1-12-2022 2:00 PM

Mapping The Functional Organization of Human Frontoparietal Cortex With fMRI

Geetika Gupta, *The University of Western Ontario*

Supervisor: Mur, Marieke, *The University of Western Ontario*
Co-Supervisor: Khan, Ali R., *The University of Western Ontario*
A thesis submitted in partial fulfillment of the requirements for the Master of Science degree in Neuroscience

Abstract

Higher-order cognitive functions, such as working memory, attention, and decision making, depend strongly on the functional integrity of frontal and parietal cortices. However, the internal workings of the frontoparietal network (FPN) are not well understood. A major contributor to this knowledge gap is our limited understanding of the intrinsic functional organization of the FPN. In order to address this gap, we examine task-dependent reconfigurations of functional connectivity (FC) within the FPN. We analyzed fMRI task-state data from 924 individuals from the Human Connectome Project Young Adult study. Our results show that FC within the FPN is highly stable across time within individuals. Furthermore, FC within the FPN is more consistent within than between individuals and more consistent within than between tasks. Overall, our findings indicate that human individuals exhibit a partially unique fine-grained functional organization within the FPN, and that this organization contains a task-specific component.

Keywords: functional magnetic resonance imaging, frontoparietal network, functional connectivity, cognitive control, individual variability

Summary for Lay Audience

Our ability to plan ahead and flexibly adapt our behaviour to novel situations relies heavily on a network of frontal and parietal brain regions. These brain regions become active when we remember the recent past, make and execute plans, or become aware of mistakes. How do these brain regions support this diverse range of cognitive functions?

Research on human frontoparietal function has focused on identifying one-to-one mappings between brain regions and cognitive functions. This approach has not been as successful as originally anticipated. The majority of frontoparietal brain regions are not involved in just one but many cognitive functions, including working memory, decision making, and error detection. This suggests that a change in approach is necessary: instead of attempting to link activity of whole brain regions to specific cognitive functions, we should examine how local processing within the frontoparietal network supports behaviour during diverse cognitive tasks.

In order to understand local processing we need to understand the fine-grained functional organization of the frontoparietal network. In other words, we need to understand how information flows through the network during the execution of cognitive tasks. In this project, we aim to map the fine-grained functional organization of the frontoparietal network using functional magnetic resonance imaging (fMRI) data acquired in a large group of individuals. Our results indicate that the frontoparietal network indeed shows a fine-grained functional organization, and that this organization contains an individually unique component. Furthermore, the functional organization contains a task-general as well as a task-specific component.

Results will greatly enhance our understanding of the neural basis of human cognition and provide a starting point for exploring functional organization level targets to study causes of cognitive dysfunction.

Acknowledgements

I would like to thank my supervisors Drs. Marieke Mur and Ali Khan, without whose consistent support, guidance and assistance, this thesis project wouldn't have been possible. I am thankful for having been a recipient of Marieke's knowledge and also admire her work ethic and positive attitude. Although there were many ups and downs throughout the project, Marieke's positivity inspired me to continue working hard. I am also grateful for all of the opportunities Ali has provided over the past couple of years. Ali was my first mentor at Western and gave me the opportunity to get involved in neuroimaging research. Thank you for fostering an environment that provided the intellectual resources and support for me and other members of the lab.

I would also like to thank my advisory committee members Drs. Julio Martinez-Trujillo and Jörn Diedrichsen for their expertise and guidance. Thank you for providing constructive feedback during our meetings and taking the time to mentor me in crucial areas of my project.

This project would also not have been possible without members of the Mur and Khan labs. I would especially like to thank Derrick Xiang for contributing his computational expertise during the data analysis stage. A special thanks to colleagues and friends - Ehsan Tousi, Reebal Rafeh, Ladan Shahshahani, David Mekhaiel, Patrick Park, Kajal Patel, Raveena Berde, and Dolly Dhaliwal. Although there were many tough times over the past two years, especially with the pandemic, I was lucky to have great company along the way. Each of you have provided an abundance of emotional support and encouragement of my career goals, which has made me a stronger and better person.

Lastly, I am ever grateful for the support my parents and siblings have provided me. Everything I have today would not have been possible without the support and sacrifices they have made for my education.

Contents

List of Figures

List of Tables

List of Abbreviations

- **CSF** Cerebrospinal Fluid
- **DTI** Diffusion Tensor Imaging
- **EPI** Echo Planar Imaging
- **FC** Functional Connectivity
- **FIR** Finite Impulse Response
- **fMRI** Functional Magnetic Resonance Imaging
- **FPN** Frontoparietal Network
- **FWHM** Full Width at Half Maximum
- **GE** Gradient Echo
- **GLM** General Linear Model
- **GSR** Global Signal Regression
- **GVC** Global Variability Coefficient
- **HCP** Human Connectome Project
- **HRF** Hemodynamic Response Function
- **Hz** Hertz

- **ICA** Independent Component Analysis
- **ITI** Inter-task Interval
- **LFP** Local Field Potential
- **MDS** Multi-dimensional Scaling
- **MNI** Montreal Neurological Institute
- **MRI** Magnetic Resonance Imaging
- **MSM** Multimodal Surface Matching
- **RF** Radio frequency
- **ROI** Region of Interest
- **T** Tesla
- **TE** Echo Time
- **TR** Repetition Time
- **T1w** T1 weighted
- **T2w** T2 weighted
- **WM** Working Memory

Chapter 1

Introduction

1.1 Motivation

Our ability to plan ahead and flexibly adapt our behaviour to novel situations relies heavily on a network of frontal and parietal brain regions. Lesions to the frontoparietal network (FPN) yield impairments across a broad array of cognitive functions, including our ability to remember the recent past, make and execute plans, or become aware of mistakes (Dias et al., 1996; Goldman-Rakic, 1988; Luria, 1966; Passingham, 1975; Stuss et al., 2000; Woolgar et al., 2010). Neuroimaging studies in healthy human individuals have shown that the FPN becomes active during the execution of a broad range of cognitive tasks (Duncan, 2010) (figure 1.1). The network appears to function as a flexible hub that orchestrates sensory and motor regions in the brain when executing cognitive tasks (Cole et al., 2013). Despite its well-established role in human cognition, the internal workings of the FPN are not well understood.

To better understand the internal workings of the FPN, this thesis aims to map the network's functional organization using functional magnetic resonance imaging (fMRI). Functional mapping commonly involves measuring brain responses during the execution of cognitive tasks, and creating brain maps of response preferences for stimulus and task parameters. In sensory

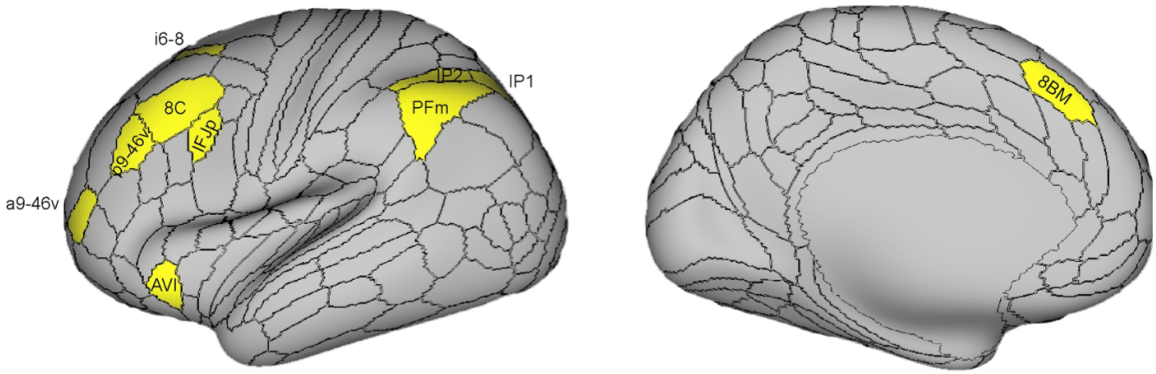

Figure 1.1: The frontoparietal network (FPN). The figure shows core regions of the FPN, which is recruited by a wide variety of cognitive demands, including working memory and decision making (Assem et al., 2020). The network is indicated in yellow on an inflated cortical surface showing the multimodal cortical parcellation from the Human Connectome Project (HCP) (Glasser et al., 2016).

and motor cortices, functional mapping has played an important role in understanding how the brain supports perception and action. For example, sensory brain regions contain interpretable maps of visual and auditory stimulus properties, including retinotopic maps in visual cortex and tonotopic maps in auditory cortex (Sereno et al., 1995; Formisano et al., 2003). These maps reveal how the brain represents the outside world and also provide clues about how brain regions are organized into larger systems. For example, retinotopic maps in the visual cortex have revealed a hierarchical system of brain regions that supports object recognition (Sereno et al., 1995; Levy et al., 2001). Despite these successes in the sensory domain, creating interpretable functional maps of frontoparietal cortices has proven challenging.

