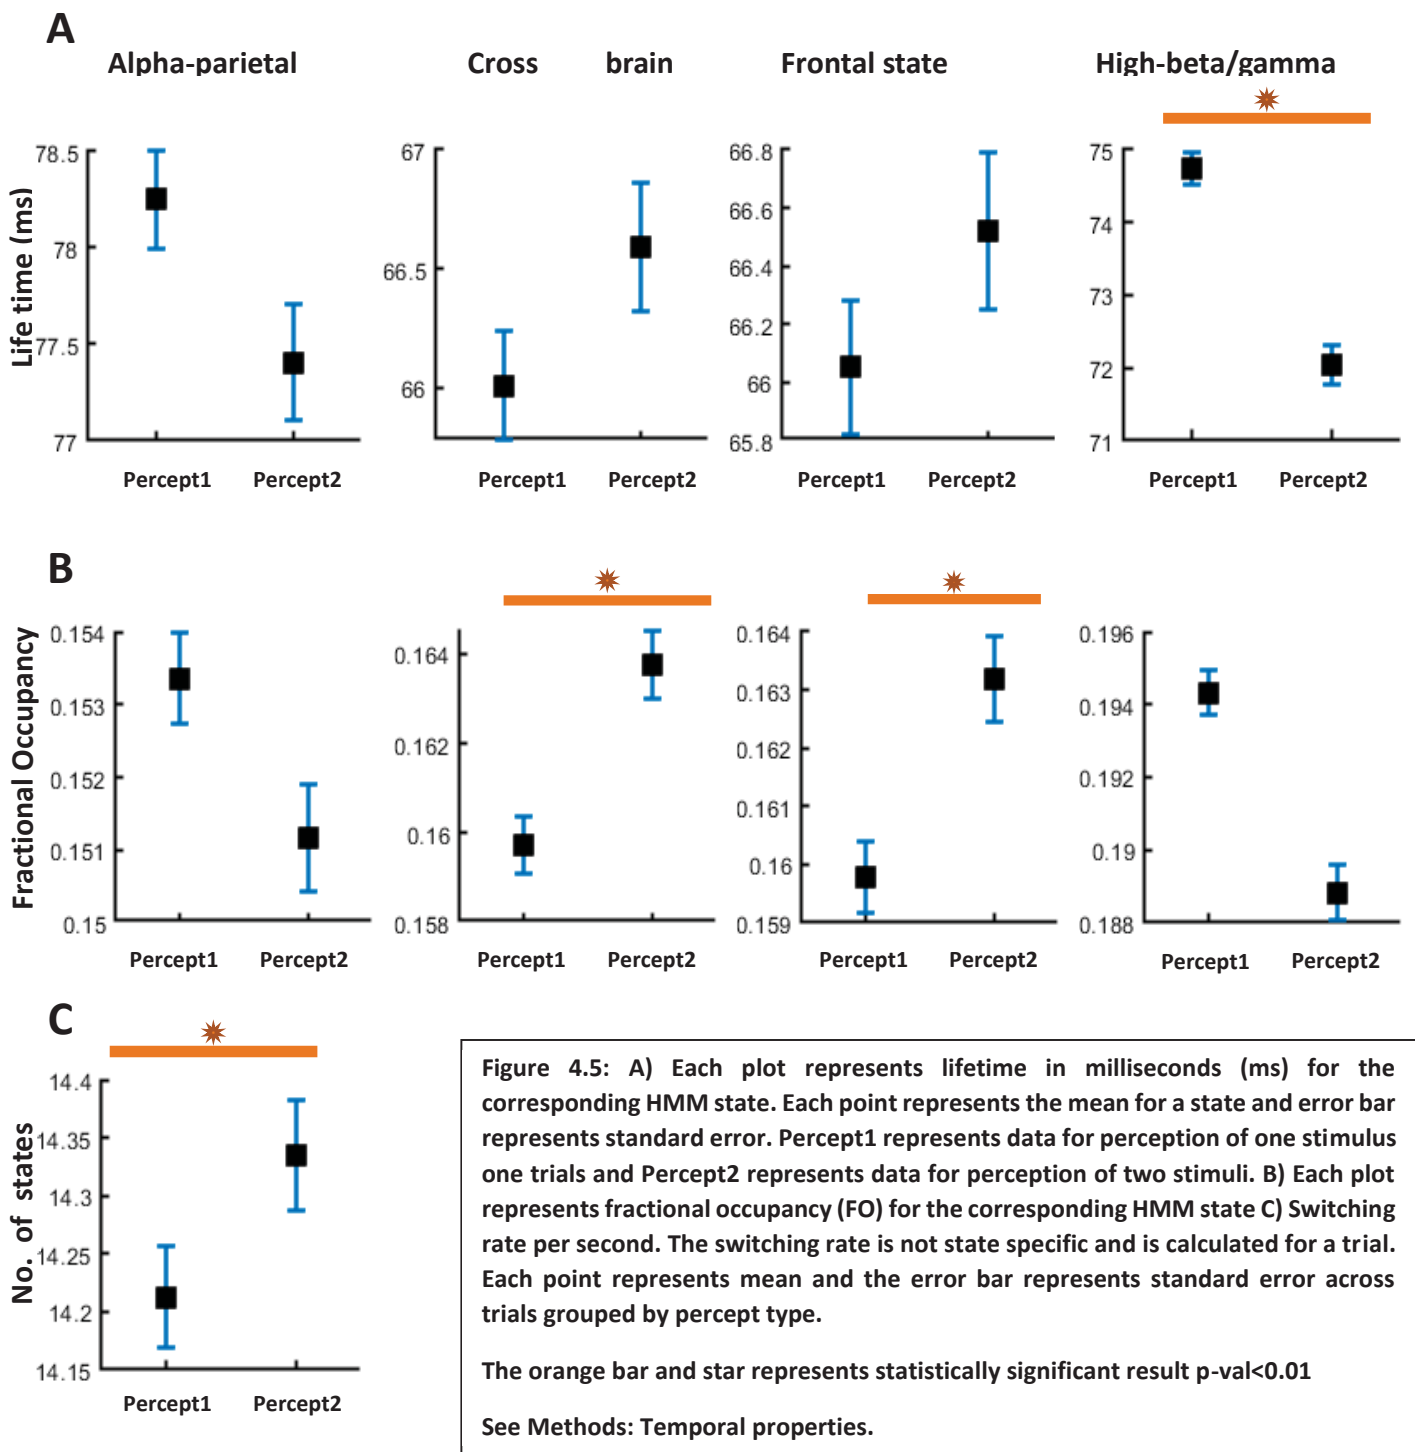
The transition results require a more refined interpretation of the HMM states. The spectrally specific states are plausibly functionally constrained whereas the broadband states either distribute information (frontal state) or assimilate information (cross-brain state) from different constrained networks. Hence, the cross-brain state would require broadband coherence to synchronize incoming information from different networks. Finally, information

assimilation might reconfigure and constrain the connectivity manifold of the cross-brain state to perceptually relevant regions. Therefore, for the percept of two stimuli it is relevant that transitions are balanced between states with functionally constrained connectivity (alpha-parietal, high-beta gamma) and subsequently transition to broadband distributive or assimilative states (frontal, cross-brain).

4.3.4 Faster transitions lead to correct tactile perception

During the pre-stimulus period, some state transitions have a higher likelihood of perceiving one or two stimuli – which already indicates that whole-brain network interactions are relevant for the percept of these ambiguous stimuli. Besides the transition probability between networks, it is also important to analyze the temporal properties of the different networks which could be further factors influencing the percept. Network-specific temporal properties analyzed include the lifetime of a network (duration spent in a network during a visit), and fractional occupancy (the average network probability across time for a trial). For each trial we also calculated a switching rate, which refers to the number of times network transitions happen per second within a trial.

In terms of lifetime, the alpha-parietal state had the highest lifetime for both percepts (alpha-parietal > all other states; $p < 0.001$) but its lifetime was unchanged between the two percepts ($p = 0.59$). For the high-beta/gamma state the lifetime and fractional occupancy significantly decreased for the percept of two stimuli compared to one ($p < 0.001$). For both the cross-brain and the frontal state FO increased for the percept of two stimuli ($p < 0.001$). In addition, the overall rate of switching between different networks was significantly increased for the percept of two stimuli (Fig 5 D). Therefore, during the pre-stimulus period for the percept of two stimuli the high-beta/gamma state should be occupied for shorter durations while the frontal and cross-brain states should have a higher temporal occupation. Finally, network-specific temporal changes should be accompanied by faster overall inter-network transitions.



4.4 Discussion

Perception of an incoming stimulus in an uncertain, non-cued environment requires pre-stimulus spontaneous brain activity to set up dynamics such that the brain can optimally make use of the incoming stimuli to reach a decision. In our work, we demonstrate how these dynamics play out in terms of transitions between functionally distinct whole-brain networks: Pre-stimulus activity in a tactile perception task was characterized by spontaneous switching between four brain networks. The duration of those brain networks being active and the

switching between networks was flexibly adapted and led to the percept of either one or two stimuli. More specifically for the percept of two stimuli as compared to one, the transitions from the alpha-parietal network and high-beta/gamma network increased to the frontal network as well as the cross-brain network and reciprocal transitions decreased. Finally, faster inter-network switching combined with shortened activation (lifetime) of the high-beta/gamma network lead to the percept of two stimuli. Therefore, for optimal perception a balance between cortical gating networks and activation of task-relevant brain regions, as well as improved information dissemination by the frontal network and assimilation of network activity by the cross-brain state in the pre-stimulus period, are essential.

4.4.1 Connectivity manifolds for network descriptions

For functionally characterizing the HMM states we here used both coherence as well as distance metrics on the diffusion map, which yields a manifold of each network based on data-driven network extraction. This is a critical difference compared to previous studies: Commonly functional relevance is ascribed to spectral changes and coherence by correlating them with a behavioral outcome. The spectral approach is based on the notion that coherence serves as a mechanism of neural communication between different brain regions (Fries, 2015). However, Schneider and colleagues demonstrated that coherence is an outcome of neural communication and not its mechanism (Schneider et al., 2021). Thus, coherence between two brain regions can exist even when there is no relationship of information transfer between the sender and the receiver. Furthermore, the calculation of coherence is itself based on the concept of neural oscillations which recently have been questioned regarding their mechanistic role in the brain (Singer, 2018). It has been suggested that to resolve ambiguity regarding the functional role of neural oscillations they should be interpreted in the context of neural circuits producing them (Ibarra-Lecue I et al., 2022). Hence considering a collection of coherent brain regions as a network and correlating them to behavior is not sufficient to establish their functional roles. Also, when calculating coherence in a specific portion of the spectrum, such as 13-30Hz, one cannot claim that the two brain regions are exclusively using this spectrum for coherence-based communication. In view of these limitations concerning neural oscillations and spectral coherence, we here obtained for the HMM states the connectivity manifold based on diffusion maps, which provides a further complementary network description of the HMM states. This approach allowed us to disentangle for the alpha-parietal state pair-wise alpha band coherent interactions into a sparse connectivity manifold. Similarly, a densely connected manifold was not expected based on the broadband activity of the frontal state. The connectivity manifolds were also critical to distinguish the relevance of broadband spectral coherence in the cross-brain and frontal states. Still, the connectivity manifolds do not provide a circuit-level description of spectral activity. However, because of its independence of spectral approaches, it provides a complementary framework within which we could connect whole-brain networks and spectral activity with perception. The connectivity manifolds might prove useful tools to understand the complex properties of connectivity in the brain.

4.4.2 Transitions between higher-order cognitive networks for perception

Before perceiving two stimuli the brain networks switched to the frontal and cross-brain states and away from the high-beta/gamma state. On the other hand, when perceiving only one stimulus brain activity transitioned to the alpha-parietal state and/or was trapped in the high-beta/gamma state. For the percept of two stimuli, neither the high-beta/gamma state nor the alpha-parietal state alone were the relevant networks. Moreover, only when brain activity reciprocally switched between these two states and these states transitioned to the frontal and/or cross-brain states were the participants able to perceive two stimuli. This is in line with the spectral activity of the frontal state reflecting computation within the anterior higher-order cognitive regions and its ability to transmit information based on its connectivity manifold. Hence, HMM states responsible for tracking different aspects of the task should send their locally computed activity to the frontal state for communication. However, if the frontal state transitions to the alpha-parietal state it would lead to excessive inhibitory pruning leading to a loss in information transfer causing the percept to be one pulse. Similarly, transitions from the frontal to the functionally constrained high-beta/gamma state would be detrimental to perception.

The cross-brain state had broadband spectral activity. Moreover, unlike the frontal state, broadband activity was not restricted spatially in the brain. Yet the manifold for the cross-brain state was sparsely connected. In addition, the transition from the alpha-parietal state to the cross-brain state was helpful for the perception of two stimuli. Given the spectral characteristics, transition patterns, and IP-r – motor cortex connectivity we consider the cross-brain state to represent a reconfigurable network, transitions to which are critical for assimilating information from different functionally constrained networks. For a reconfigurable network, the spatially non-specific broadband activity becomes important enabling it to incorporate modulations brought through various other network's information. Plausibly, these reconfigurations are coordinated by parietal cortex activity, especially the IP-r, which previous research has implicated to be involved in such actions in the human brain (Husain and Nachev, 2007; Royal et al., 2016; Singh-Curry and Husain, 2009). Therefore, we consider one spectrally broadband state (frontal) to distribute information whereas the other broadband state (cross-brain state) to incorporate information.

For the high-beta/gamma state the left post- and precentral regions were embedded within a network of prefrontal areas. This reflects a conglomeration of previously reported results in literature: In the sensorimotor system, beta activity is thought to provide inhibitory roles (Donner et al., 2009; Pogosyan et al., 2009) while in the prefrontal regions beta activity tends to indicate reactivation of neuronal ensembles for further processing (Haegens et al., 2011b; Miller et al., 2018; Spitzer and Haegens, 2017). Of note is the presence of the right IP-r in the cluster of prefrontal and left sensory-motor regions. It is known that IP-r in the human brain is involved in detecting salient stimuli within sequences as well as controlling attention over time (Royal et al., 2016; Singh-Curry and Husain, 2009). Hence, the high-beta/gamma state

includes an ensemble of fronto-parietal cortical regions that plausibly aid in activating perceptually relevant sensory-motor regions. Further functional insights about the high-beta/gamma state can be obtained from the altered transition patterns of this state for the percept of two stimuli, which includes increased reciprocal transitions between itself and the alpha-parietal state as well as increased transitions to the distributive frontal state. The high-beta/gamma state is also the only state that despite showing the largest FO across both responses reduces its lifetime for the percept of 2 stimuli. This highlights the critical role of duration when it comes to the action of high-beta/gamma activity in the brain.