Early neuroimaging work in humans focused on identifying one-to-one mappings between frontoparietal brain regions and cognitive functions (Rypma et al., 1999; Smith et al., 2009). This approach assumes that one region controls one function. However, given that the FPN as a whole becomes active during the execution of a broad range of cognitive tasks (Duncan and Owen, 2000; Duncan, 2010; Poldrack, 2011), this approach has not been as fruitful as

originally anticipated. Prior work in nonhuman primates is largely consistent with these findings. While some degree of regional functional specialization in frontoparietal cortex is expected given cytoarchitectonic and connectomic differences (Petrides and Pandya, 1999; Cavada and Goldman-Rakic, 1989), differences in neural response properties across the network tend to be relatively subtle (Duncan, 2001, 2010, 2013). At any moment in time, the majority of neurons within the FPN appears to be engaged in coding information relevant to the ongoing cognitive task, but evidence for spatial clustering of neurons with similar response properties is sparse (Duncan, 2001; Machens et al., 2010; Leavitt et al., 2018). Large-scale spatial clustering according to response properties may not be the most efficient solution given the computational goals of the FPN, which emphasize flexible integration of information across domains as opposed to domain-specific representation of information as observed in sensory cortices (Miller and Cohen, 2001; Duncan, 2001; Rigotti, 2010; Duncan, 2010). The above findings, together with the limited ability of fMRI to capture dynamic neural interactions at a fast time scale, may explain the lack of a clear consensus in the human literature on the functional organization of prefrontal cortex (Duncan, 2013; Badre and Nee, 2018; Bhandari et al., 2018).

In investigating the functional organization of frontoparietal cortex with human fMRI, it may be more fruitful to think about the FPN as consisting of groups of neurons that work together to complete the task, and that may flexibly reconfigure themselves, as well as their connections to other parts of the brain, depending on the task (Miller and Cohen, 2001; Cole et al., 2013). Task-dependent reconfigurations can be examined using functional connectivity (FC) approaches. Here we use the term FC to refer to temporal co-fluctuations in activity measured in different brain areas (Friston et al., 1993). These co-fluctuations are often taken to be indicative of communication between brain areas. FC can be used to define functional networks, which are groups of brain areas whose activity co-fluctuates over time. The FPN is an example of such a network (Fox and Raichle, 2007). Networks may adjust, or reconfigure, their FC with other brain networks to best accommodate changing task demands (Cole et al., 2013). For example, FC between visual and frontoparietal brain networks may be stronger during a visual than an auditory attention task. Such reconfigurations can be thought of as reflecting a

modulation of input sensitivity in the networks, which changes the pattern of brain information flow.

Prior work in both humans and nonhuman primates indeed shows evidence for task-dependent network reconfigurations at the whole-brain level, with the FPN being more flexible than other networks in its connectivity to other parts of the brain (Miller and Cohen, 2001; Krienen et al., 2014; Gratton et al., 2018; Cole et al., 2021). These studies also suggest that task-dependent reconfigurations tend to be unique to the individual and predictive of individual differences in task activation and behaviour on higher-order cognitive tasks (Gratton et al., 2018; Cole et al., 2021; Schultz and Cole, 2016; Finn et al., 2015). While this work is starting to show the importance of network reconfigurations for human cognition, it leaves open if the FPN reconfigures itself *internally* to meet changing task demands. If we find evidence for task-dependent reconfigurations at a finer spatial scale within the FPN, this opens up a window into the network's internal functional organization. We address this question in the current thesis by analyzing open source fMRI data measured during the execution of a range of cognitive tasks from a large group of healthy young adults (Glasser et al., 2013, 2016).

1.2 Functional Magnetic Resonance Imaging

1.2.1 Principles of MRI

Magnetic Resonance Imaging (MRI) enables researchers and clinicians to obtain high resolution images of soft brain tissue in a non-invasive manner. MRI can be used to differentiate between brain tissue types, including normal and abnormal tissue, and therefore provides information about morphology and is used as a diagnostic tool to detect disease. The MRI scanner creates a strong magnetic field in which the radio frequency (RF) electromagnetic fields interact with the atomic nuclei in the body tissue, forming a series of stacked two-dimensional images.

MRI measures the nuclear magnetic resonance signal from atomic nuclei with an odd

number of protons and neutrons since they have an angular momentum and create small magnetic dipoles that spin around their axis of rotation. The scanner detects changes in precession from a large number of spins. Since water and fat contain a majority of the single protons found in the nuclei of hydrogen atoms, they are the most frequent dipoles in the brain.

In a normal environment without an external magnetic field, the magnetic fields of the spins in the human body are oriented randomly and have no net dipole. However, in the presence of the scanner's strong magnetic field, the spins will become aligned with the field and precess around the axis of the external magnetic field. The precession frequency of the protons is directly proportional to the strength of the external magnetic field and is referred to as the Larmor frequency. The Larmor frequency is the frequency at which a particular type of atom precesses at a particular field strength (i.e. it differs across tissue types and field strengths). It is the product of the gyromagnetic ratio of an atom and the magnetic field strength. Resonance occurs when the applied electromagnetic pulse has the same frequency as the proton's precession frequency resulting in excitation of the protons. Excitation occurs since the protons absorb the transmitted energy and causes the spins to flip from a lower energy state (parallel) to higher energy state (antiparallel). This enables the excited protons to rotate in phase, the spins create a net magnetic field. This creates a current flow which is detected by the MRI scanner (receiver coil).

There are also smaller magnetic field gradients formed in orthogonal directions which adjust the Larmor frequency, enabling detection of the signal origin.

Once the RF transmitter is turned off, the spin precession is not stable because the protons interact with the magnetic fields and begin to decay. As a result, these spin-spin interactions lead to varying local magnetic field strengths and slightly different Larmor frequencies, resulting in phase shifts between the precessing spins, which is referred to as dephasing or transversal relaxation.

T1 relaxation is referred to as spin-lattice relaxation. This is where the net magnetization

returns to the original value in the direction of the applied magnetic field. T2 relaxation is referred to as spin-spin relaxation. This is the dephasing of magnetization vectors because of differences in Larmor frequencies.

The T2* parameter (measurements of changing local magnetic field inhomogeneities) provides indirect measurements of local neuronal activity. A major factor contributing to local inhomogeneities is due to the presence of deoxygenated hemoglobin. T2* relaxation is shorter in deoxyhemoglobin than in oxyhemoglobin blood.

1.2.2 Principles of fMRI

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive technique employed in studying the activation of brain regions. MRI provides a static structural view of brain tissue whereas fMRI extends this function by capturing functional changes in the brain caused by neuronal activity. Neuronal activity consumes energy which is delivered by oxygen in the blood. Regions that are highly active in the brain display a high signal due to a higher ratio of oxygenated blood.

The blood consists of hemoglobin molecules which bind and deliver oxygen to the various tissues and organs in the human body. Hemoglobin tends to have varying magnetic properties based on whether it is bound to oxygen. Oxyhemoglobin is weakly diamagnetic and therefore has a small effect on magnetic fields. However, deoxyhemoglobin is strongly paramagnetic and therefore has a larger effect on magnetic fields. This difference in magnetic properties allows for an improved MR signal since the deoxygenated blood creates distortions in the magnetic field (causing the nuclei there to lose magnetization faster) which lead to an accelerated T2* relaxation time. When the local blood vessels have more oxyhemoglobin compared to deoxyhemoglobin, it results in a longer T2* relaxation time which produces a larger signal from that region.

When there is local brain activity, there is an increase in neuronal activity which utilizes

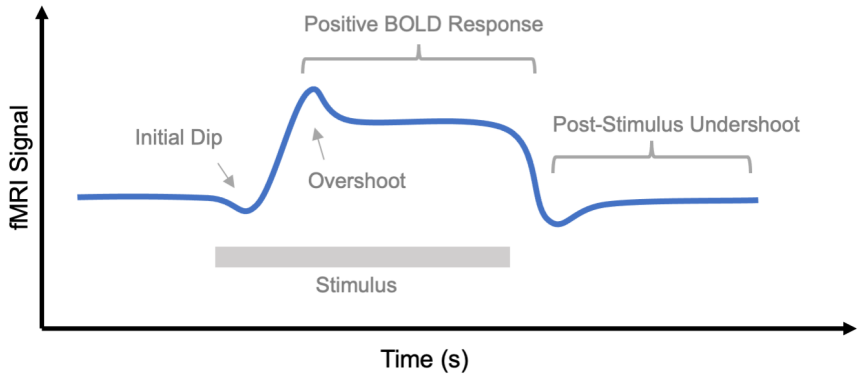

Figure 1.2: Hemodynamic response function. At the beginning of the stimulus presentation, there is an initial dip in fMRI signal, followed by a positive BOLD response and then a post-stimulus undershoot.

more oxygen and therefore increases the ratio of deoxygenated hemoglobin to oxygenated hemoglobin. The brain detects that there is an energy need when it detects metabolites (indicative of neural activity) after which it increases blood flow (functional hyperemia), which leads to an increase in oxygen delivery. After a short interval of approximately 3 seconds (after neural activity was induced), there is a strong increase in local blood flow, and associated changes in blood volume and oxygen extraction, all of which dynamically affect local deoxyhemoglobin levels, and therefore the measured signal (Buxton, 2012). The measured BOLD response shows a characteristic shape as it unfolds over time (Boynton et al., 1996). It peaks about 4-6 seconds after which it slowly goes back to baseline. Often, there is an initial dip and a post-stimulus undershoot, possibly due to temporal offsets between neural activity and the changes in blood flow, blood volume, and oxygen extraction. This has been referred to as the hemodynamic response function (HRF) (Figure 1.2).

Although fMRI does not directly measure the spiking activity of neurons, the blood oxygenation dependent levels (BOLD) are still tightly linked to neural activity. Studies have shown that changes in the BOLD signal reflect changes in the underlying local field potentials (LFPs) (Logothetis et al., 2001; Logothetis, 2008). LFPs are the summation of inhibitory and excitatory postsynaptic potentials from a large number of neurons in the recording site.

Therefore, LFPs are a reflection of information processing in local neural populations.

1.3 Resting State fMRI

Resting state fMRI indicates a scan which is performed while no stimulus is being presented to the subject. This method of scanning can be beneficial for understanding how different brain regions are functioning and communicating with one another during resting-state. In the early stages of fMRI research, many studies focused on performing fMRI scans while engaging the subject in a task. At this time, there were many concerns that the brain activity would vary unpredictably during resting state as it would be difficult to monitor what the subject would be thinking about. However, Biswal et al. (1995) measured brain activity from subjects at rest and demonstrated that different brain regions interacted with the motor cortex, suggesting that correlations in resting state activity can provide insight into the function of neural systems even though they are not actively engaged. Other studies have also shown that there are groups of brain regions that become more active during resting-state and activity decreases during a demanding cognitive task (task-negative networks) (Raichle et al., 2001; Shulman et al., 1997). These studies suggest that there is ongoing information processing and functional connections between regions during rs-fMRI. Overall, these studies have advanced the concept of resting state fMRI as a functionally relevant paradigm and paved the way for its acceptance in the neuroscience field.

1.3.1 Functional Connectivity

Functional connectivity (FC) is the study of relationships between areas of the brain based on co-varying fluctuations of activity (Fox and Raichle, 2007). FC can be determined by measuring the similarity between BOLD signals of various brain regions. If different areas have similar signals, it is taken to suggest that the regions are communicating with one another. Spontaneous fluctuations in the signal are used to investigate the FC during resting state since

the subject is not performing any cognitive task. Therefore these spontaneous fluctuations provide useful information about the similarity between brain regions and can be used to study brain connectivity.

It is important to acknowledge that there are different types of connectivity. For example, FC does not describe the directionality (or causality) of the signal. This type of connectivity is called effective connectivity. Also, FC does not indicate a direct physical connection between the interacting brain regions. This is referred to as anatomical (or structural) connectivity, which relies on white matter tracts. White matter tracts are known to be responsible for carrying information across brain regions and enable functional signals to be transferred between spatially separated regions. Recent studies have suggested a direct relationship between structural and FC (Hermundstad et al., 2014; Deco et al., 2014). These studies employ structural diffusion tensor imaging (DTI) which is an MRI technique that maps white matter tracts in the brain.

FC analyses come in multiple forms. Most approaches quantify the zero-lag similarity of regional activity fluctuations during rest and use this information to group brain regions into larger-scale functional networks. Functional networks can be estimated from a whole-brain connectivity matrix, which contains pairwise correlations between regional time courses, using clustering techniques (Fox and Raichle, 2007), or by applying independent component analysis (ICA) to the voxel time courses (Smith et al., 2009). Prior work using these approaches has consistently identified the FPN (Fox and Raichle, 2007; Smith et al., 2009). The brain networks discovered with FC analyses correspond to networks that become active during the execution of specific tasks (Smith et al., 2009). This suggests that the brain's functional organization is relatively stable across task-driven and resting states, supporting the idea of an inherent functional organization. Only a few fMRI studies have examined FC at a more fine grained spatial scale within the FPN (Waskom and Wagner, 2017; Stiers and Goulas, 2018). These studies also showed that task tuning is correlated with FC: the time courses of voxels with similar tuning profiles also co-fluctuate during rest.

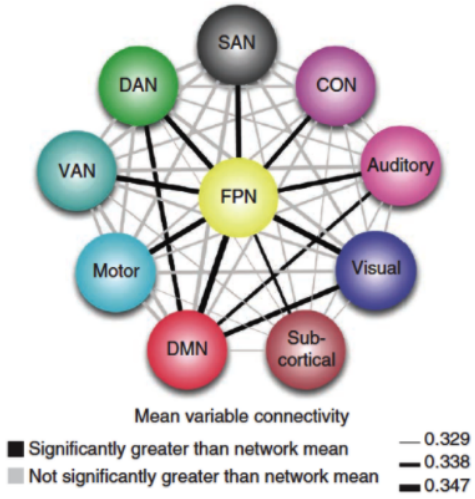

Figure 1.3: FPN mean variable connectivity. Figure above shows the pairwise mean variable connectivity between the ten networks shown. Black lines indicate that the mean variable connectivity between two networks is significantly greater than the mean, whereas grey lines indicate that they are not significantly greater than the mean. The thickness of the lines represents the strength of the variable connectivity. Overall, the FPN variable connectivity is the highest across the ten networks. (Adapted from Cole et al., 2013)

These studies leave open if and how strongly FC varies across tasks. For example, Cole et al. (2013) have shown that the FPN shows task-dependent changes in connectivity with other networks in the brain (Figure 1.3). Their study demonstrates that the FC pattern of the FPN shifted more significantly compared to other networks across a variety of tasks. The variability in FC was estimated at a whole-brain level and was referred to as the global variability coefficient (GVC). However, a gap in this research remains as the FC of the FPN has not been investigated at a finer spatial scale yet. This is one aspect that we aim to address in this thesis.

1.3.2 Limitations of Resting State fMRI

Although rs-fMRI has many applications and provides useful insights about brain function, there are several limitations that should be taken into consideration when interpreting resting-state FC results.

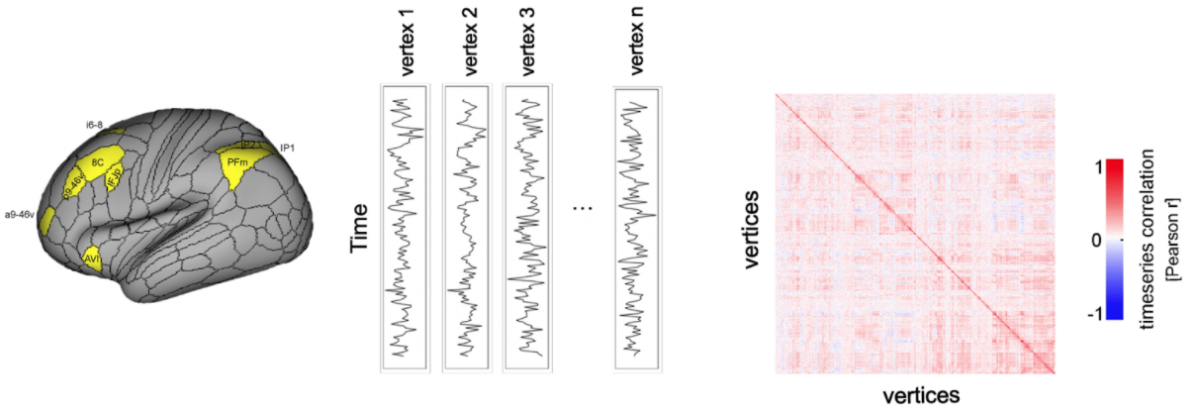

Figure 1.4: Extraction of functional timeseries data and FC matrices.

Spatial Limitations

The spatial resolution of fMRI refers to how well it distinguishes differences between nearby locations of activity in the brain, which depends on a few factors, the main two being the voxel size and the hemodynamic point spread function (PSF) (Uğurbil et al., 2003; Parkes et al., 2005). The PSF describes the spatial spread of the BOLD signal due to oxygenated blood entering the venous blood supply beyond the spatial location of neural activity.

In this thesis, we analyze 3T gradient-echo (GE) echo-planar-imaging (EPI) fMRI data with a voxel resolution of 2 mm isotropic (Barch et al., 2013), meaning that the slice thickness is equal to the in-plane resolution and the voxels are cubic. The hemodynamic PSF for 3T GE EPI fMRI data is approximately 3-4 mm (Parkes et al., 2005). Taken together, our data has a spatial resolution in the range of 2-4 mm. At this level, microstructural features such as cortical columns and sub-nuclei are not clearly identified. For example, if there are structures similar to cortical columns in the prefrontal cortex, these are expected to be on the scale of 0.7 mm (Goldman-Rakic, 1988). 3T GE EPI fMRI may therefore be sensitive to clusters of columns, but not individual ones. Despite these spatial limitations, fMRI is a useful tool for measuring brain activity across a larger spatial area compared to neural recording methods. Regardless of the measurement tool used, there will be a trade-off between spatial resolution and spatial

coverage (Sejnowski et al., 2014).

In addition to the above two factors, which are inherent to how fMRI samples neural activity, fMRI preprocessing methods may also reduce the spatial resolution of the data. For example, it is common to spatially smooth the data by applying a low-pass filter to the image (Goebel, 2007). Signals with larger spatial scales benefit from this approach because it increases the signal-to-noise ratio. While this may improve the power of statistical tests and comparisons between subjects, it may lead to reduced spatial resolution. In this thesis, we do not apply spatial smoothing to the data given that we aim to assess FC at a fine grained spatial scale within the FPN.

Temporal Limitations

Temporal resolution refers to the ability to distinguish changes in the fMRI signal over time. Given that neural activity during task execution can fluctuate rapidly (Hebart et al., 2018), it is important to measure the timing of brain activity with high precision. The temporal resolution of fMRI is limited by a few factors, including the hemodynamic response function (HRF; (Boynton et al., 1996)) and the repetition time (TR).