The temporal properties of the alpha-parietal state did not significantly change and the lifetime of the alpha-parietal state was the longest amongst all states irrespective of the response by the participants. This at first seems to be in contrast with sustained increases or decreases in alpha activity depending on task performance within pre-stimulus epochs (Haegens et al., 2011b, 2011a; Thut et al., 2006; van Kerkoerle et al., 2014). However, depending on the percept the transitions to and from the alpha-parietal state changed. These results indicate that alpha activity like high-beta/gamma activity exerts its effect on brain networks in short epochs. Furthermore, based on the connectivity manifold of the alpha-parietal state the effect of alpha on the whole-brain network level is to restrict connectivity to task-specific brain regions. Thus, for optimal tactile perception, the brain does not suppress alpha activity but instead balances it by incorporating transitions to different whole-brain networks with specific spectral properties and network connectivity that help in perception. Therefore, the classical framework of alterations in alpha activity to accomplish perceptual tasks needs to be updated to a transition-based framework.

To summarize we demonstrate a nuanced role of spectrally specific coherence in cognitive brain networks. We distinguished networks characterized by broadband spectral coherence in terms of distributive and assimilative roles as demonstrated by their connectivity manifolds. A balance between functionally constrained and spectrally specific states is important in the pre-stimulus period for perception. Furthermore, the transitions to the spectrally broadband networks should be considered equally, if not more, relevant for optimal perception of tactile stimuli.

4.4.3 Limitations of the study

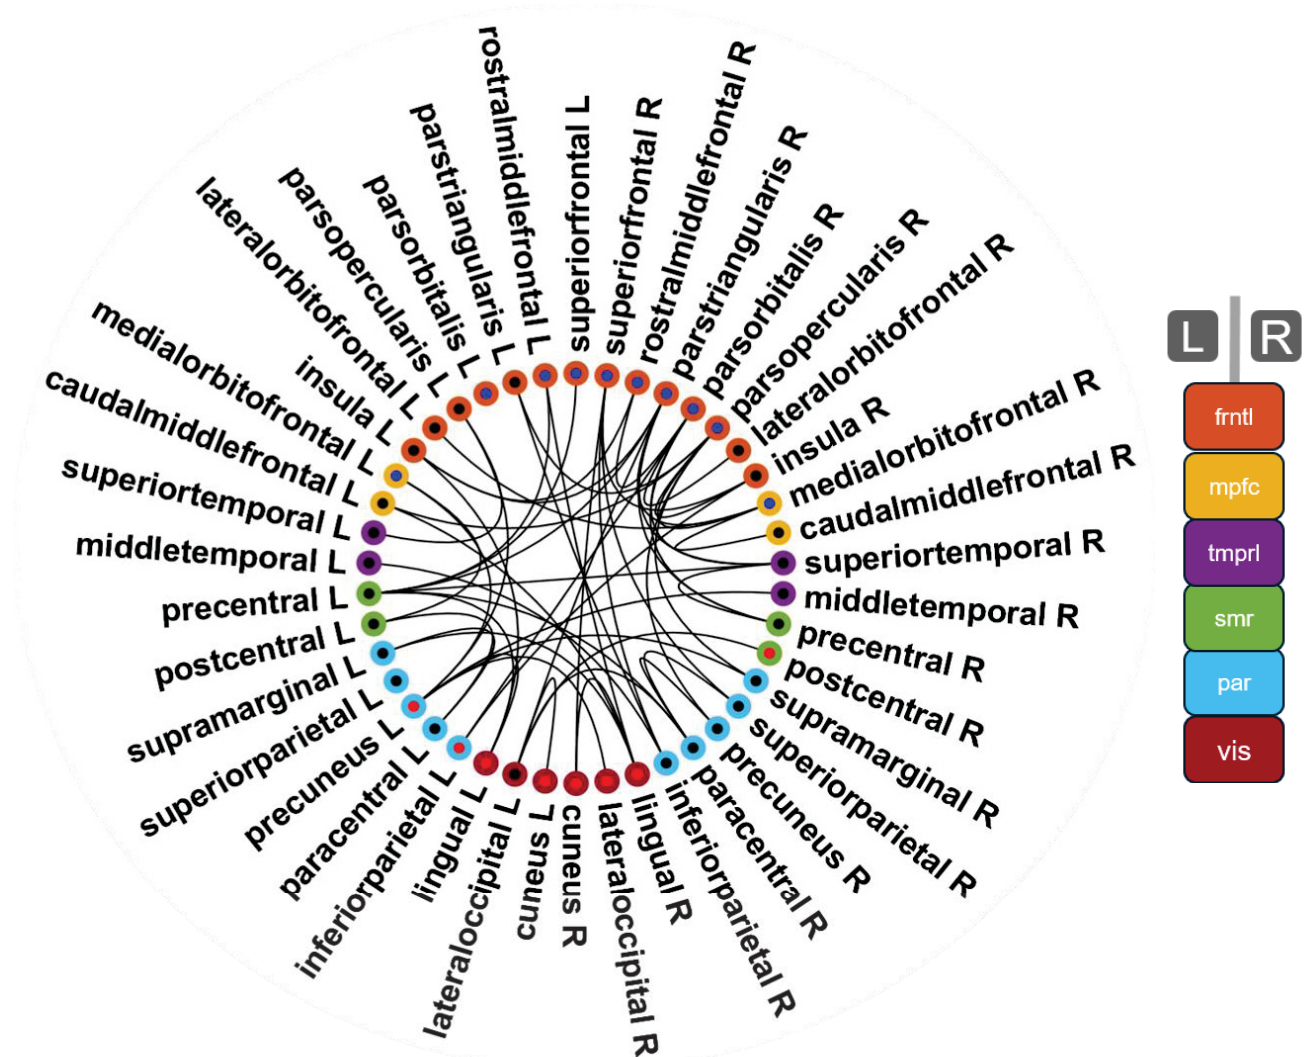
HMM is an unsupervised method, completely blind to the underlying biology of the time series it is fit on. For any HMM model, the analysis of the states helps assess the characteristics of the HMM states. Hence, the connection between biology and the HMM states is post-hoc and the algorithm is not built to extract physiologically interpretable features. For the posthoc analysis, we chose spectral coherence and diffusion maps to produce connectivity manifolds. Manifolds provide an independent view of connectivity from spectral coherence, but they do not provide further information regarding circuit mechanisms of spectral coherence. Furthermore, we are unable to answer the question of whether the connectivity observed on the manifold is responsible for producing the spectral qualities or vice-versa. This remains a crucial question for future work.

Our main result concerns the transitions between different networks. Given the spectral and connectivity manifold characteristics we are able to make statements about the networks' purported functionality. But due to a lack of measures, such as pupil diameter or amount of distractor suppression we are unable to provide concrete evidence regarding the functional roles of the time-resolved networks in the pre-stimulus epoch. Finally, neurophysiologically we do not know what governs the transitions between the networks. Better characterization of the transitions themselves in terms of interpretable neuroscientific properties remains a key challenge in future studies.

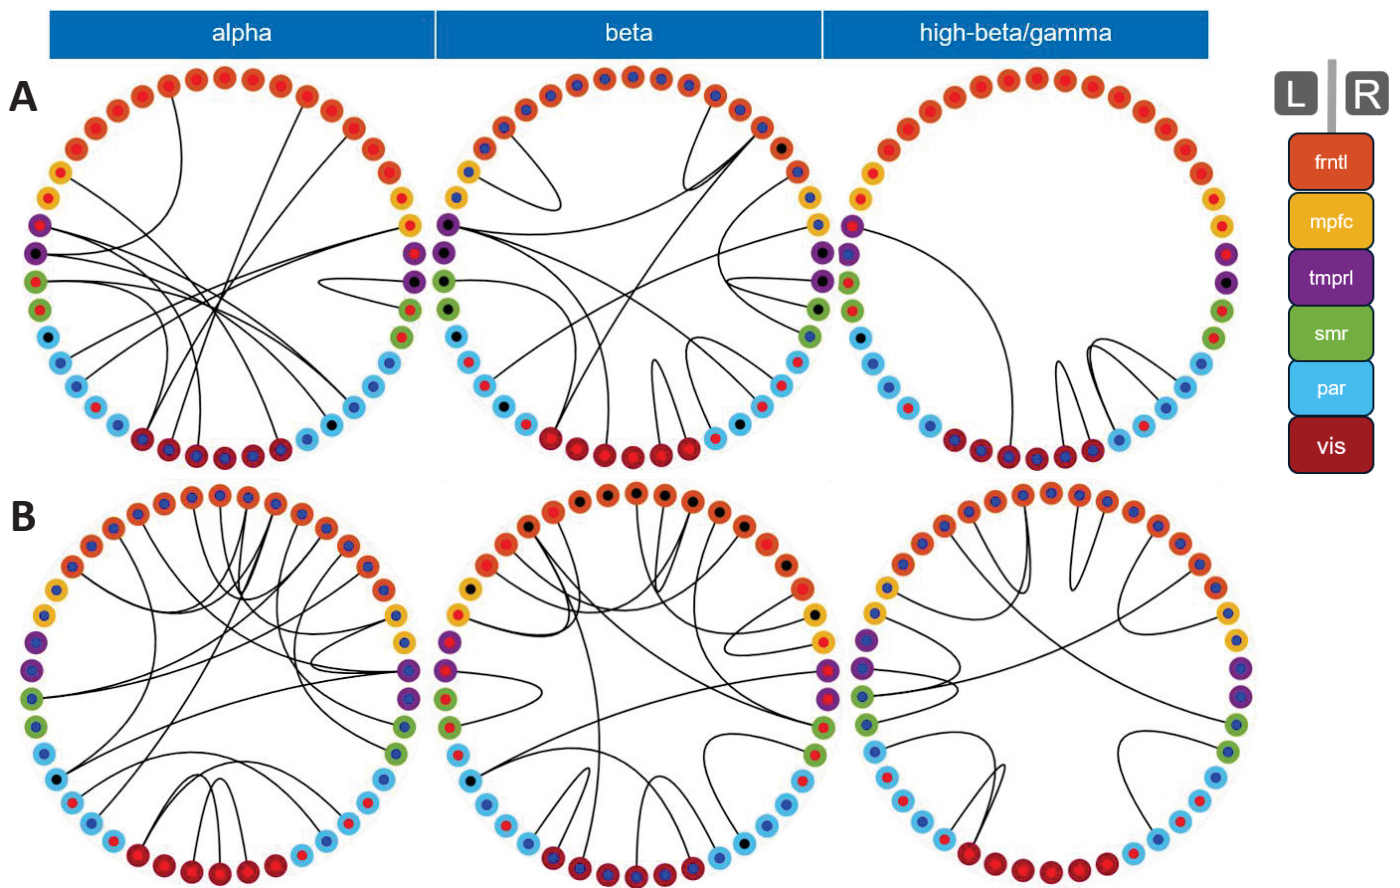
4.5 Conclusion

We here establish a link between spontaneous whole-brain dynamics between multiple networks and the percept of an ambiguous tactile stimulus in a non-informative environment. Faster transitions between cognitive states are necessary to properly perceive incoming stimuli. Hence, temporal patterning of spectrally specific and broadband cognitive states is critical for behavior. Furthermore, we show, to our knowledge for the first time, whole-brain evidence for how pair-wise spectral coherence between brain areas actually manifests on a connectivity manifold between multiple brain regions. This enhances our perspective on the role of oscillatory activity on the cortical level and its relationship to behavior.