The HRF acts as a low-pass filter, making the BOLD fMRI signal predominantly sensitive to neural signals in the range of 0.01-1 Hz (He et al., 2008). Signals in this range may predominantly reflect neural communication at larger spatial distances (Matsui et al., 2016). This raises the question of whether the 3T HCP fMRI data analyzed here is suitable for assessing FC at a relatively fine grained spatial scale within the FPN. Prior reports from the monkey electrophysiology literature suggest the answer is yes. Kiani and colleagues identified functional subnetworks within the lateral prefrontal cortex within the millimetre range from a signal that was temporally broadband, measured in the range of 0.01 Hz to 16.67 Hz (Kiani et al., 2015). The detected functional structure was relatively stable across the measured frequency range. This provides support for the feasibility of identifying a similar structure from fMRI data in humans (Fox and Raichle, 2007).

The repetition time (TR) is the time it takes to acquire one functional image of the brain. In technical terms, is the length of time between corresponding consecutive points on a repeating series of pulses and echoes, where each series of pulses and echoes corresponds to one image. The TR defines the sampling rate of fMRI data. Commonly used TRs for acquiring a whole-brain image are $\sim$ 1 second, and the smallest possible TR with reasonable spatial coverage ($\sim$ 4 cm thick slab) and voxel size (2.5 mm isotropic) is $\sim$ 250 ms for 3T fMRI data (Lewis et al., 2016). The HCP 3T fMRI data were acquired with whole-brain coverage at a relatively short TR of 720 ms. This corresponds to a sampling rate of $\sim$ 1.4 Hz. The HCP protocols use multiband acquisition (multiband factor 8) (Moeller et al., 2010; Barch et al., 2013) which increases the temporal resolution by exciting and reading out multiple slices simultaneously. Taken together, the strongest signal in the data will be in the <1 Hz frequency band.

Noise: Slow Drifts

Slow drifts refer to low-frequency fluctuations in the baseline BOLD signal over time. The presence of these slow changes contributes to noise components in fMRI and they can affect the reliability and accuracy of the data and produce false positives. Therefore, these noise components are usually removed before data analysis (Yan et al., 2009). Slow drifts may arise from different sources. For example, some studies observed that they may be due to scanner instabilities (Smith et al., 1999), whereas others suggest that they may reflect spontaneous neuronal events related to fluctuations in metabolic events (Biswal et al., 1995; Greicius et al., 2003).

Noise: Head Motion

Head motion alters the uniformity of the magnetic field which is set at the beginning of the scan with shimming. This can lead to distortions in the signal. Head motion can also result in spin-history effects which refers to the change in excitation and signal intensity of a voxel as a

particular slice may correspond to a different part of the brain after some movement. This form of motion may be corrected using methods such as ICA or spin-history corrections.

Noise: Physiological Noise

The 0.01 - 1 Hz frequency band that fMRI is most sensitive to contains physiological noise components, driven by cardiac and respiratory rhythms that produce periodic fluctuations approximately in the 1-1.3 Hz and 0.25-0.45 Hz range, respectively. Relative changes in the cardiac rhythm and respiration can additionally impact other factors that impact the BOLD signal such as blood flow, blood oxygenation, blood CO2 levels and vasodilation. Since activity fluctuations across the brain are tightly coupled to the cardiac and respiratory variations, it may be difficult to separate this noise from the neural signal (Chen et al., 2020; Caballero-Gaudes and Reynolds, 2017). If the sampling rate is fast enough, i.e. more than twice the frequency of the noise which is approximately greater than 1 Hz for respiratory noise and greater than 2 Hz for cardiac noise, it may be possible to characterize and minimize these confounds (Agrawal et al., 2020). However, fMRI usually does not sample the signal at a frequency high enough to directly measure these fluctuations. Since the heart rate and respiration are usually unmodeled, these fluctuations contribute to the measured time course through aliasing and lead to temporal autocorrelation in the signal.

Reducing Noise

Slow drifts can be attenuated by using high-pass temporal filtering. This method removes low-frequency trends from the data. Since the head is a rigid body, the displacement of the voxels in space can be expressed in terms of translation and rotation along each of the three axes. In order to correct for head motion, prepossessing includes volume realignment where the displacement of the head is estimated and the fMRI volumes are realigned.

The noise components discussed above may be regressed out using ‘nuisance regressors’. Physiological noise can be identified by using a pulse oximeter to measure the cardiac rhythm

and a respiratory belt can measure the respiratory rate. Alternatively, time courses of the white matter and cerebral spinal fluid (CSF) may be regressed out as these time courses include physiological or other noise components but not neural signals. This is the approach that we followed in this thesis. Although physiological regressors may be considered to be closer to the data, white matter and CSF regressors also reflect these noise components substantially (Cole et al., 2019). For example, the cardiac rhythm contributes to a large portion of the variance around large vessels and in areas with CSF (Dagli et al., 1999).

Noise may also be removed using ICA-based methods (Beckmann and Smith, 2004). One commonly used method is ICA-FIX, which first identifies network components from resting-state data and then applies an automatic algorithm that identifies and removes noise components from the data (Salimi-Khorshidi et al., 2014). The HCP includes fMRI data preprocessed with ICA-FIX in later releases.

Although these these techniques are beneficial in reducing noise, they carry the risk of also removing some signal of interest. Also, these measurements are not perfect since noise may be present in other frequency bands and may remain unidentified. Thus, we always hope to remove more noise than signal.

1.4 Task fMRI

1.4.1 Introduction

While rs-fMRI is useful for many applications such as characterizing FC, the researcher has relatively weak control over the cognitive state of the subject. Even though there are no stimuli or tasks being presented to subjects during rest, there are still many activities occurring in the brain during resting state. The brain is constantly engaged in many distinct tasks that can range from regulating homeostasis of the body to memory consolidation of behaviour or previous thoughts (Bijsterbosch, 2017). As a result, the brain is continuously engaged in mental activities beyond

the researcher's control while the subject is in a 'resting state'. While this may be beneficial for understanding the functional organization of the brain at rest, we should also engage subjects in concerted tasks to assess brain activity and functional organization in response to external events and goals.

Task-state fMRI is beneficial for multiple reasons as it facilitates research on an array of cognitive tasks. It enables researchers to identify activity patterns, functional relationships between brain regions, and residual activity fluctuations that occur during different cognitive tasks. Task state fMRI is versatile because researchers can present a broad range of sensory stimuli, including auditory, tactile, and visual stimuli, as well as complex cognitive tasks of varying difficulty (Sereno et al., 1995; Formisano et al., 2003; Fedorenko et al., 2013). However, these applications present their own unique challenges including difficulty with selection and design of cognitive tasks for detecting functionally specialized brain regions, especially in the FPN (Lorenz et al., 2018).

Typically, experiments are designed as either a block, event-related or mixed block/event related format (Henson, 2007). Block designs were developed earlier and are a relatively simple approach as it involves presenting consecutive stimuli as a series of epochs or blocks. There are alternating on-off periods where the stimulus is presented (i.e. on-period) and where it is removed to allow a state of rest (i.e. off-period). An advantage of block designs is that the fMRI signal increases in response to a longer stimulation period (Boynton et al., 1996). Since the events are concentrated within task blocks, it enables block designs to have better power than event-related designs. However, since block designs only present one condition in each block, it is not possible to randomize the stimulus conditions. This means that subjects may be able to predict the next stimulus condition, which may not always be desired.

Event-related designs differ from block designs in that they evoke transient rather than continuous task-based neuronal activity. An underlying assumption of an event-related design is that the neural activity of interest occurs for discrete intervals. The stimuli that lead to the

neural activity are known as events. Event-related designs often require rapid and randomized subject responses. This prevents subjects from adopting cognitive strategies or heuristics while performing the task. It also permits a post hoc sorting of events into correct and incorrect trials, which may be beneficial for further analyses. Slow event-related designs also enable estimation of the HRF after each stimulus presentation. This is beneficial because by identifying the timing and properties of the HRF, inferences can be made about the relative timing of neural activity during distinct processes of a trial (Formisano and Goebel, 2003). The main disadvantage of event-related designs is their reduced statistical power relative to block designs. The HCP tasks therefore used a block design (Barch et al., 2013).

In order to analyze the task-state fMRI data, we use a general linear model (GLM). In general, a GLM is used for modelling the signals obtained from the data. The purpose of a multiple regression analysis is to find the linear combination of the regressors that produces a timeseries that is "best fit" to the data and explains as much variance as possible. In a GLM, data from one voxel is modelled as a linear combination of a model (X) consisting of a set of regressors (X1 and X2 - blue lines in Figure 1.5), which are the independent variables. The regressors may describe variance in the data, such as the BOLD signal. The dependent variable can be the preprocessed BOLD data. The output is a beta value (β) which represents amplitude and is calculated for each regressor in the model (i.e. β_1 for X_1 and β_2 for X_2). It represents the contribution of the regressor to the overall linear combination of regressors. The component of the data that cannot be explained by the regressors is the error, also known as the residual activity or noise. This is the difference between the data and the line of best fit. The residuals can be calculated by subtracting the sum of each regressor multiplied by its beta value from the data. The model with the best fit is the one with the smallest residuals. In conventional approaches, the residuals are considered to be noise and tend to be discarded.