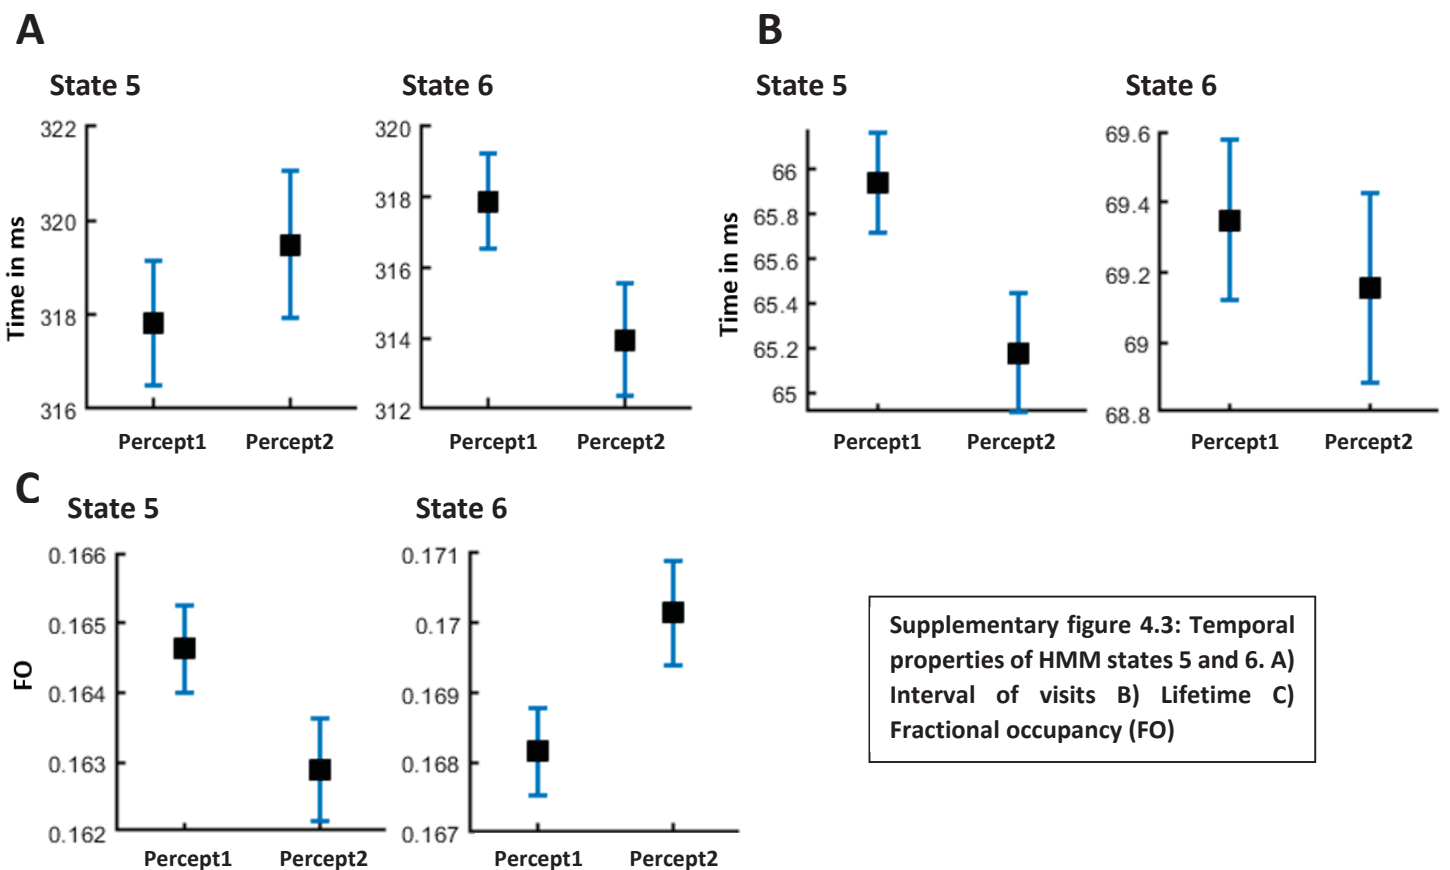
4.6 Supplementary figures




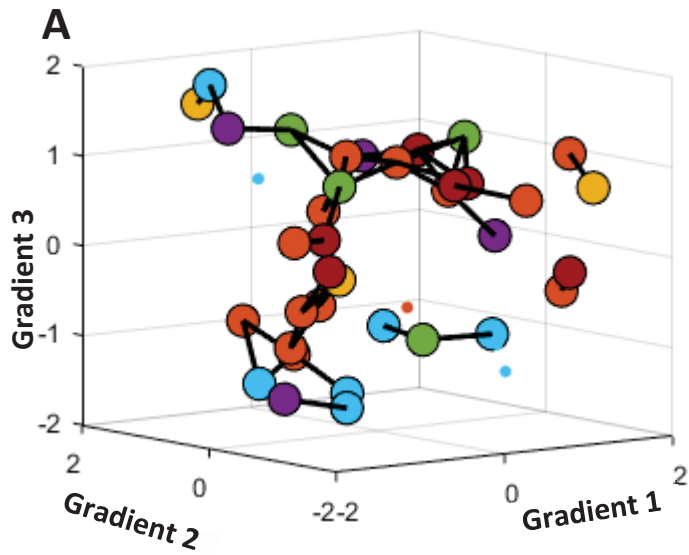
Supplementary figure 4.1: Brain region labels for the ring figures displayed in figures 4.2-4.6. The brain regions are based on the Mindboggle atlas.




Supplementary figure 4.2 A and B: Coherence ring figures for HMM states 5 and 6



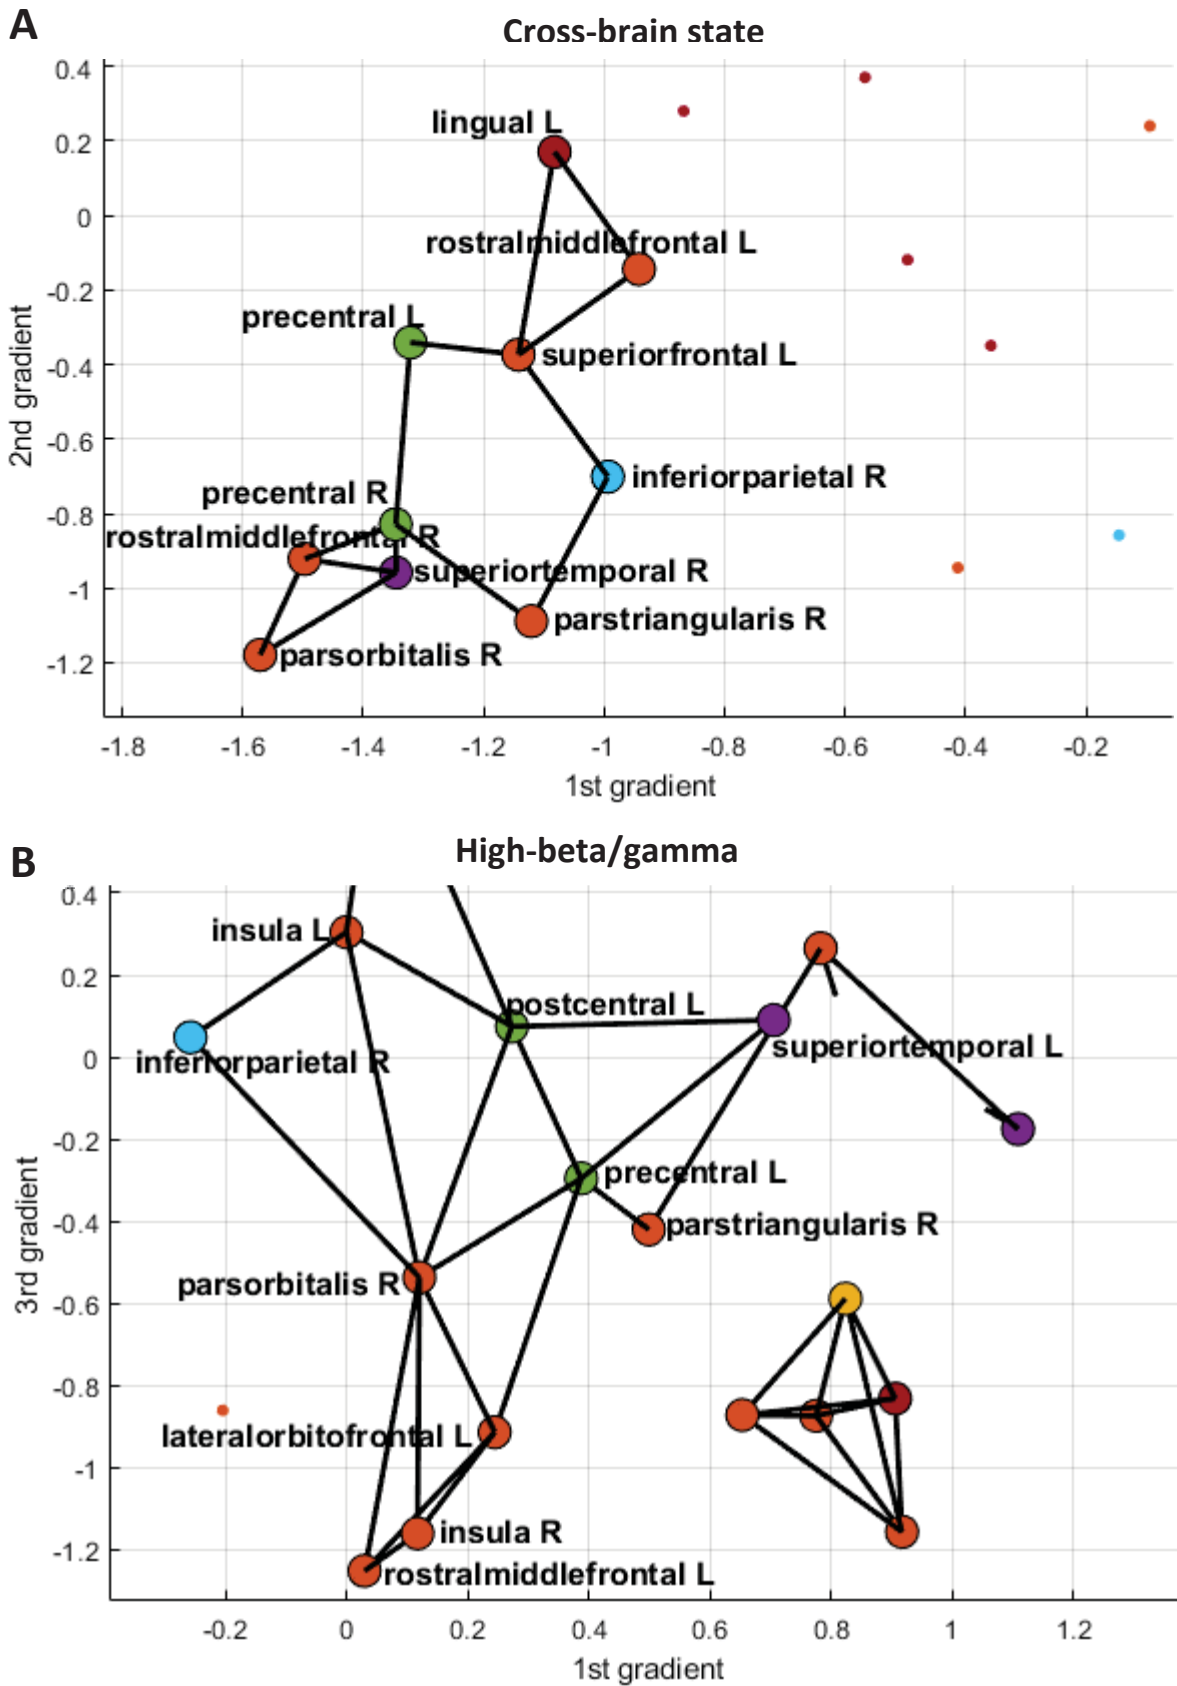
Supplementary figure 4.3: Temporal properties of HMM states 5 and 6. A) Interval of visits B) Lifetime C) Fractional occupancy (FO)



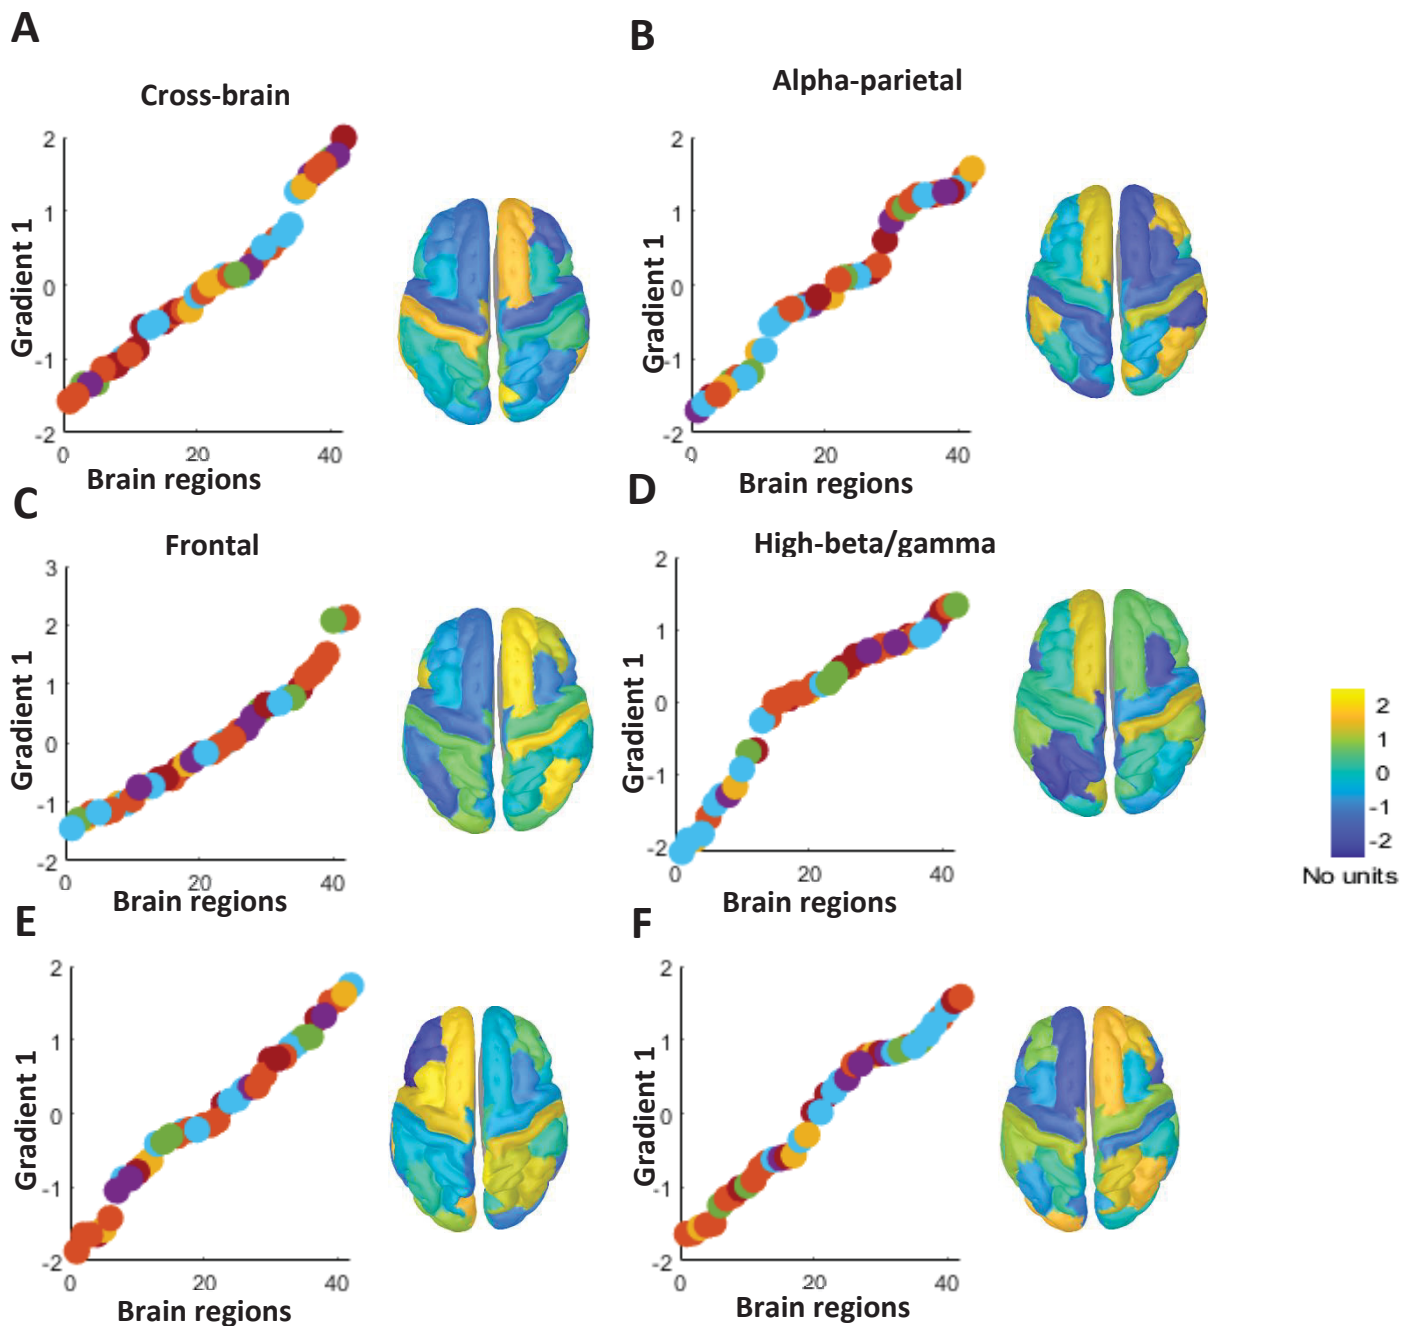
Supplementary figure 4.4: A) HMM state 5 B) HMM



Supplementary figure 4.5: Each plot represents intervals of visits (IOVs) in milliseconds (ms) for the corresponding HMM state.



Supplementary figure 4.6: Zoomed version of the connectivity manifold. A) Cross-brain state B) High-beta/gamma state. In both the figures the inferior parietal R is connected to sensory motor regions via two frontal areas.



Supplementary figure 4.7: Principal gradient representation of HMM states. A) Parietal alpha state B) Frontal state C) Cross brain state D) High-beta/gamma state. E) State 5. F) State 6. Each plot contains two parts. The first part is a plot where coloured nodes represent brain regions from the Mindboggle atlas and follow the same color scheme as previous figures. Here the brain regions have been reorganised on the basis of their value on the principal gradient (gradient 1, see Methods: Diffusion map analysis). Part two of the plot is the representation of the principal gradient (gradient 1) value for Mindboggle atlas regions projected on the standard MNI 152 brain.

Supplementary table 4.1

The tables contain the summary of ANOVA results for tests described in the section Methods: Uniqueness of the state specific embedding.

1A) Eigenvector 1

State	F(df1,df2)	F-value	p-value
Cross-brain	594,5	19492.4	1e-15
Alpha-parietal	594,5	2049.5	1e-15
Frontal	594,5	38733.8	1e-15
High-beta/gamma	594,5	3154.2	1e-15
State 5	594,5	27312.2	1e-15
State 6	594,5	11568.8	1e-15

1B) Eigenvector 2

State	F(df1,df2)	F-value	p-value
Cross-brain	594,5	10889.8	1e-15
Alpha-parietal	594,5	4006.7	1e-15
Frontal	594,5	43415.6	1e-15
High-beta/gamma	594,5	5235.0	1e-15
State 5	594,5	15358.4	1e-15
State 6	594,5	7667.7	1e-15

1C) Eigenvector 3

State	F(df1,df2)	F-value	p-value
Cross-brain	594,5	7147.3	1e-15
Alpha-parietal	594,5	847.57	1e-15
Frontal	594,5	18951.9	1e-15
High-beta/gamma	594,5	1749.3	1e-15
State 5	594,5	8740.5	1e-15
State 6	594,5	5600.9	1e-15

5 Thesis discussion

The aim of the thesis was to explore how UML can be used to unravel the properties of spontaneous large-scale e-phys human brain signals. We emphasized how the biophysical qualities of the neurons and the characteristics of inter-regional connections together affect the statistical aspects of extracellular brain potentials as a justification for the use of unsupervised learning. Through our publications, we showed how UML could be used to show how resting-state brain potentials in different situations have varied spatial, temporal, and spectral properties.

5.1 Summary of publications

5.1.1 Comparison of EEG and fMRI based resting brain networks in humans

In our first publication, we extracted resting-state brain networks based on EEG datasets. Functional magnetic resonance imaging, or fMRI, has been used as the gold standard to characterize resting-state networks. An intensely researched question in the field of human electrophysiology for resting state networks has been regarding the frequency specificity of resting state functional networks. E-phys methods are suitable to address the question of frequency specificity due to their high temporal resolution. ICA and MEG-PAC were applied to the EEG dataset, and the extracted EEG networks were compared to fMRI networks from the same set of subjects. We found that EEG was able to recover resting-state networks comparable to fMRI. As part of our pipeline, we were able to find the optimal filter setting required to extract the Hilbert envelope from the EEG dataset. When coupled to ICA, we found narrow-band finite impulse response filters to be optimal for extracting canonical resting state networks. We were also able to show that ICA was less susceptible to artifact noise than MEG-PAC. In order to establish the frequency specificity of our results and calculate inferential statistics, we used a bootstrapping approach for our ICA results and found that resting state networks did not show frequency specificity.

There were several challenges that we encountered in our study. We previously mentioned that the different patterns that emerge after unsupervised analysis of human e-phys datasets are open to varied interpretations. When we applied ICA to extract spatial components from EEG data, we had to rely on comparison to existing fMRI brain networks as well as visual inspection to confirm that the findings from the EEG dataset were not erroneous. Out of the 20 independent components we extracted from the EEG data, we could utilize only 7 as "brain networks." Normally, in dimensionality reduction applications, the unused components are often discarded as external "noise." In our case, we could not discard the rest of the components as noise since ICA was applied post-artefact cleaning. Artefact cleaning removes any external noise, such as eye blinks, heartbeats, and signal distortion caused by participant head movement. We also could not interpret the remaining components as physiologically relevant due to a) the absence of task behavior and b) no interpretation under the given physiological framework. Furthermore, we cannot rank the components according to any

information criterion, unlike in principal component analysis. In principal component analysis, components can be ranked based on the amount of variance explained by each component. The amount of variance can be calculated based on the eigenvalues of the extracted eigenvectors that span the principal component space. However, in ICA, we cannot measure the amount of variance explained by a component because that is not the objective of the underlying optimization used to find the independent components. Furthermore, we cannot rank the independent components by the amount of "independence" that a component captures. ICA is also criticized for the randomness of its results. The randomness is not a shortcoming of the algorithm but a necessity for the optimization procedure and is a common strategy in many MLPA algorithms. Furthermore, while randomization for optimization does not necessarily translate into missing components, it does lead to a shuffle of component identities. Therefore, the first independent component on one iteration of ICA would not necessarily be the same as the first component on the next ICA iteration. Instead of ranking order, we can cluster the ICA outputs on different iterations to find the most stable components (Himberg et al., 2004). The clustering approaches are similar to the bootstrapping approach we applied to establish frequency specificity in our publication. However, even if we account for the inherent randomness of the optimization algorithm (Hyvarinen, 1999) used by ICA, interpreting the physiological significance of the independent components remains a challenge. Next, as with any MLPA, we adjust hyperparameters in ICA. For the FastICA algorithm used in our work, this includes the number of components, initialization scheme, and non-linearity. In our opinion, nonlinearity is the most critical parameter since it is directly related to the measure of non-Gaussianity that the FastICA algorithm tries to maximize in order to obtain independent components (Hyvärinen and Oja, 2000). FastICA is an iterative algorithm. For example, when it tries to find a single independent component it performs the optimisation to determine the direction for a weight vector that maximizes the projection's $w^T X$ non-Gaussianity, where $X \in \mathbb{R}^{N \times T}$ is the time series matrix fed to the ICA algorithm. In FastICA negentropy is used as a measure of non-Gaussianity. The projection $w^T X$ is passed through a non-linearity to assess the negentropy of the projection. Based on the maximum entropy principle, two nonlinear functions have been suggested to approximate negentropy (Hyvärinen, 1997). The two nonlinear functions suggested are $G1(u) = \log(\cosh(a1*u))/a1$ and $G2(u) = -\exp(-u^2/2)$ (Hyvärinen and Oja, 2000). The functions are also often referred to as "contrast functions" in ICA algorithms. Here $G1$ and $G2$ are arbitrary function names and u is the projection input $w^T X$. For our implementation, we used $G1(u)$, where $a1$ is a constant between 1 and 2. We used $G1(u)$ as it is recommended to be a good general-purpose contrast function (Hyvarinen, 1999). Furthermore, while super or sub Gaussian independent components can be a criterion to choose different contrast functions, we did not explicitly test for them in our dataset (Hyvarinen, 1999). There is no existing study to suggest which nonlinearity should be used to obtain a more physiologically relevant ICA output. Finally, in our study, MEG-PAC, a neurophysiologically grounded method, could not be used due to elevated noise levels in the EEG dataset. We conclude that EEG can be used to reliably extract resting state networks from human recordings, but the spatial resolution

of the extracted EEG networks is lower than that of comparable fMRI networks from the same set of subjects.

5.1.2 Neuromodulation of spontaneous cortico–subthalamic activity in Parkinson’s disease

In the second publication (Sharma et al., 2021), we used a time delay embedded (TDE) hidden Markov model (Vidaurre et al., 2016) to parse the differential effects of dopamine on resting state cortico-STN connectivity. We characterized spatial, temporal, and spectral features of whole brain resting state networks on and off dopamine (on and off medication conditions, respectively). Each HMM state accounted for a unique effect dopamine could have on resting-state networks. Out of the six HMM states, we could interpret three within the current pharmacological framework of dopamine action in Parkinson’s disease. Dopamine’s differential effect allowed us to distinguish between pathological and physiological spectral connectivity. The cortico-cortical state provided e-phys evidence in the delta-theta band for dopaminergic medication overdose in Parkinson’s disease. Prior to our study, evidence was obtained solely through task-based or functional magnetic resonance imaging (fMRI) studies (Cools et al., 2002; MacDonald et al., 2011; Ray and Strafella, 2010). On medication, the cortico-STN state revealed concurrent cortico-cortical and STN-STN interactions as well as spectrally and spatially specific cortico-STN interactions. Furthermore, on medication, a fronto-parietal motor network was present, indicating a shift from STN-mediated motor connectivity to cortical motor connectivity. These findings have not been reported in previous studies. The local STN-STN delta oscillations could only be slightly altered by dopaminergic treatment in the STN-STN condition. Additionally, under the influence of dopaminergic medication, our study demonstrated significant changes in the temporal characteristics of the connection profiles, including lifespan and fractional occupancy. Future modifications to drug and DBS-based therapy techniques may depend on this understanding.

Consistent with our arguments about the difficulty of interpreting the outputs of UML methods, we could only interpret three of the six HMM states within a physiological framework. As with the ICA components, we could not rule out the remaining states as random noise. Also, we saw that adding or taking away PD patients, whose DBS electrodes were mistakenly implanted outside of the planned anatomical target, changed the model, even though the results of several optimization iterations of the HMM model fitting were consistent.

A critical element in our pipeline was delay-embedding the time series. Technically, delay embedding is done to recover the underlying phase space in which the dynamics of a system emerge (Packard et al., 1980). We embedded the time series for our HMM model to bring our approach closer to the realm of Gaussian processes (Rasmussen, 2004). Using delay embedding, we could define covariance simultaneously over spatial and temporal dimensions. The multivariate Gaussian covariance was then used to define the HMM emission distributions. Yet despite delay embedding, we were restricted to Gaussian emission

distributions that could have limited the complexity of the structure of our dataset model. Another limitation of our method was that we used separate HMM models to characterize medication-on and medication-off datasets. The two-model approach was primarily adopted due to the heterogeneity of the patient's response to medication. However, using two models warranted a post-hoc analysis to find correspondence between the states of the two HMM models. In future work, we would like to train a single HMM with an increased number of states to see if they can capture different features of the dataset.

5.1.3 Transitions between latent cognitive networks track somatosensory perceptual variability

For the third publication, our HMM pipeline was able to find a significant relationship between state transition probabilities and the behavioral responses of different subjects. In order to better interpret the output from an unsupervised HMM model, we applied diffusion maps and spectral analysis to state-specific data. We found that each HMM state captured spectrally-specific whole brain network activity. With the help of diffusion maps, we identified a functional manifold for states that were associated with endogenously driven cognitive demands. These states also exhibited spatio-spectral properties that have been previously associated with distinct forms of top-down cognition during perceptual tasks (Dijkstra et al., 2017; Long and Kuhl, 2018; Makino, 2019; Witt and Stevens, 2013).

We demonstrated that transitions between different types of cognitive states are essential for accurate perception of incoming inputs. Therefore, behavior depends on the temporal patterning of spectrally distinct cognitive states. Additionally, to our knowledge, we provide the first whole-brain evidence of how pair-wise spectral coherence between brain regions actually manifests on the functional connectivity manifold between various brain regions. This broadens our understanding of the function of oscillatory activity on the cortical level and its connection to behavior. Despite including analyses of transition probabilities, diffusion map analysis, and pairwise spectral coherence, we could only interpret a subset of transitions and HMM states. A shortcoming of our approach was that during the pre-stimulus resting period, we did not record physiological parameters or measures such as pupil diameter, head movement, or hand or leg muscle movement. This limited our ability to interpret the HMM states. We also attempted to have the trained HMM model classify out-of-sample behavioral responses, but we were unsuccessful in our classification approach. We could not use the Viterbi sequence (Viterbi, 1967) of a trial to decode behavioral responses.

In this study, as with our previous HMM modeling attempt, we used multivariate Gaussian distributions to model the HMM states, which limited our ability to model more complex statistical properties of our dataset. Next, we used time delay embedding to fit our HMM model. The delay embedding scheme used in our analysis pipeline used data vectors from both the past and future of the current time point. This prevented us from using pre-stimulus spontaneous activity to predict stimulus responses or changes in neural activity for a future time point. Furthermore, we did not explicitly take into account the response type to control

state switching for the HMM. This may have limited our ability to find a predictive state sequence for a specific response type.