Task-based fMRI has been used widely to map the functional organization of the brain, and has especially been successful at doing so at the whole-brain level (Smith et al., 2009) and for sensory systems (Sereno et al., 1995; Kanwisher et al., 1997; Formisano et al., 2003). Although

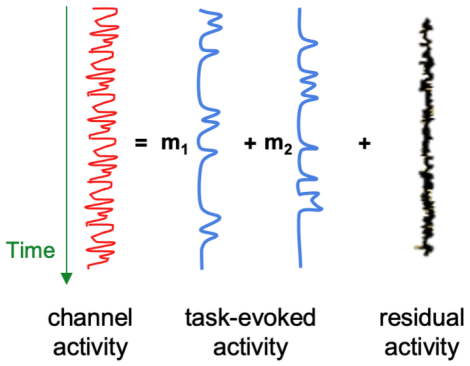

Figure 1.5: Schematic of general linear modelling of fMRI data. Red indicates the BOLD signal time course, blue indicates predicted BOLD signal by each regressor in the model, and black indicates the residuals unexplained by the model.

resting-state FC and task-based fMRI are quite different approaches, there is evidence in the literature that the two converge on what they infer about the brain's functional organization (Smith et al., 2009; Greicius et al., 2003; Waskom and Wagner, 2017; Haak et al., 2018). For example, Smith and colleagues showed that major brain networks, including the FPN, correspond between rest and task (Smith et al., 2009). In fact, resting-state FC can be used to predict task-based activations (Tavor et al., 2016; Cole et al., 2016).

1.4.2 Functional Connectivity

FC analyses are most commonly applied to resting-state fMRI data, but can also be applied to task-based fMRI data (Cole et al., 2019). The analyses can either be applied to the full time courses as measured during task performance, or separately to the task-predicted and residual components of the time courses (Fox and Raichle, 2007; Cole et al., 2019, 2021; Kiani et al., 2015). FC analyses on the task-predicted time courses reveal structure in co-fluctuations between voxels that can be explained by the task design. The same analyses applied to the residual time courses reveal structure in co-fluctuations between voxels that cannot be explained by the task design. The residual time courses can be thought of as the task-based 'equivalent' of resting-state time courses.

Functional structure in the residuals has been shown to be relatively similar to functional structure revealed by resting-state data (Greicius et al., 2003; Cole et al., 2014; Gratton et al., 2018). This suggests that there is an inherent functional organization that exists during both rest and task. However, there are differences between the two as well and these differences appear to be meaningful for cognition (Schultz and Cole, 2016; Cole et al., 2021). This suggests that network configurations are adjusted to best accommodate the task at hand (Varela et al., 2001; Cole et al., 2013). Recent work in humans suggests that task-driven network reconfigurations are largely unique to the individual as opposed to shared across individuals (Gratton et al., 2018).

1.4.3 Limitations of Task State fMRI

Many of the limitations described for resting-state fMRI also apply to task state fMRI such as noise, spatial and temporal limitations. However, task state fMRI can be difficult to regulate as task performance across subjects and studies may be inconsistent. As a result, testing must be designed and implemented systematically to avoid poor test-retest variability and reduce confounds. For example, tasks must be considerably engaging so that the subject's attention is maintained throughout the experiment. Also, the tasks should strive to minimize variations in behavioural performance and cognitive strategy since some subjects may use heuristics (i.e. selecting the option that seems most familiar to them), whereas others may adopt a more analytic approach (i.e. comparing the costs and benefits of every option).

1.5 Human Connectome Project

Thousands of neuroimaging studies over the past two decades have identified major functional brain networks, such as the FPN (Fox and Raichle, 2007; Smith et al., 2009; Duncan, 2010; Cole et al., 2013). Traditionally, the majority of these studies have focused on localizing brain function at a relatively coarse spatial scale using a single source of data, for example resting-state

fMRI. However, recently, researchers have developed improved multi-modal pipelines that take advantage of the different properties of the brain to provide rich information about its various structures and functions (Glasser et al., 2016).

The data used in this thesis are obtained from the Human Connectome Project (HCP) which is an NIH-funded collaboration between the University of Minnesota and Washington University (Van Essen et al., 2012). The specific dataset that we use is the S1200 Subject Release which has data publicly available from approximately 1200 young healthy adults. The HCP follows a surface-based alignment instead of a traditional volume-based approach (Glasser et al., 2013). This is motivated by studies which show that cortical folds are unreliable landmarks for brain alignment (Coalson et al., 2018). Instead, it is more reliable to use cortical features such as myelin maps and areal patches defined based on resting state FC (referred to as multimodal surface matching (MSM)) (Coalson et al., 2018). Since the HCP follows a surface-based alignment, our analyses were performed on the cortical surface using vertices, rather than voxels. Studies show that a vertex-wise approach may yield better multivariate classification performance than voxel-wise information decoding (Oosterhof et al., 2011).

The HCP S1200 Subject Release consists of 3T and 7T data. We chose to analyze the 3T data since there are seven tasks that engage subjects in cognitive tasks. The seven tasks are: working memory, gambling, motor, language, social cognition, relational processing, and emotion processing. We did not analyze the 7T data in this thesis because it only consists of two tasks: retinotopic mapping and movie-watching. Past studies have also shown the utility of 3T task data in analyzing cognitive processes in the FPN (Assem et al., 2020; Cole et al., 2016).

1.5.1 Glasser Parcellation

The HCP data comes with the multi-modal Glasser parcellation (Glasser et al., 2016), which we used for defining the FPN of interest. A brain parcellation is a delineation of spatial non-overlapping parts of the brain (parcels, also referred to as regions of interests (ROIs)). The

parcellation is created by a variety of imaging techniques which take advantage of the similarities between regions such as cytoarchitecture, structural or FC and task-related activity. Researchers have created a variety of maps based on these different features. Despite recent advances, there is still no universal atlas of the human brain due to the myriad of challenges involved. Nevertheless, brain parcellations are useful in providing information about the organization and features of the brain and allow streamlined methods of analyzing brain function (Glasser et al., 2016; Arslan et al., 2018).

Different ways of creating a brain parcellation have been developed, such as an atlas-based parcellation, connectivity-based parcellation and multimodal parcellations (Arslan et al., 2018; Zhi et al., 2021). An atlas-based parcellation subdivides the brain by employing an anatomical template. This approach has shown to be limited because it does not adequately address the variations between subjects since the divisions are based on an averaged dataset (Tzourio-Mazoyer et al., 2002; Fischl, 2004; Desikan et al., 2006). Another method is connectivity-based parcellation, which divides the brain regions by grouping voxels that are similar in their connectivity patterns measured from fMRI (Gordon et al., 2016; Thomas Yeo et al., 2011). Another parcellation method employs multimodal techniques which partition the brain by employing information from a variety of sources. Multimodal parcellations have shown to be a promising technique in defining brain regions more accurately (Glasser et al., 2016).

In this work, we use the Glasser surface-based parcellation which is created based on multimodal techniques (Glasser et al., 2016) (Figure 1.6). It uses four different neuroimaging modalities. The atlas identifies 180 symmetric cortical parcels per hemisphere (360 total ROIs) based on architectural, functional, connective and topographic features. This parcellation was created from the multimodal dataset from the HCP and an objective semi-automated neuroanatomical procedure. There were notable differences between neighbouring areas in terms of microstructural architecture, functional specialization, connectivity to other regions, and internal topographic organization. Due to the multi-modal approach used in its development, this parcellation is different from most others. The architectural measurements included cortical

thickness and cortical myelin content which was computed from the T1 and T2-weighted structural MRI. Cortical function was assessed using task state fMRI contrasts for seven different tasks. FC was quantified from resting-state fMRI.

The Glasser parcellation was constructed through rigorous automated algorithms and manual neuroanatomical approaches (Glasser et al., 2016). Initially, an observer-independent semi-automated neuroanatomical approach was created to delineate the boundaries, which were then reviewed by expert neuroanatomy specialists. The cortical areas in subjects were subsequently identified using an automated algorithm. The multimodal parcellation was achieved by training a machine-learning classifier; it was repeated in new subjects and studies with a high degree of reproducibility. After these steps, the HCP research group was able to achieve an average subject's parcellation by averaging the results from data of 210 subjects in the healthy young adult dataset.

1.6 Thesis Outline

1.6.1 Thesis Question

In this thesis, we ask whether there is a fine-grained functional organization within the human FPN. A better understanding of the network's internal functional organization will contribute to a better understanding of its role in higher-order cognition. Higher-order cognition requires flexible integration of information across domains. The FPN may therefore adhere to somewhat different organizational principles than sensory cortices, which show orderly and stable maps of functional selectivity to stimulus features. Instead, we hypothesize that the FPN contains neurons that multi-task and that flexibly reconfigure themselves to best meet the current task demands. The task-dependent reconfigurations are expected to provide clues about the internal functional organization of the network. We address the research question by examining task-dependent reconfigurations of FC within the human FPN as measured with fMRI. This approach