5.1.4 A general note on Hidden Markov Model fitting

An HMM is a probabilistic model that requires optimization of an objective function to find suitable model parameters for the dataset. A critical question regarding any probabilistic model is the selection of the function to optimize. According to Bayes' theorem, we have the likelihood and posterior functions for any probabilistic model that we can choose to optimize (Bilmes, 1998). Usually a maximization approach is used to optimize either the likelihood or the posterior functions. The approach used to optimize the likelihood function is called maximum likelihood estimation (MLE). Two different approaches can be used to optimize the posterior function: maximum a posteriori (MAP) estimation or a fully Bayesian approach. MLE and MAP assume that the model parameters are fixed-point estimates. In contrast, a fully Bayesian approach assumes model parameters to be random variables, and hence the Bayesian approach results in a full probability distribution describing the model parameters. From a machine learning perspective, MLE and MAP approaches lead to an overfit model. Intuitively, if any value fails to occur during the MLE/MAP training phase (parameter estimation phase) of the algorithm, then the corresponding probability will be zero. Furthermore, if this specific value shows up during model deployment, then the MLE/MAP-trained model will not be able to handle it. Hence, we used a fully Bayesian approach to train our models.

In addition to being a probabilistic model, the HMM is a latent variable model. Direct optimization of any probabilistic objective function, likelihood or posterior, is difficult. Therefore, in practice, an iterative alternating scheme for latent variable models is used to optimize parameters for an HMM. The standard algorithm for finding model parameters for many latent variable models is the expectation maximization (EM) algorithm (Dempster et al., 1977). The EM algorithm is a coordinate ascent algorithm that contains two parts. The expectation (E) step calculates the expected value of the complete joint likelihood of both the data and the latent variables with respect to the posterior (latent variable) distribution. The E step uses the current estimate of the model parameters. The M (maximization) step maximizes the model parameters with respect to the complete joint data and likelihood distribution of the latent variables. These two steps are used alternately until the expected value in the E step converges. It is expected that the E step's expectation will continue to increase with every iteration until convergence. An important point about the EM algorithm is that the E step requires explicit knowledge of the posterior distribution of the latent variable. For most probabilistic models, this posterior distribution is not computable analytically. Hence, an alternate approach for model estimation involves using an analytical approximation for the posterior distribution. This alternate approach is called the variational Bayes approximation. We can extend the EM to the variational Bayes method to approximate the model parameters. When applying EM to optimize the variational Bayes objective for parameter estimation, it can also be referred to as "variational EM" or "coordinate ascent

variational inference." Given the analytically intractable latent variable posterior distribution for the HMM, we used the variational Bayes approach.

Finally, a problem with EM (with or without variational equations) is that in order to calculate the E step at every iteration, the full dataset is used for estimation. This means that each iteration of the EM algorithm scales in computational cost with the size of the dataset used. Recently, a method has been developed to perform variational Bayesian inference using a stochastic, gradient-based approach called stochastic variational inference (SVI). In this optimization approach, inference can be performed by computing function gradients to be optimized using fast, noisy approximations from subsampled data. Hence, we do not need to go through the entire dataset for each iteration of parameter updates. For technical details regarding SVI, see Hoffman et al., 2013.

In the case of HMMs, there are two kinds of subsampling when applying SVI. First, each time series in its entirety is considered an individual sample (Johnson and Willsky, 2014). Second, a single time series is subsampled into shorter temporal epochs, and parameters are updated (Foti et al., 2014). Let us take the example of patient data with three 10-minute recording sessions per subject. If we have 10 patients, then our recording paradigm would lead to a dataset containing 30 different time series in total. In the first subsampling scheme, we would randomly select a few complete time series at every iteration of SVI and infer the model parameters. In the second scheme, the SVI algorithm would subsample short temporal segments from all 30 different time series and update model parameters until convergence. We used the first subsampling scheme to fit our models in study numbers two and three. This was done because we did not encounter extremely long or high sampling rate time series, which would have made the first subsampling scheme computationally expensive to apply. It is important to remember that we have to properly initialize the parameters for any stochastic optimization algorithm. We seeded the initialization phase with random parameters and then used a single variational inference pass through the full dataset to set the SVI scheme's initial starting point. Overall, our inference schemes, the objective functions, and the initialization scheme were the ones that have been used in practice for a wide variety of time series or independent sample mixture model problems (Ashwood et al., 2022; Batty et al., 2019; Paun et al., 2022; Regazzoni et al., 2022; Vidaurre et al., 2016).

5.1.5 Contributions of our research

In this thesis, our main objective was to use UML to analyze e-phys signals recorded non-invasively from the human brain. Starting with a static, non-time-resolved method was important for proving that methods like ICA, which are based on statistics but don't take biology into account, can be used to study spontaneous brain networks. We evaluated several filter settings that might be optimally combined with ICA and found that FIR narrow band filters were most successful in recovering canonical brain networks. We were able to reproduce earlier neuroimaging and e-phys findings, providing a suitable beginning point for our brain network research. This led to the next set of time-resolved analyses, which were

done to find out more about the role of spontaneous "resting state" epochs, which can be changed by a neurotransmitter in the brain or by behavior.

The move from a static to a time-resolved method revealed UML's potential for neural time series and neuroscience. We reasoned that different forms of brain-region interactions affecting the time series would result in nonstationarities. Sharma et al. (2021) used the HMM to partition the nonstationarities into distinct HMM states. Each HMM state represented a brain network. We reproduced multiple Parkinson's disease and dopamine-related cortical brain network findings that had been previously proven using time-averaged (static analysis) techniques. In these brain network results, we also showed how frequency-specific activity was a transient, rather than sustained, feature of brain networks. In this study, we presented a novel analysis pipeline for elucidating the complex effects of neuromodulators using statistically grounded approaches.

In our third investigation, we expanded the HMM pipeline's neuroscientific utility in two different ways. First, we used behavioral data to demonstrate how distinct HMM states (brain networks) were related to one another. We were able to demonstrate a correlation between the somatosensory perception of the stimulus and the way transitions between HMM states take place. The link between perception and state transitions justifies a change in perspective from one that stresses sustained frequency-specific rhythms to one that prioritizes transitory dynamical brain networks. These transitory, dynamical brain networks may or may not have spectrally specific properties while carrying out perceptual tasks. Most theories about how important frequency-specific activity is for perception are based on time-averaged analyses that look at rhythmic activity over time and see it as a neural correlate of perception. Hence, our HMM-based time-resolved dynamical results are a significant contribution to the area of MEG studies exploring sensory perception.

Next, we look at how spectral coherence and network connectivity are related by using the HMM's ability to describe states in terms of zero-mean multivariate Gaussian distributions. A large number of studies often characterize brain networks using spectral coherence. After that, conclusions are made about how different coherence-based brain networks work, especially in MEG studies that focus on tasks. However, spectral coherence calculation is based on Fourier analysis, which does not take into account the noise that affects high-dimensional brain signals even after considerable quality control. Also, it has been shown that spectral coherence is not the mechanism itself but rather the result of communication between different parts of the brain (Schneider et al., 2021). So we used diffusion maps as a way to figure out how networks are formed in the brain. Diffusion maps use the diffusion operator to perform a random walk on a connected graph and use the local geometry of the dataset to show how two brain regions are connected (Coifman et al., 2005; Singer and Coifman, 2008), which is something that PCA and ICA don't do. The characteristics of the data that control the diffusion process on the graph are a critical concern when using diffusion maps. It has been demonstrated that if a dataset or time series is nonstationary, the probability distribution will most likely control the diffusion process rather than the data itself

(Carter et al., 2009; Carter et al., 2011; Lian et al., 2015). In our analysis pipeline, we factored the nonstationarity of the time series by first parsing the underlying time series with a probabilistic discrete state space model, like the HMM. Then, to get state-specific connectivity manifolds, we looked at state-specific covariance matrices that were calculated between different brain regions. This pipeline lets us make neuroscientific conclusions from the state manifolds without using spectral coherence analysis. Our novel pipeline, which is made up of diffusion maps and HMM, shows how useful it is to use a non-linear dimensionality reduction method by first addressing the fact that the time series is not stationary. Using only data-driven methods, this HMM-diffusion map combination makes it possible to draw new neuroscientific conclusions from the time series. Lastly, the diffusion map technique brings information geometry into the field of brain connectivity, which has been dominated by Fourier analysis up until now. However, it is important to point out that, at the core of both techniques, we used the covariance between brain regions as the base measure of connectivity between different brain regions. Humans, non-human primates, and rodents can all be recorded using the recording techniques used in neuroscience today. This makes it possible to record large time-series datasets. Until better recording techniques that can capture dynamics at the neurotransmitter or synaptic level are used to study brain activity, covariance is the best estimation we have of neural communication. It will therefore be crucial to keep learning about the statistical characteristics of neural signals and creating techniques that can effectively use these statistics to understand how the brain works.

In conclusion, we showed the potential of UML to shed light on dynamic fluctuations in brain activity and the potential links to neuromodulation and behavior. Future studies will focus on improving HMM models by incorporating better probability distributions to model e-phys datasets (Ashwood et al., 2022; Smith et al., 2021). Furthermore, we would include behavioral parameters as auxiliary variables to improve decodability in HMM models (Hälvä and Hyvarinen, 2020; Hyvarinen et al., 2019). Understanding and exploiting the statistical structure of electrophysiology datasets will lead to better insights into aberrant neural activity and into interventions that can modify activity for improved clinical results. By utilizing UML algorithms as well as electrical and pharmacological perturbations to study human brain function, the field of cognitive and computational neuroscience hopes to answer some of the most pressing questions about how we think, move, and feel.

References

- Akam T, Kullmann DM. 2010. Oscillations and filtering networks support flexible routing of information. *Neuron* **67**:308–320.
- Ariani G, Pruszynski JA, Diedrichsen J. 2022. Motor planning brings human primary somatosensory cortex into action-specific preparatory states. *eLife* **11**:e69517.
- Ashwood ZC, Roy NA, Stone IR, Urai AE, Churchland AK, Pouget A, Pillow JW. 2022. Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience* **25**:201–212.
- Bagi B, Brecht M, Sanguinetti-Scheck JI. 2022. Unsupervised discovery of behaviorally relevant brain states in rats playing hide-and-seek. *Current Biology* **32**:2640-2653.e4.
- Baillet S. 2017. Magnetoencephalography for brain electrophysiology and imaging. *Nature Neuroscience* **20**:327–339.
- Baker A, Kalmbach B, Morishima M, Kim J, Juavinett A, Li N, Dembrow N. 2018. Specialized Subpopulations of Deep-Layer Pyramidal Neurons in the Neocortex: Bridging Cellular Properties to Functional Consequences. *Journal of Neuroscience* **38**:5441–5455.
- Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Probert Smith PJ, Woolrich M. 2014. Fast transient networks in spontaneous human brain activity. *eLife* **3**:e01867.
- Basile GA, Bertino S, Nozais V, Bramanti A, Ciurleo R, Anastasi GP, Milardi D, Cacciola A. 2022. White matter substrates of functional connectivity dynamics in the human brain. *NeuroImage* **258**:119391.
- Batty E, Whiteway M, Saxena S, Biderman D, Abe T, Musall S, Gillis W, Markowitz J, Churchland A, Cunningham JP, Datta SR, Linderman S, Paninski L. 2019. BehaveNet: nonlinear embedding and Bayesian neural decoding of behavioral videos. *Advances in Neural Information Processing Systems*. pp. 15706–15717.
- Baumgarten TJ, Schnitzler A, Lange J. 2015. Beta oscillations define discrete perceptual cycles in the somatosensory domain. *Proceedings of the National Academy of Sciences* **112**:12187–12192.
- Bell PT, Shine JM. 2016. Subcortical contributions to large-scale network communication. *Neuroscience & Biobehavioral Reviews* **71**:313–322.
- Benchenane K, Tiesinga PH, Battaglia FP. 2011. Oscillations in the prefrontal cortex: a gateway to memory and attention. *Current Opinion in Neurobiology* **21**:475–485.
- Bilmes J. 1998. A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models. *CTIT technical reports series*.
- Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS. 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine* **34**:537–541.

Boon LI, Geraedts VJ, Hillebrand A, Tannemaat MR, Contarino MF, Stam CJ, Berendse HW. 2019. A systematic review of MEG-based studies in Parkinson's disease: The motor system and beyond. *Human Brain Mapping* **40**:2827–2848.

Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D. 2020. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*. pp. 1877–1901.

Buschman TJ, Denovellis EL, Diogo C, Bullock D, Miller EK. 2012. Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron* **76**:838–846.

Buschman TJ, Miller EK. 2007. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* **315**:1860–1862.

Buzsáki G. 2006. *Rhythms of the Brain*. Oxford University Press.

Buzsáki G, Anastassiou CA, Koch C. 2012. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience* **13**:407–420.

Buzsáki G, Moser EI. 2013. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience* **16**:130–138.

Cagnan H, Duff EP, Brown P. 2015. The relative phases of basal ganglia activities dynamically shape effective connectivity in Parkinson's disease. *Brain* **138**:1667–1678.

Carter K, Raich R, Finn W, Hero A. 2011. Information-Geometric Dimensionality Reduction. *Signal Processing Magazine IEEE* **28**:89–99.

Carter KM, Raich R, Finn WG, Hero III AO. 2009. FINE: Fisher Information Nonparametric Embedding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **31**:2093–2098.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012. Neural population dynamics during reaching. *Nature* **487**:51–56.

Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. 2010. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* **68**:387–400.

Cifre I, Miller Flores MT, Penalba L, Ochoa JK, Chialvo DR. 2021. Revisiting Nonlinear Functional Brain Co-activations: Directed, Dynamic, and Delayed. *Frontiers in Neuroscience* **15**:700171.

Coifman RR, Lafon S, Lee AB, Maggioni M, Nadler B, Warner F, Zucker SW. 2005. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences* **102**:7426–7431.

Colclough GL, Brookes MJ, Smith SM, Woolrich MW. 2015. A symmetric multivariate leakage correction for MEG connectomes. *Neuroimage* **117**:439–448.