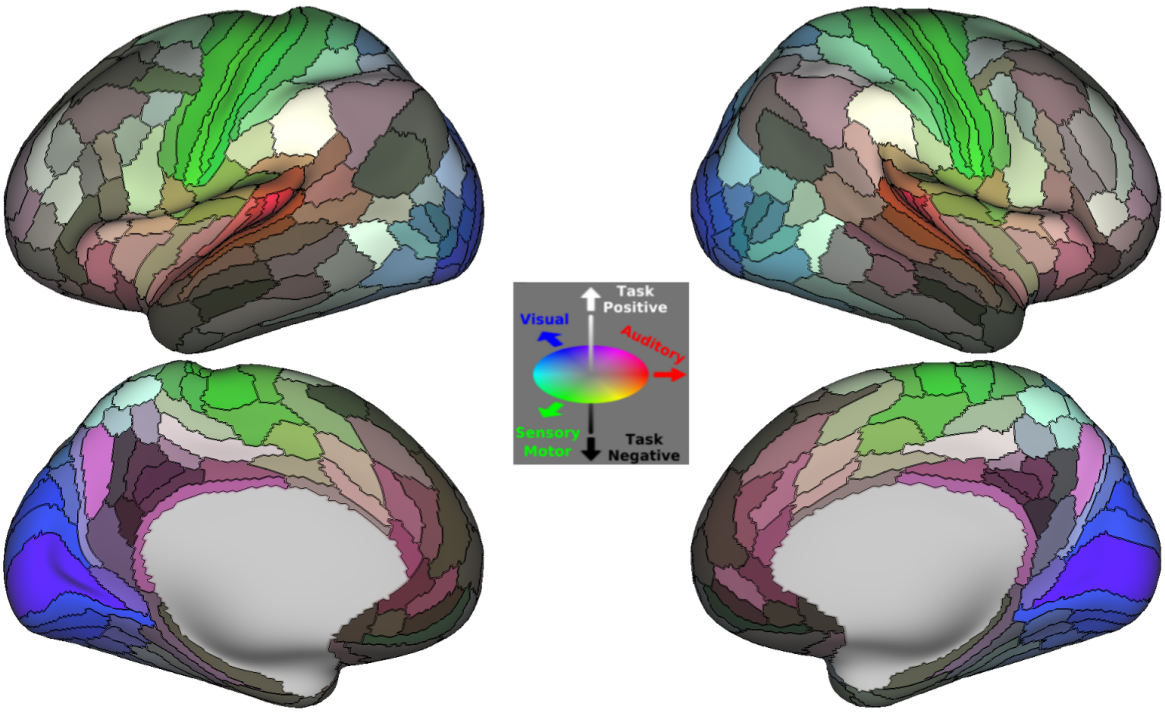

Figure 1.6: Glasser Parcellation. Colour scheme represents the likelihood of a region responding during the execution of a particular cognitive task (Glasser et al., 2016).

assumes that the reconfigurations happen at spatial and temporal scales accessible to 3T fMRI (Kiani et al., 2015).

1.6.2 Thesis Objectives

Our overall aim of mapping the internal functional organization of the human FPN translates into three objectives. First, we will characterize the functional organization of the human FPN using FC approaches. To visualize the spatial structure of the functional organization, we will map the organization onto the cortical surface using dimensionality reduction techniques. Second, we will test whether FC profiles within the FPN are replicable within individuals over time, and whether they are more consistent within than between individuals. If we confirm these hypotheses, this provides evidence for the existence of a fine-grained functional organization within the FPN that is (partially) unique to the individual. Third, we will assess if the individual FC profiles are affected by task. If they are affected by task, this provides evidence for task-dependent internal reconfigurations of the FPN. In Chapter 4, we will elaborate on how these reconfigurations may shed further light on the internal organization of the FPN.

Chapter 2

Methods

2.1 Human Connectome Project

2.1.1 Participants

We analyzed 3 Tesla (T) fMRI data from the Washington University of Minnesota Consortium of the Human Connectome Project (HCP) database. This open-source database contains both resting-state and task fMRI data for 1206 healthy young adults. The dataset is available in the Human Connectome Project repository (https://www.humanconnectome.org/study/hcpyoung-adult). We only included individuals with complete data sets yielding a total of 924 individuals for analysis (22-37 years old) (mean age 28 +/- 3.72 standard deviation, 487 females, 437 males). Of those individuals, 100 participants were genetically unrelated, while the remaining 824 participants consisted of sets of siblings, including dizygotic and monozygotic twins.

2.1.2 Experimental Design

Each participant performed seven tasks twice in the scanner over two days. The task-based fMRI data available from the HCP include: Working Memory, Gambling, Motor, Language, Social

Cognition, Relational Processing, and Emotion Processing. The first three tasks were performed on one day, the remaining four tasks were performed on another day. The experimental design was described in Barch et al. (2013), where further details can be found. We summarize the most relevant experimental design features below. Stimuli were projected onto a computer screen behind the participant's head within the imaging chamber. The screen was viewed using a mirror positioned approximately 8 cm above the subject's face.

Working Memory: Participants were presented with two runs of 8 task blocks (10 trials of 2.5s each, for 25s) for each run and 4 fixation blocks (15 s each). Within each run, 4 blocks used a 2-back working memory task and 4 blocks used a 0-back working memory task. A 2.5s cue indicated the task type (and target for 0-back) at the start of the block. On each trial, the stimulus was presented for 2 seconds, followed by a 500 ms inter-task interval (ITI). Block stimuli consisted of images of places, tools, faces and body parts. In each block there were 2 targets, and (in the case of the 2-back task) 2–3 non-target lures (repeated items in the wrong n-back position, either 1-679 back or 3-back).

Gambling: This task was adapted from Delgado et al. (2000). Participants were presented with a card guessing game where they guessed the number on a mystery card in order to win or lose money. Participants were presented with blocks of 8 trials that are either mostly reward (6 reward trials pseudo randomly interleaved with either 1 neutral and 1 loss trial, 2 neutral trials, or 2 loss trials) or mostly loss (6 loss trials pseudo randomly interleaved with either 1 neutral and 1 reward trial, 2 neutral trials or 2 reward trials). In each of the two runs, there were 2 mostly reward and 2 mostly loss blocks, interleaved with 4 fixation blocks (15s each).

Motor: To map motor areas, participants were instructed to either tap their left or right fingers, squeeze their left or right toes, or move their tongue. Each block of movement started with a 3s cue followed by 12s of 10 movements. There were three 15s fixation blocks per run. There were two runs, each consisted of 13 blocks with 2 tongue movements, 4 hand movements,

and 4 foot movements.

Language: The language task consisted of two runs interleaved with 4 blocks of the story task and 4 blocks of the math task, where each block was approximately 30s. Participants were presented with a story task and a math task. The story blocks consisted of short auditory stories followed by a 2-alternative forced-choice question that asked participants about the topic of the story. The math blocks presented trials aurally and instructed participants to perform addition and subtraction problems presented as a 2-alternative forced-choice question.

Social Cognition: The social cognition task consisted of two task runs with 5 video blocks and 5 fixation blocks (15s each). Participants were presented with brief video clips of objects that interacted or moved randomly. After each video clip, participants were asked whether the objects had a mental interaction.

Relational Processing: In this task, there were 2 runs, where each one consisted of 3 relational condition blocks, 3 control matching blocks, and 3 16s fixation blocks. In the relational condition, there were 5 trials per block where stimuli were presented for 3500ms with a 500ms ITI. In the matching condition, there were 5 trials per block where stimuli were presented for 2800ms with a 400ms ITI. Each block was 18s in duration. In the relational condition, participants were presented with 2 pairs of different shapes with varying textures. They were asked how the pairs differed in terms of shape or texture. In the control matching condition, participants were shown two objects at the top of the screen and one object at the bottom of the screen and asked whether either matched in terms of shape or texture.

Emotion Processing: This task was adapted from Hariri et al. (2002). The emotion processing task consisted of two runs, each with 3 face blocks and 3 shape blocks, and ending with 8s fixation blocks. Each trial started with a 3000ms task cue (shape or face) and was presented in blocks of 6 trials of the same task where the stimulus was presented for 2000ms and a 1000ms ITI. Overall, each block was 21s in duration including the cue. Participants were presented with blocks of trials where they performed a matching task on images of fearful faces

or neutral shapes.

2.1.3 MRI Measurements

Participants took part in two measurement sessions on separate days. They performed two resting-state runs and three of the seven tasks in one session, and two resting-state runs and four of the seven tasks in the other session. For each task, data were acquired for two runs. Within each session, acquisitions alternated between phase encoding in a right-to-left (RL) direction in one run and phase encoding in a left-to-right (LR) direction in the other run. Data were acquired on a 3-Tesla Siemens Siemens Skyra "Connectom" MRI scanner with a 32-channel RF receive head coil. Data were acquired using 2D gradient-echo echo-planar-imaging (GE EPI) with whole-brain coverage and the following parameters: repetition time (TR) = 720 ms, echo time (TE) = 33.1 ms, multiband factor = 8, 72 slices at oblique axial orientation, 2.0 mm isotropic voxels. Each resting-state run consisted of 1200 volumes (14 min and 33 s per run). Participants kept their eyes open with a relaxed fixation on a bright cross-hair on a dark background projection. Task runs were 3-5 minutes in duration, depending on the task. Cardiac and respiratory signals were measured using a standard Siemens pulse oximeter placed on the fingertip and a breathing belt placed around the chest, with a 400 Hz sampling rate. Further details can be found in past papers (Smith et al., 2013; Van Essen et al., 2012; Uğurbil et al., 2013; Sotiropoulos et al., 2013).

2.2 fMRI Preprocessing

We used preprocessed data made available on the HCP ConnectomeDB data management platform. These data were processed using the minimal preprocessing pipeline for the HCP dataset, which is described in detail in Glasser et al. (2013). Below, we describe the most relevant steps of this pipeline.