- Cools R. 2006. Dopaminergic modulation of cognitive function-implications for L-DOPA treatment in Parkinson's disease. *Neuroscience Biobehavioral Reviews* **30**:1–23.
- Cools R, Barker RA, Sahakian BJ, Robbins TW. 2001. Enhanced or Impaired Cognitive Function in Parkinson's Disease as a Function of Dopaminergic Medication and Task Demands. *Cerebral Cortex* **11**:1136–1143.
- Cools R, Clark L, Owen AM, Robbins TW. 2002. Defining the Neural Mechanisms of Probabilistic Reversal Learning Using Event-Related Functional Magnetic Resonance Imaging. *Journal of Neuroscience* **22**:4563–4567.
- Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME. 2001. Frequencies Contributing to Functional Connectivity in the Cerebral Cortex in "Resting-state" Data. *American Journal of Neuroradiology* **22**:1326–1333.
- Cruz AV, Mallet N, Magill PJ, Brown P, Auerbach BB. 2009. Effects of dopamine depletion on network entropy in the external globus pallidus. *Journal of Neurophysiology* **102**:1092–1102.
- Dempster AP, Laird NM, Rubin DB. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society Series B (Methodological)* **39**:1–38.
- Diesmann M, Gewaltig M-O, Aertsen A. 1999. Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**:529–533.
- Ecker AS, Berens P, Tolias AS, Bethge M. 2011. The effect of noise correlations in populations of diversely tuned neurons. *Journal of Neuroscience* **31**:14272–14283.
- Engel AK, Fries P. 2010. Beta-band oscillations—signalling the status quo? *Current Opinion in Neurobiology, Cognitive neuroscience* **20**:156–165.
- Fields RD, Stevens-Graham B. 2002. New Insights into Neuron-Glia Communication. *Science* **298**:556–562.
- Finn ES, Huber L, Jangraw DC, Molfese PJ, Bandettini PA. 2019. Layer-dependent activity in human prefrontal cortex during working memory. *Nature Neuroscience* **22**:1687–1695.
- Fleischhauer K. 1974. On different patterns of dendritic bundling in the cerebral cortex of the cat. *Zeitschrift für Anatomie und Entwicklungsgeschichte* **143**:115–126.
- Florin E, Baillet S. 2015. The brain's resting-state activity is shaped by synchronized cross-frequency coupling of neural oscillations. *NeuroImage* **111**:26–35.
- Florin E, Watanabe M, Logothetis NK. 2015. The role of sub-second neural events in spontaneous brain activity. *Current Opinion in Neurobiology* **32**:24–30.
- Förstner W, Moonen B. 2003. A Metric for Covariance Matrices In: Grafarend EW, Krumm FW, Schwarze VS. *Geodesy-The Challenge of the 3rd Millennium*. pp.299–309.
- Foti NJ, Xu J, Laird D, Fox EB. 2014. Stochastic Variational Inference for Hidden Markov Models. *Advances in Neural Information Processing Systems*. pp.3599- 3607

- Foxe JJ, Snyder AC. 2011. The Role of Alpha-Band Brain Oscillations as a Sensory Suppression Mechanism during Selective Attention. *Frontiers in Psychology* **2**:154.
- Fraser AM. 2008. Hidden Markov Models and Dynamical Systems. *USA: Society for Industrial and Applied Mathematics*.
- Fries P. 2015. Rhythms for Cognition: Communication Through Coherence. *Neuron* **88**:220–235.
- Gallivan JP, Chapman CS, Wolpert DM, Flanagan JR. 2018. Decision-making in sensorimotor control. *Nature Reviews Neuroscience* **19**:519–534.
- George JS, Strunk J, Mak-McCully R, Houser M, Poizner H, Aron AR. 2013. Dopaminergic therapy in Parkinson's disease decreases cortical beta band coherence in the resting state and increases cortical beta band power during executive control. *Neuroimage Clinical* **3**:261–270.
- Gerfen CR, Economo MN, Chandrashekar J. 2018. Long Distance Projections of Cortical Pyramidal Neurons. *Journal of Neuroscience Res* **96**:1467–1475.
- Gershman SJ, Uchida N. 2019. Believing in dopamine. *Nature Reviews Neuroscience* **20**:703–714.
- Gertz H, Fiehler K. 2015. Human posterior parietal cortex encodes the movement goal in a pro-/anti-reach task. *Journal of Neurophysiology* **114**:170–183.
- Gidon A, Zolnik TA, Fidzinski P, Bolduan F, Papoutsi A, Poirazi P, Holtkamp M, Vida I, Larkum ME. 2020. Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science* **367**:83–87.
- Gill JV, Lerman GM, Zhao H, Stetler BJ, Rinberg D, Shoham S. 2020. Precise Holographic Manipulation of Olfactory Circuits Reveals Coding Features Determining Perceptual Detection. *Neuron* **108**:382-393.e5.
- Glaser J, Whiteway M, Cunningham JP, Paninski L, Linderman S. 2020. Recurrent Switching Dynamical Systems Models for Multiple Interacting Neural Populations. *Advances in Neural Information Processing Systems*. pp. 14867–14878.
- Gogulski J, Boldt R, Savolainen P, Guzmán-López J, Carlson S, Pertovaara A. 2015. A segregated neural pathway for prefrontal top-down control of tactile discrimination. *Cerebral Cortex* **25**:161–166.
- Golding NL, Spruston N. 1998. Dendritic Sodium Spikes Are Variable Triggers of Axonal Action Potentials in Hippocampal CA1 Pyramidal Neurons. *Neuron* **21**:1189–1200.
- Goris RLT, Movshon JA, Simoncelli EP. 2014. Partitioning neuronal variability. *Nature Neuroscience* **17**:858–865.
- Greicius MD, Supekar K, Menon V, Dougherty RF. 2009. Resting-State Functional Connectivity Reflects Structural Connectivity in the Default Mode Network. *Cerebral Cortex* **19**:72–78.

Gross RE, Krack P, Rodriguez-Oroz MC, Rezai AR, Benabid A-L. 2006. Electrophysiological mapping for the implantation of deep brain stimulators for Parkinson's disease and tremor. *Movement Disorders* **14**:259-283.

Guan S, Jiang R, Bian H, Yuan J, Xu P, Meng C, Biswal B. 2020. The Profiles of Non-stationarity and Non-linearity in the Time Series of Resting-State Brain Networks. *Frontiers in Neuroscience* **14**:493.

Haegens S, Cousijn H, Wallis G, Harrison PJ, Nobre AC. 2014. Inter- and intra-individual variability in alpha peak frequency. *NeuroImage* **92**:46–55.

Haegens S, Händel BF, Jensen O. 2011a. Top-Down Controlled Alpha Band Activity in Somatosensory Areas Determines Behavioral Performance in a Discrimination Task. *Journal of Neuroscience* **31**:5197–5204.

Haegens S, Luther L, Jensen O. 2012. Somatosensory anticipatory alpha activity increases to suppress distracting input. *Journal of Cognitive Neuroscience* **24**:677–685.

Haegens S, Nacher V, Luna R, Romo R, Jensen O. 2011b. α -Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceeding of the National Academy of Sciences* **108**:19377–19382.

Haegens S, Vergara J, Rossi-Pool R, Lemus L, Romo R. 2017. Beta oscillations reflect supramodal information during perceptual judgment. *Proceeding of the National Academy of Sciences* **114**:13810–13815.

Halgren M, Ulbert I, Bastuji H, Fabó D, Erőss L, Rey M, Devinsky O, Doyle WK, Mak-McCully R, Halgren E, Wittner L, Chauvel P, Heit G, Eskandar E, Mandell A, Cash SS. 2019. The generation and propagation of the human alpha rhythm. *Proceedings of the National Academy of Sciences* **116**:23772–23782.

Hälvä H, Hyvarinen A. 2020. Hidden Markov Nonlinear ICA: Unsupervised Learning from Nonstationary Time Series. *Proceedings of Machine Learning Research*. pp. 939–948.

Hammond C, Bergman H, Brown P. 2007. Pathological synchronization in Parkinson's disease: networks, models and treatments. *Trends in Neuroscience* **30**:357–364.

Hanslmayr S, Aslan A, Staudigl T, Klimesch W, Herrmann CS, Bäuml K-H. 2007. Prestimulus oscillations predict visual perception performance between and within subjects. *Neuroimage* **37**:1465–1473.

Harmony T. 2013. The functional significance of delta oscillations in cognitive processing. *Frontiers in Integrative Neuroscience* **7**:83.

Helfrich RF, Huang M, Wilson G, Knight RT. 2017. Prefrontal cortex modulates posterior alpha oscillations during top-down guided visual perception. *Proceeding of the National Academy of Sciences* **114**:9457–9462.

Helfrich RF, Knight RT. 2016. Oscillatory Dynamics of Prefrontal Cognitive Control. *Trends in Cognitive Science* **20**:916–930.

- Helmstaedter M, Sakmann B, Feldmeyer D. 2009. The relation between dendritic geometry, electrical excitability, and axonal projections of L2/3 interneurons in rat barrel cortex. *Cerebral Cortex* **19**:938–950.
- Hennig JA, Golub MD, Lund PJ, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Yu BM, Chase SM. 2018. Constraints on neural redundancy. *eLife* **7**:e36774.
- Henson RN, Mattout J, Phillips C, Friston KJ. 2009. Selecting forward models for MEG source-reconstruction using model-evidence. *Neuroimage* **46**:168–176.
- Hershey S, Chaudhuri S, Ellis DPW, Gemmeke JF, Jansen A, Moore C, Plakal M, Platt D, Saurous RA, Seybold B, Slaney M, Weiss R, Wilson K. 2017. CNN Architectures for Large-Scale Audio Classification. *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 131-135.
- Himberg J, Hyvärinen A, Esposito F. 2004. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage* **22**:1214–1222.
- Hirschmann J, Özkurt TE, Butz M, Homburger M, Elben S, Hartmann CJ, Vesper J, Wojtecki L, Schnitzler A. 2013. Differential modulation of STN-cortical and cortico-muscular coherence by movement and levodopa in Parkinson’s disease. *Neuroimage* **68**:203–213.
- Hlushchuk Y, Hari R. 2006. Transient suppression of ipsilateral primary somatosensory cortex during tactile finger stimulation. *Journal of Neuroscience* **26**:5819–5824.
- Hoffman MD, Blei DM, Wang C, Paisley J. 2013. Stochastic Variational Inference. *Journal of Machine Learning Research* **14**:1303–1347.
- Holt AB, Kormann E, Gulberti A, Pötter-Nerger M, McNamara CG, Cagnan H, Baaske MK, Little S, Köppen JA, Buhmann C, Westphal M, Gerloff C, Engel AK, Brown P, Hamel W, Moll CKE, Sharott A. 2019. Phase-Dependent Suppression of Beta Oscillations in Parkinson’s Disease Patients. *Journal of Neuroscience* **39**:1119–1134.
- Huang MX, Mosher JC, Leahy RM. 1999. A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. *Physics in Medicine and Biology* **44**:423.
- Humphries M. 2021. *The Spike: An Epic Journey Through the Brain in 2.1 Seconds*. Princeton University Press.
- Huntenburg JM, Bazin P-L, Goulas A, Tardif CL, Villringer A, Margulies DS. 2017. A Systematic Relationship Between Functional Connectivity and Intracortical Myelin in the Human Cerebral Cortex. *Cerebral Cortex* **27**:981–997.
- Husain M, Nachev P. 2007. Space and the parietal cortex. *Trends in Cognitive Science* **11**:30–36.
- Hyvarinen A. 1999. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks* **10**:626–634.

- Hyvarinen A. 1997. One-unit contrast functions for independent component analysis: a statistical analysis. *Proceedings of the 1997 IEEE Signal Processing Society Workshop*. pp. 388–397.
- Hyvärinen A, Oja E. 2000. Independent component analysis: algorithms and applications. *Neural Networks* **13**:411–430.
- Hyvarinen A, Sasaki H, Turner R. 2019. Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning. *Proceeding in Machine Learning Research*. pp. 859–868.
- Ibarra-Lecue I, Haegens S, Harris AZ. 2022. Breaking Down a Rhythm: Dissecting the Mechanisms Underlying Task-Related Neural Oscillations. *Frontiers in Neural Circuits* **16**:846905.
- Iemi L, Busch NA, Laudini A, Haegens S, Samaha J, Villringer A, Nikulin VV. 2019. Multiple mechanisms link prestimulus neural oscillations to sensory responses. *eLife* **8**:e43620.
- Indic P, Pratap R, Nampoori VP, Pradhan N. 1999. Significance of time scales in nonlinear dynamical analysis of electroencephalogram signals. *International Journal of Neuroscience* **99**:181–194.
- Jensen O, Mazaheri A. 2010. Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Frontiers in Human Neuroscience* **4**:186.
- Johnson M, Willsky A. 2014. Stochastic Variational Inference for Bayesian Time Series Models. *Proceedings of Machine Learning Research*. pp. 1854–1862.
- Jones DT, Vemuri P, Murphy MC, Gunter JL, Senjem ML, Machulda MM, Przybelski SA, Gregg BE, Kantarci K, Knopman DS, Boeve BF, Petersen RC, Jack CR. 2012. Non-Stationarity in the “Resting Brain’s” Modular Architecture. *PLoS One* **7**:e39731.
- Jones SR. 2016. When brain rhythms aren’t “rhythmic”: implication for their mechanisms and meaning. *Current Opinion in Neurobiology* **40**:72–80.
- Jubault T, Monetta L, Strafella AP, Lafontaine A-L, Monchi O. 2009. L-Dopa Medication in Parkinson’s Disease Restores Activity in the Motor Cortico-Striatal Loop but Does Not Modify the Cognitive Network. *PLoS One* **4**:e6154.
- Kanayama N, Tamè L, Ohira H, Pavani F. 2012. Top down influence on visuo-tactile interaction modulates neural oscillatory responses. *Neuroimage* **59**:3406–3417.
- Kaplan AYa, Fingelkurts Andrew A., Fingelkurts Alexander A., Borisov SV, Darkhovsky BS. 2005. Nonstationary nature of the brain activity as revealed by EEG/MEG: Methodological, practical and conceptual challenges. *Signal Processing* **85**:2190–2212.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2014. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience* **17**:440–448.
- Kawaguchi Y, Karube F, Kubota Y. 2006. Dendritic Branch Typing and Spine Expression Patterns in Cortical Nonpyramidal Cells. *Cerebral Cortex* **16**:696–711.

- Keitel A, Gross J. 2016. Individual Human Brain Areas Can Be Identified from Their Characteristic Spectral Activation Fingerprints. *PLoS Biology* **14**:e1002498.
- Kelly C, de Zubicaray G, Di Martino A, Copland DA, Reiss PT, Klein DF, Castellanos FX, Milham MP, McMahon K. 2009. L-Dopa Modulates Functional Connectivity in Striatal Cognitive and Motor Networks: A Double-Blind Placebo-Controlled Study. *Journal of Neuroscience* **29**:7364–7378.
- Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, Shenoy KV. 2008. Detecting Neural-State Transitions Using Hidden Markov Models for Motor Cortical Prostheses. *Journal of Neurophysiology* **100**:2441–2452.
- Khambhati AN, Sizemore AE, Betzel RF, Bassett DS. 2018. Modeling and interpreting mesoscale network dynamics. *NeuroImage* **180**:337–349.
- Kim Y, Hsu C-L, Cembrowski MS, Mensh BD, Spruston N. 2015. Dendritic sodium spikes are required for long-term potentiation at distal synapses on hippocampal pyramidal neurons. *eLife* **4**:e06414.
- Klimesch W. 2012. α -band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Science* **16**:606–617.
- Kornblith S, Buschman TJ, Miller EK. 2016. Stimulus Load and Oscillatory Activity in Higher Cortex. *Cerebral Cortex* **26**:3772–3784.
- Kreiman G, Serre T. 2020. Beyond the feedforward sweep: feedback computations in the visual cortex. *Annals of the New York Academy of Sciences* **1464**:222–241.
- Kriegeskorte N, Douglas PK. 2019. Interpreting encoding and decoding models. *Current Opinion in Neurobiology, Machine Learning, Big Data, and Neuroscience* **55**:167–179.
- Krizhevsky A, Sutskever I, Hinton GE. 2012. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*. Curran Associates, Inc. pp. 1097-1105.
- Lalo E, Thobois S, Sharott A, Polo G, Mertens P, Pogosyan A, Brown P. 2008. Patterns of bidirectional communication between cortex and basal ganglia during movement in patients with Parkinson disease. *Journal of Neuroscience* **28**:3008–3016.
- Layton OW, Mingolla E, Yazdanbakhsh A. 2014. Neural dynamics of feedforward and feedback processing in figure-ground segregation. *Frontiers in Psychology* **5**:972.
- Leguey I, Benavides-Piccione R, Rojo C, Larrañaga P, Bielza C, DeFelipe J. 2018. Patterns of Dendritic Basal Field Orientation of Pyramidal Neurons in the Rat Somatosensory Cortex. *eNeuro* **5**:ENEURO.0142-18.2018.
- Li Y, Ward MJ, Richardson RM, G'Sell M, Ghuman AS. 2020. Endogenous activity modulates stimulus and circuit-specific neural tuning and predicts perceptual behavior. *Nature Communication* **11**:4014.

- Lian W, Talmon R, Zaveri H, Carin L, Coifman R. 2015. Multivariate time-series analysis and diffusion maps. *Signal Processing* **116**:13–28.
- Linkenkaer-Hansen K, Nikulin VV, Palva S, Ilmoniemi RJ, Palva JM. 2004. Prestimulus Oscillations Enhance Psychophysical Performance in Humans. *Journal of Neuroscience* **24**:10186–10190.
- Little S, Tan H, Anzak A, Pogosyan A, Kühn A, Brown P. 2013. Bilateral functional connectivity of the basal ganglia in patients with Parkinson’s disease and its modulation by dopaminergic treatment. *PLoS One* **8**:e82762.
- Litvak V, Florin E, Tamás G, Groppa S, Muthuraman M. 2021. EEG and MEG primers for tracking DBS network effects. *NeuroImage* **224**:117447.
- Litvak V, Jha A, Eusebio A, Oostenveld R, Foltynie T, Limousin P, Zrinzo L, Hariz MI, Friston K, Brown P. 2011. Resting oscillatory cortico-subthalamic connectivity in patients with Parkinson’s disease. *Brain* **134**:359–374.
- Litwin-Kumar A, Doiron B. 2012. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience* **15**:1498–1505.
- London M, Häusser M. 2005. Dendritic computation. *Annual Reviews in Neuroscience* **28**:503–532.
- Lundqvist M, Herman P, Warden MR, Brincat SL, Miller EK. 2018. Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature Communication* **9**:394.
- MacDonald PA, MacDonald AA, Seergobin KN, Tamjeedi R, Ganjavi H, Provost J-S, Monchi O. 2011. The effect of dopamine therapy on ventral and dorsal striatum-mediated cognition in Parkinson’s disease: support from functional MRI. *Brain* **134**:1447–1463.
- Macdonald PA, Monchi O. 2011. Differential effects of dopaminergic therapies on dorsal and ventral striatum in Parkinson’s disease: implications for cognitive function. *Parkinsons Disease* **2011**:572743.
- Marco-Pallarés J, Münte TF, Rodríguez-Fornells A. 2015. The role of high-frequency oscillatory activity in reward processing and learning. *Neuroscience Biobehavioral Reviews* **49**:1–7.
- Margulies DS, Ghosh SS, Goulas A, Falkiewicz M, Huntenburg JM, Langs G, Bezgin G, Eickhoff SB, Castellanos FX, Petrides M, Jefferies E, Smallwood J. 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceeding of the National Academy of Sciences* **113**:12574–12579.
- Marinelli L, Quartarone A, Hallett M, Frazzitta G, Ghilardi MF. 2017. The many facets of motor learning and their relevance for Parkinson’s disease. *Clinical Neurophysiology* **128**:1127–1141.
- Marreiros AC, Cagnan H, Moran RJ, Friston KJ, Brown P. 2013. Basal ganglia–cortical interactions in Parkinsonian patients. *Neuroimage* **66**:301–310.

- Martínez-Vázquez P, Gail A. 2018. Directed Interaction Between Monkey Premotor and Posterior Parietal Cortex During Motor-Goal Retrieval from Working Memory. *Cerebral Cortex* **28**:1866–1881.
- Mazaheri A, van Schouwenburg MR, Dimitrijevic A, Denys D, Cools R, Jensen O. 2014. Region-specific modulations in oscillatory alpha activity serve to facilitate processing in the visual and auditory modalities. *NeuroImage* **87**:356–362.
- Mckeown MJ, Makeig S, Brown GG, Jung T-P, Kindermann SS, Bell AJ, Sejnowski TJ. 1998. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping* **6**:160–188.
- Mellem MS, Wohltjen S, Gotts SJ, Ghuman AS, Martin A. 2017. Intrinsic frequency biases and profiles across human cortex. *Journal of Neurophysiology* **118**:2853–2864.
- Miller EK, Lundqvist M, Bastos AM. 2018. Working Memory 2.0. *Neuron* **100**:463–475.
- Miska NJ, Richter LM, Cary BA, Gjorgjieva J, Turrigiano GG. 2018. Sensory experience inversely regulates feedforward and feedback excitation-inhibition ratio in rodent visual cortex. *eLife* **7**:e38846.
- Montani F, Ince RAA, Senatore R, Arabzadeh E, Diamond ME, Panzeri S. 2009. The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **367**:3297–3310.
- Moran A, Bar-Gad I, Bergman H, Israel Z. 2006. Real-time refinement of subthalamic nucleus targeting using Bayesian decision-making on the root mean square measure. *Movement Disorders* **21**:1425–1431.
- Moses DA, Metzger SL, Liu JR, Anumanchipalli GK, Makin JG, Sun PF, Chartier J, Dougherty ME, Liu PM, Abrams GM, Tu-Chan A, Ganguly K, Chang EF. 2021. Neuroprosthesis for Decoding Speech in a Paralyzed Person with Anarthria. *New England Journal of Medicine* **385**:217–227.
- Munck JC de, Wolters CH, Clerc M. 2012. EEG and MEG: forward modeling. *Handbook of Neural Activity Measurement*. pp. 192-256.
- Naud R, Sprekeler H. 2018. Sparse bursts optimize information transmission in a multiplexed neural code. *Proceeding of the National Academy of Sciences* **115**:E6329–E6338.
- Okun MS. 2012. Deep-Brain Stimulation for Parkinson’s Disease. *New England Journal of Medicine* **367**:1529–1538.
- Olde Dubbelink KTE, Stoffers D, Deijen JB, Twisk JWR, Stam CJ, Berendse HW. 2013a. Cognitive decline in Parkinson’s disease is associated with slowing of resting-state brain activity: a longitudinal study. *Neurobiology of Aging* **34**:408–418.
- Olde Dubbelink KTE, Stoffers D, Deijen JB, Twisk JWR, Stam CJ, Hillebrand A, Berendse HW. 2013b. Resting-state functional connectivity as a marker of disease progression in Parkinson’s disease: A longitudinal MEG study. *Neuroimage Clinical* **2**:612–619.

- Olejniczak P. 2006. Neurophysiologic Basis of EEG. *Journal of Clinical Neurophysiology* **23**:186–189.
- Oswal A, Beudel M, Zrinzo L, Limousin P, Hariz M, Foltynie T, Litvak V, Brown P. 2016a. Deep brain stimulation modulates synchrony within spatially and spectrally distinct resting state networks in Parkinson's disease. *Brain* **139**:1482–1496.
- Oswal A, Brown P, Litvak V. 2013. Movement related dynamics of subthalamo-cortical alpha connectivity in Parkinson's disease. *Neuroimage* **70**:132–142.
- Oswal A, Jha A, Neal S, Reid A, Bradbury D, Aston P, Limousin P, Foltynie T, Zrinzo L, Brown P, Litvak V. 2016b. Analysis of simultaneous MEG and intracranial LFP recordings during Deep Brain Stimulation: a protocol and experimental validation. *Journal of Neuroscience Methods* **261**:29–46.
- Packard NH, Crutchfield JP, Farmer JD, Shaw RS. 1980. Geometry from a Time Series. *Physical Review Letters* **45**:712–716.
- Paquola C, Wael RVD, Wagstyl K, Bethlehem RAI, Hong S-J, Seidlitz J, Bullmore ET, Evans AC, Misisic B, Margulies DS, Smallwood J, Bernhardt BC. 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLOS Biology* **17**:e3000284.
- Paun I, Husmeier D, Hopcraft JGC, Masolele MM, Torney CJ. 2022. Inferring spatially varying animal movement characteristics using a hierarchical continuous-time velocity model. *Ecology Letters* **25**:2726:2738.
- Pesaran B, Vinck M, Einevoll GT, Sirota A, Fries P, Siegel M, Truccolo W, Schroeder CE, Srinivasan R. 2018. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nature Neuroscience* **21**:903–919.
- Phillips CG, Zeki S, Barlow HB. 1984. Localization of function in the cerebral cortex. Past, present and future. *Brain* **107**:327–361.
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. 2008. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**:995–999.
- Podvalny E, Flounders MW, King LE, Holroyd T, He BJ. 2019. A dual role of prestimulus spontaneous neural activity in visual object recognition. *Nature Communication* **10**:3910.
- Poirazi P, Papoutsi A. 2020. Illuminating dendritic function with computational models. *Nature Reviews Neuroscience* **21**:303–321.
- Rajagovindan R, Ding M. 2011. From prestimulus alpha oscillation to visual-evoked response: an inverted-U function and its attentional modulation. *Journal of Cognitive Neuroscience* **23**:1379–1394.
- Raposo D, Kaufman MT, Churchland AK. 2014. A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience* **17**:1784–1792.

- Rasmussen CE. 2004. Gaussian Processes in Machine Learning In: Bousquet O, von Luxburg U, Rätsch G, editors. *Advanced Lectures on Machine Learning: ML Summer Schools 2003*, Lecture Notes in Computer Science: pp. 63–71.
- Rassi E, Wutz A, Müller-Voggel N, Weisz N. 2019. Prestimulus feedback connectivity biases the content of visual experiences. *Proceedings of the National Academy of Sciences* **116**:16056–16061.
- Ray N, Strafella AP. 2010. Dopamine, Reward, and Frontostriatal Circuitry in Impulse Control Disorders in Parkinson’s Disease: Insights from Functional Imaging. *Clinical EEG Neuroscience* **41**:87–93.
- Recanatesi S, Pereira-Obilinovic U, Murakami M, Mainen Z, Mazzucato L. 2022. Metastable attractors explain the variable timing of stable behavioral action sequences. *Neuron* **110**:139-153.e9.
- Regazzoni F, Salvador M, Dede’ L, Quarteroni A. 2022. A machine learning method for real-time numerical simulations of cardiac electromechanics. *Computer Methods in Applied Mechanics and Engineering* **393**:114825.
- Reuter M, Schmansky NJ, Rosas HD, Fischl B. 2012. Within-subject template estimation for unbiased longitudinal image analysis. *Neuroimage* **61**:1402–1418.
- Reyes AD. 2003. Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nature Neuroscience* **6**:593–599.
- Roelfsema PR, Holtmaat A. 2018. Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience* **19**:166–180.
- Royal I, Vuvan DT, Zendel BR, Robitaille N, Schönwiesner M, Peretz I. 2016. Activation in the Right Inferior Parietal Lobule Reflects the Representation of Musical Structure beyond Simple Pitch Discrimination. *PLOS ONE* **11**:e0155291.
- Sacchet MD, LaPlante RA, Wan Q, Pritchett DL, Lee AKC, Hämäläinen M, Moore CI, Kerr CE, Jones SR. 2015. Attention drives synchronization of alpha and beta rhythms between right inferior frontal and primary sensory neocortex. *Journal of Neuroscience* **35**:2074–2082.
- Samaha J, Lemi L, Haegens S, Busch NA. 2020. Spontaneous Brain Oscillations and Perceptual Decision-Making. *Trends in Cognitive Sciences* **24**:639–653.
- Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. 2006. A high-performance brain–computer interface. *Nature* **442**:195–198.
- Schnitzler A, Gross J. 2005. Normal and pathological oscillatory communication in the brain. *Nature Reviews Neuroscience* **6**:285–296.
- Seidlitz J, Váša F, Shinn M, Romero-Garcia R, Whitaker KJ, Vértes PE, Wagstyl K, Kirkpatrick Reardon P, Clasen L, Liu S, Messinger A, Leopold DA, Fonagy P, Dolan RJ, Jones PB, Goodyer IM, Raznahan A, Bullmore ET. 2018. Morphometric Similarity Networks Detect Microscale Cortical Organization and Predict Inter-Individual Cognitive Variation. *Neuron* **97**:231-247.