Structural data preprocessing was performed using the *PreFreeSurfer*, *FreeSurfer* and

PostFreeSurfer HCP pipelines. In brief, structural images (T1w and T2w) were corrected for gradient nonlinearity distortion and registered to the Montreal Neurological Institute (MNI) space template using an initial rigid-body alignment (6 degrees of freedom) followed by linear (FLIRT) and non-linear (FNIRT) volume registration. The structural images were then warped back into native volume space, and corrected for readout distortion and B1 bias. The T1w structural image in native volume space was passed on to Freesurfer's recon-all algorithm for segmentation. Freesurfer outputs, including white matter and pial surface definitions, were converted to NIFTI format. The native-mesh surfaces were subsequently registered to the Conte69 population-average surfaces, which brings the surfaces into alignment with the Freesurfer standard surface mesh. The HCP processing pipeline uses a downsampled version of this standard mesh, which is referred to as 32k_fs_LR space. This space has an average vertex spacing of 2 mm on the midthickness surface. The data analyzed in this thesis reside in the 32_fs_LR space.

Functional data preprocessing was performed using the *fMRIVolume* and *fMRISurface* HCP pipelines. In brief, these pipelines implement gradient nonlinearity distortion correction, motion correction (6 degrees of freedom: 3 translation and 3 rotation parameters), EPI distortion correction, EPI to T1w registration using a combination of FSL's FLIRT and FreeSurfer's BBRegister tools, native volume to MNI nonlinear registration, and mild high-pass (200s) temporal filtering. Resting state and task state functional images were mapped from volume to surface space with ribbon-constrained volume to surface mapping. A standard CIFTI grayordinate space was created by including both subcortical voxel data and cortical surface data. Grayordinates consist of surface vertices and subcortical voxels contained in a CIFTI file. Data were smoothed with a 2mm FWHM kernel in the grayordinate space that did not mix data across gyral banks for surface data and areal borders for subcortical data.

2.3 fMRI Analyses

2.3.1 Extraction of Residual Time Courses

First, we selected a subset of 10 parcels from the Glasser parcellation. The selected parcels are core regions of the FPN as assessed by Assem et al. (2020) (see Figure 1.1). The selected parcels show stronger activation with increasing working memory, relational processing, and arithmetic task demands. Only parcels showing these effects for at least two of the three task demands were included. The 10 parcels include seven frontal parcels (8C, IFJp, a9-46v, p9-46v, i6-8, AVI, 8BM) and three parietal parcels (IP1, IP2, PFm) (details listed in table A.1). We extracted time series data from these parcels for both left and right hemispheres (3745 vertices in total) and for resting as well as task states.

We then prepared nuisance regressors for both resting-state and task data and used these to model variability due to head motion and physiological noise using linear regression. We used 12 motion regressors per run, consisting of 3 translation and 3 rotation regressors, plus their first-order derivatives. The physiological noise was modelled with white matter and cerebral spinal fluid (CSF) regressors. These regressors were created by extracting time course data from voxels covering the white matter and ventricles, respectively, and by averaging these data across voxels. This process yielded one white matter regressor and one CSF regressor per participant per run. We also included a linear trend regressor. This yielded at total of 15 nuisance regressors per run. The first 5 frames of each run were removed before fitting the GLM to the data. Later versions of the pipeline have implemented ICA+FIX. In our analyses, we did not use this approach. Instead we used the nuisance regressors based on the white matter and CSF time courses.

We did not perform global signal regression (GSR) since this may introduce negative correlations between time series (Murphy et al., 2009). This would bias the FC analyses and could introduce artificial differences in FC between resting and task states (Cole et al., 2019).

We did not perform low-pass temporal filtering because there may be task signals at higher frequencies (e.g. relative to slow resting-state fluctuations).

Next, we prepared task regressors based on the seven tasks, modelling 24 task conditions in total. The emotion processing task consisted of two conditions, face and shape. The gambling task consisted of two conditions as well, punishment and reward. The language task consisted of a story condition and a math condition. The motor task had six conditions: cue, right hand trials, left hand trials, right foot trials, left foot trials, and tongue trials. For the relational processing task, the conditions were matched objects or unmatched objects. The social cognition task consisted of two conditions, interacting objects (theory of mind) or objects randomly moving. For the working memory task, there were four conditions for each 0-back and 2-back load: body, face, tool, place.

We created predictors by applying a finite impulse response (FIR) function to the task timings. We created one regressor per TR, resulting in 25-50 predictors per condition. Since all tasks used block designs, each TR for each block was modelled separately for each task condition based on the FIR model, with a lag extending up to 25 TRs after task block offset. The FIR approach allows us to estimate the HRF as precisely as possible and thus regress out task-predicted activation as best as possible. A FIR function does not assume an HRF and therefore it is better than the standard HRF because it avoids including task related information in the residuals. Prior work shows that the FIR model reduces both false positives and false negatives in the identification of FC estimates (Cole et al., 2019). This is due to the FIR model's ability to flexibly fit the average task-evoked response as it unfolds over time. Removing the averaged evoked response of a task condition is useful for separating the task-predicted from the residual activity, each of which may have different FC profiles (Norman-Haignere et al., 2012).

We fit the GLMs to the resting-state and task time series data of each individual vertex within the FPN. For resting-state data, the GLMs consisted of nuisance regressors only. For task data, the GLMs consisted of both nuisance and task regressors. After fitting, we computed the

residuals and used these for further analyses.

2.3.2 Estimation of Functional Connectivity Matrices

We estimated FC by computing pairwise Pearson correlation coefficients between the residual time series of the frontoparietal vertices. We stored the pairwise correlations in a FC matrix (Figure 2.1a). We computed these matrices for each resting-state and task run, yielding a total of 18 matrices per subject for further analysis. These 18 matrices were based on four resting-state runs and two runs for each of the seven tasks. The correlation matrices serve as a starting point for examining the network's internal functional organization.

2.3.3 Visualization of Functional Connectivity on the Cortical Surface

Dimensionality refers to the number of coordinate values used to identify a point in space. When data are collected in a high-dimensional space, for example a space spanned by the number of time points in the fMRI data, they are challenging to interpret. In order to better visualize the data, we can perform dimensionality reduction. Dimensionality reduction is a statistical method which converts the high-dimensional data into a low-dimensional space while capturing the dimensions where most of the variance occurs (Cunningham and Yu, 2014; Williamson et al., 2016). These methods have revealed evidence of the neural mechanisms underlying cognitive functioning such as integrating sensory information and decision making (Kaufman et al., 2014; Mante et al., 2013).

In our analyses, we use dimensionality reduction techniques to enable projection of the FC results onto the cortical sheet (Figure 2.1) (Kiani et al., 2015). By visualizing the results on the cortical sheet, we can get an impression of the spatial structure of the network's internal functional organization. We first converted the correlation coefficients in the FC matrices into correlation distances and then applied 2-dimensional (2D) multidimensional scaling (MDS) (criterion: metric stress) to the matrices. MDS is a nonlinear dimensionality reduction technique

that converts distances between points in a high-dimensional space to distances between points in a low-dimensional space with the least possible distortion (Borg and Groenen, 2005; Kruskal and Wish, 1978; Torgerson, 1958; Shepard, 1980). The 2D MDS representation consists of points (in our case vertices) on a plane, where smaller distances between vertices reflect a stronger positive correlation in their residual time courses. The MDS plots illustrate the time course correlations between vertices and facilitate visualization of the data structure.

To visualize FC on the cortical surface, we next colour coded the vertices according to their location in the 2D MDS space, with hue coding for polar angle and saturation coding for eccentricity (Figure 2.1b). Next, we projected these colours to the cortical surface and displayed them using Connectome Workbench (Figure 2.1c). Vertices with similar colours have similar time series; vertices with dissimilar colours have dissimilar time series. Given that MDS is a dimensionality reduction technique, it naturally induces distortions. To assess the degree of distortion, we computed the Pearson correlation coefficient between the distances in the original high dimensional space and the distances in MDS space. For displaying multiple MDS solutions, for example multiple task runs within the same subject, we used procrustes alignment to ensure the best possible alignment between the MDS solutions. We selected one MDS solution as the reference, and aligned the others to this reference.

2.3.4 Consistency of Functional Connectivity Within and Between Individuals

We first estimated the replicability of the FC matrices over time within individuals. For each individual, we computed the Pearson correlation coefficient between pairs of connectivity matrices. Pairs consisted of two different runs of the same type acquired on the same day, for example two resting state runs or two working memory runs. We computed pairwise correlations for the following mental states: rest day 1, rest day 2, and each of the seven tasks. We next estimated the consistency of the FC matrices between individuals. For each pair of individuals, we computed the Pearson correlation coefficient between pairs of connectivity matrices. Pairs