- Semedo JD, Jasper AI, Zandvakili A, Krishna A, Aschner A, Machens CK, Kohn A, Yu BM. 2022. Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communication* **13**:1099.
- Semedo JD, Zandvakili A, Machens CK, Yu BM, Kohn A. 2019. Cortical Areas Interact through a Communication Subspace. *Neuron* **102**:249–259.
- Shadlen MN, Newsome WT. 1998. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience* **18**:3870–3896.
- Shamir M, Sompolinsky H. 2004. Nonlinear population codes. *Neural Computation* **16**:1105–1136.
- Sharma A, Vidaurre D, Vesper J, Schnitzler A, Florin E. 2021. Differential dopaminergic modulation of spontaneous cortico–subthalamic activity in Parkinson’s disease. *eLife* **10**:e66057.
- Sherman MA, Lee S, Law R, Haegens S, Thorn CA, Hämäläinen MS, Moore CI, Jones SR. 2016. Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. *Proceeding of the National Academy of Sciences* **113**:E4885–4894.
- Sheth SA, Nemoto M, Guiou M, Walker M, Pouratian N, Toga AW. 2004. Linear and nonlinear relationships between neuronal activity, oxygen metabolism, and hemodynamic responses. *Neuron* **42**:347–355.
- Shimamoto SA, Ryapolova-Webb ES, Ostrem JL, Galifianakis NB, Miller KJ, Starr PA. 2013. Subthalamic nucleus neurons are synchronized to primary motor cortex local field potentials in Parkinson’s disease. *Journal of Neuroscience* **33**:7220–7233.
- Shine JM. 2019. Neuromodulatory Influences on Integration and Segregation in the Brain. *Trends in Cognitive Sciences* **23**:572–583.
- Shohamy D, Myers CE, Grossman S, Sage J, Gluck MA. 2005. The role of dopamine in cognitive sequence learning: evidence from Parkinson’s disease. *Behavioural Brain Research* **156**:191–199.
- Siegel M, Buschman TJ, Miller EK. 2015. Cortical information flow during flexible sensorimotor decisions. *Science* **348**:1352–1355.
- Singer W. 2018. Neuronal oscillations: unavoidable and useful? *European Journal of Neuroscience* **48**:2389–2398.
- Singh-Curry V, Husain M. 2009. The functional role of the inferior parietal lobe in the dorsal and ventral stream dichotomy. *Neuropsychologia* **47**:1434–1448.
- Smith J, Linderman S, Sussillo D. 2021. Reverse engineering recurrent neural networks with Jacobian switching linear dynamical systems. *Advances in Neural Information Processing Systems*. Curran Associates, Inc. pp. 16700–16713.

- Smith ST. 2005. Covariance, subspace, and intrinsic Cramer-Rao bounds. *IEEE Transactions on Signal Process* **53**:1610–1630.
- Spitzer B, Fleck S, Blankenburg F. 2014a. Parametric alpha- and beta-band signatures of supramodal numerosity information in human working memory. *Journal of Neuroscience* **34**:4293–4302.
- Spitzer B, Gloel M, Schmidt TT, Blankenburg F. 2014b. Working memory coding of analog stimulus properties in the human prefrontal cortex. *Cerebral Cortex* **24**:2229–2236.
- Spitzer B, Haegens S. 2017. Beyond the Status Quo: A Role for Beta Oscillations in Endogenous Content (Re)Activation. *eNeuro* **4**:ENEURO.0170-17.2017.
- Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A. 2015. Going deeper with convolutions. 2015 *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1–9.
- Tadel F, Baillet S, Mosher JC, Pantazis D, Leahy RM. 2011. Brainstorm: A User-Friendly Application for MEG/EEG Analysis. *Computational Intelligence and Neuroscience* **2011**:879716.
- Tallon-Baudry C, Bertrand O, Peronnet F, Pernier J. 1998. Induced gamma-band activity during the delay of a visual short-term memory task in humans. *Journal of Neuroscience* **18**:4244–4254.
- Thiebaut de Schotten M, Forkel SJ. 2022. The emergent properties of the connected brain. *Science* **378**:505–510.
- Thut G, Nietzel A, Brandt SA, Pascual-Leone A. 2006. α -Band Electroencephalographic Activity over Occipital Cortex Indexes Visuospatial Attention Bias and Predicts Visual Target Detection. *Journal of Neuroscience* **26**:9494–9502.
- Tinkhauser G, Pogosyan A, Tan H, Herz DM, Kühn AA, Brown P. 2017. Beta burst dynamics in Parkinson's disease OFF and ON dopaminergic medication. *Brain* **140**:2968–2981.
- Tsubokawa H, Offermanns S, Simon M, Kano M. 2000. Calcium-Dependent Persistent Facilitation of Spike Backpropagation in the CA1 Pyramidal Neurons. *Journal of Neuroscience* **20**:4878–4884.
- Uhlhaas P, Pipa G, Lima B, Melloni L, Neuenschwander S, Nikolić D, Singer W. 2009. Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience* **3**:17.
- Urai AE, Doiron B, Leifer AM, Churchland AK. 2022. Large-scale neural recordings call for new insights to link brain and behavior. *Nature Neuroscience* **25**:11–19.
- Vaillancourt DE, Schonfeld D, Kwak Y, Bohnen NI, Seidler R. 2013. Dopamine overdose hypothesis: Evidence and clinical implications. *Movement Disorders* **28**:1920-1929.

- van Bergen RS, Kriegeskorte N. 2020. Going in circles is the way forward: the role of recurrence in visual inference. *Current Opinion in Neurobiology*, Whole-brain interactions between neural circuits **65**:176–193.
- van den Heuvel MP, Yeo BTT. 2017. A Spotlight on Bridging Microscale and Macroscale Human Brain Architecture. *Neuron* **93**:1248–1251.
- van Ede F, de Lange FP, Maris E. 2014. Anticipation increases tactile stimulus processing in the ipsilateral primary somatosensory cortex. *Cerebral Cortex* **24**:2562–2571.
- van Kerkoerle T, Self MW, Dagnino B, Gariel-Mathis M-A, Poort J, van der Togt C, Roelfsema PR. 2014. Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences* **111**:14332–14341.
- van Wijk BCM, Beudel M, Jha A, Oswal A, Foltynie T, Hariz MI, Limousin P, Zrinzo L, Aziz TZ, Green AL, Brown P, Litvak V. 2016. Subthalamic nucleus phase–amplitude coupling correlates with motor impairment in Parkinson’s disease. *Clinical Neurophysiology* **127**:2010–2019.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. 2017. Attention is All You Need. *Advances in Neural Information Processing Systems*. Curran Associates, Inc. pp. 6000-6010.
- Vidaurre D, Abeysuriya R, Becker R, Quinn AJ, Alvaro-Almagro F, Smith SM, Woolrich MW. 2018a. Discovering dynamic brain networks from big data in rest and task. *Neuroimage* **180**:646–656.
- Vidaurre D, Hunt LT, Quinn AJ, Hunt BAE, Brookes MJ, Nobre AC, Woolrich MW. 2018b. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. *Nature Communication* **9**:2987.
- Vidaurre D, Quinn AJ, Baker AP, Dupret D, Tejero-Cantero A, Woolrich MW. 2016. Spectrally resolved fast transient brain states in electrophysiological data. *Neuroimage* **126**:81–95.
- Viola P, Jones M. 2001. Rapid object detection using a boosted cascade of simple features. *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. CVPR. pp 1–14.
- Viterbi A. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory* **13**:260–269.
- Vollebregt MA, Zumer JM, Ter Huurne N, Buitelaar JK, Jensen O. 2016. Posterior alpha oscillations reflect attentional problems in boys with Attention Deficit Hyperactivity Disorder. *Clinical Neurophysiology* **127**:2182–2191.
- Voon V, Fernagut P-O, Wickens J, Baunez C, Rodriguez M, Pavon N, Juncos JL, Obeso JA, Bezdard E. 2009. Chronic dopaminergic stimulation in Parkinson’s disease: from dyskinesias to impulse control disorders. *Lancet Neurol* **8**:1140–1149.
- Voytek B, Kayser AS, Badre D, Fegen D, Chang EF, Crone NE, Parvizi J, Knight RT, D’Esposito M. 2015. Oscillatory dynamics coordinating human frontal networks in support of goal maintenance. *Nature Neuroscience* **18**:1318–1324.

- Wael RV de, Benkarim O, Paquola C, Lariviere S, Royer J, Tavakol S, Xu T, Hong S-J, Langs G, Valk S, Mistic B, Milham M, Margulies D, Smallwood J, Bernhardt BC. 2020. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Communications Biology* 2020 3:1 3:1–10.
- Wang X-J. 2010. Neurophysiological and computational principles of cortical rhythms in cognition. *Physical Reviews* 90:1195–1268.
- Warling A, McDermott CL, Liu S, Seidlitz J, Rodrigue AL, Nadig A, Gur RC, Gur RE, Roalf D, Moore TM, Glahn D, Satterthwaite TD, Bullmore ET, Raznahan A. 2021. Regional White Matter Scaling in the Human Brain. *Journal of Neuroscience* 41:7015–7028.
- West T, Farmer S, Berthouze L, Jha A, Beudel M, Foltynie T, Limousin P, Zrinzo L, Brown P, Litvak V. 2016. The Parkinsonian Subthalamic Network: Measures of Power, Linear, and Non-linear Synchronization and their Relationship to L-DOPA Treatment and OFF State Motor Severity. *Frontiers in Human Neuroscience* 10:517.
- West TO, Berthouze L, Halliday DM, Litvak V, Sharott A, Magill PJ, Farmer SF. 2018. Propagation of beta/gamma rhythms in the cortico-basal ganglia circuits of the parkinsonian rat. *Journal of Neurophysiology* 119:1608–1628.
- Whalen TC, Willard AM, Rubin JE, Gittis AH. 2020. Delta oscillations are a robust biomarker of dopamine depletion severity and motor dysfunction in awake mice. *Journal of Neurophysiology* 124:312–329.
- Willett FR, Avansino DT, Hochberg LR, Henderson JM, Shenoy KV. 2021. High-performance brain-to-text communication via handwriting. *Nature* 593:249–254.
- Williams AH, Linderman SW. 2021. Statistical neuroscience in the single trial limit. *Current Opinion in Neurobiology* 70:193–205.
- Wilson HR, Cowan JD. 1972. Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. *Biophysical Journal* 12:1–24.
- Wiltschko AB, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abaira VE, Adams RP, Datta SR. 2015. Mapping Sub-Second Structure in Mouse Behavior. *Neuron* 88:1121–1135.
- Wyatte D, Curran T, O'Reilly R. 2012. The limits of feedforward vision: recurrent processing promotes robust object recognition when objects are degraded. *Journal of Cognitive Neuroscience* 24:2248–2261.
- Xing D, Yeh C-I, Shapley RM. 2009. Spatial Spread of the Local Field Potential and its Laminar Variation in Visual Cortex. *Journal of Neuroscience* 29:11540–11549.
- Yang Q, Walker E, Cotton RJ, Tolia AS, Pitkow X. 2021. Revealing nonlinear neural decoding by analyzing choices. *Nature Communication* 12:6557.
- Zavala BA, Tan H, Little S, Ashkan K, Hariz M, Foltynie T, Zrinzo L, Zaghoul KA, Brown P. 2014. Midline frontal cortex low-frequency activity drives subthalamic nucleus oscillations during conflict. *Journal of Neuroscience* 34:7322–7333.

Zhou YJ, Iemi L, Schoffelen J-M, Lange FP de, Haegens S. 2021. Alpha Oscillations Shape Sensory Representation and Perceptual Sensitivity. *Journal of Neuroscience* **41**:9581–9592.

Statutory declaration

I, **Abhinav Sharma**, hereby declare, that the work for my dissertation “Unsupervised Machine Learning for Human Neural Electrophysiology Signals” was realized independently and without using illicit help in compliance with the "Principles to ensure good scientific practice at the Heinrich-Heine-University of Düsseldorf". Only the declared sources have been used, and all direct and analogous quotes have been marked as such. The work has not been submitted in identical or similar variety at any other institution. This is my first attempt at acquiring a doctoral degree.

Düsseldorf, January 05 2023

ABHINAV SHARMA

Appendix

Contributions

Study 1 Evaluating a data driven pipeline for extracting EEG resting state networks in humans

Contributions

Abhinav Sharma: Conceptualization, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (55% contribution)

Dr. Esther Pelzer: Analysis of fMRI data

Dr. Esther Florin: Conceptualization, Software, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing

Study 2 Differential dopaminergic modulation of spontaneous cortico–subthalamic activity in Parkinson’s disease

Note: This study has been published in a peer reviewed journal (<https://elifesciences.org/articles/66057>)

Contributions

Abhinav Sharma: Conceptualization, Partial data acquisition and quality control, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (65% contribution)

Dr. Diego Vidaurre: Software, Validation, Methodology, Writing - review and editing

Dr. Jan Vesper: Resources, Investigation

Dr. Alfons Schnitzler: Resources, Investigation, Writing - review and editing

Dr. Esther Florin: Conceptualization, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing

Study 3 Latent cognitive network transitions track somatosensory perceptual variability

Contributions

Abhinav Sharma: Conceptualization, Complete data acquisition and quality control, Software, Formal analysis, Investigation, Visualization, Methodology, Writing - original draft (65% contribution)

Dr. Joachim Lange: Behavioral paradigm, Methodology

Dr. Diego Vidaurre: Software

Dr. Esther Florin: Conceptualization, Resources, Supervision, Funding acquisition, Validation, Investigation, Methodology, Project administration, Writing - review and editing