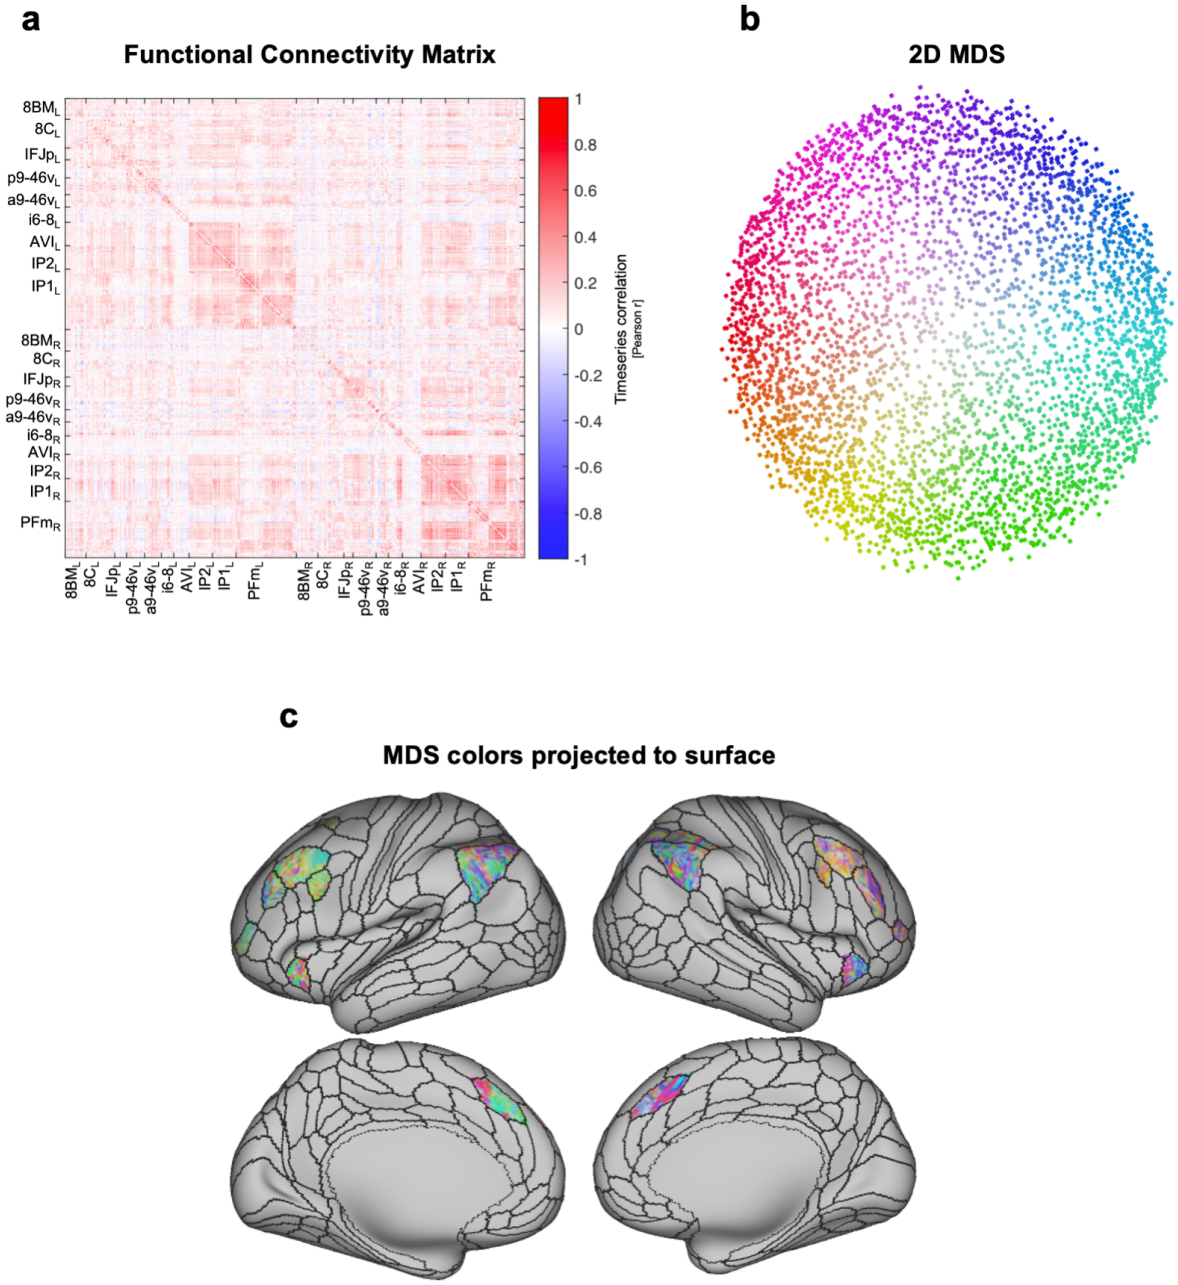

Figure 2.1: Visualizing FC within the FPN on the cortical surface. We performed 2D MDS on the FC matrices (3745 x 3745 vertices), coloured the vertices according to their location in the 2D MDS space, and projected the colours to the cortical surface. Similar colours indicate vertices with similar time series. This example is based on one resting-state run of one example subject. The Pearson correlation coefficient between the distances in the original space and the distances in the 2D MDS space is 0.65.

consisted of runs of the same type, one acquired in the first individual and the other acquired in the second individual. We computed pairwise correlations for the following mental states: rest day 1, rest day 2, and each of the seven tasks. We then averaged the within- and between-subject correlations across run pairs and subtracted the average between- from the average within-subject correlation.

We performed statistical inference on the difference in the average within- and between-subject correlations using randomization of the subject labels. We randomized the subject labels before averaging correlations across run pairs, simulating the null hypothesis of no difference in within- and between-subject consistency of FC. We performed 1,000 randomizations, each providing an estimate of the difference under the null hypothesis. If the actual difference fell within the top 5 percent of the simulated null distribution, we rejected the null hypothesis of no difference in within- and between-individual consistency of FC. We performed a one-sided test because we expected the consistency to be higher within than between subjects. We focused our analysis on the subset of 100 genetically unrelated subjects to prevent inflation of the between-subject correlations.

2.3.5 Consistency of Functional Connectivity Within and Between Tasks

To determine whether task affects FC within the FPN, we estimated the consistency of the FC matrices within and between tasks. For each individual, we computed the Pearson correlation coefficient between pairs of connectivity matrices from the same task or from different tasks. This yielded a 14-by-14 correlation matrix for each individual, with the 14 task runs on the axes (Figure 3.5a). We averaged the within- and between-task correlations across run pairs and subtracted the average between- from the average within-task correlation. We performed statistical inference on the difference in the average within- and between-task correlations using a one-sided paired t-test across individuals. We performed a one-sided test because we expected the consistency to be higher within than between tasks. We ran this analysis for all subjects (n = 924), for the related subjects (n = 824), and the unrelated subjects (n = 100).

Chapter 3

Results

3.1 Functional Connectivity is Replicable Within Individuals

We first computed FC matrices for the FPN during resting and task states in each individual (see Figure 2.1a). We then visualized the FC information on the cortical surface to get an impression of the spatial structure of the frontoparietal FC. We applied 2D MDS to the FC matrices, coloured the frontoparietal vertices according to their location in the MDS space, and projected the colours onto the cortical surface. Similar colours indicate similar residual time courses. Figure 3.1 shows results for an example participant for two resting-state runs acquired on the same day. The visualizations suggest that FC within the FPN is relatively consistent within individuals during resting state. These observations were qualitatively similar for the seven task states and for the other participants.

We next quantified these observations by assessing the replicability of FC matrices between repetitions of resting and task states within individuals. We hypothesized that FC is replicable within individuals over time. This was assessed by computing pairwise Pearson

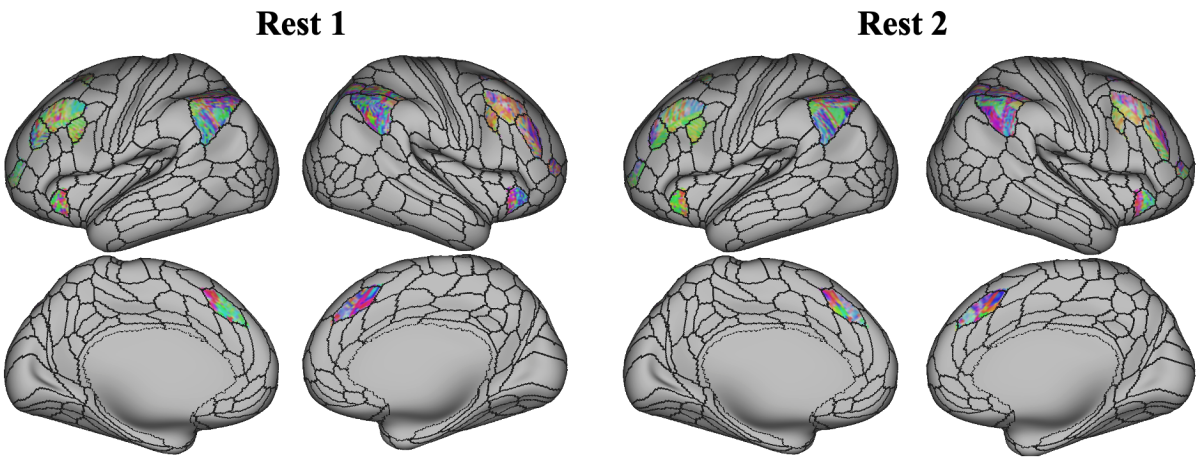

Figure 3.1: FC within the FPN visualized on the cortical surface. Vertices with similar colours exhibit similar time courses. Resting state data for run 1 and run 2 of one example subject are shown with a lateral and medial view of each hemisphere. FC appears relatively replicable over time within the same individual.

correlation coefficients between FC matrices for different runs of the same resting or task state, shown schematically in Figure 3.3a. The dark grey bars in Figure 3.3b show the average replicability of FC matrices within individuals for each task and for each day of resting-state data acquisitions. FC within the FPN indeed appears replicable within individuals. This is an important observation, because replicability of FC is a first necessary step for establishing that the FPN contains a fine-grained functional organization.

3.2 Functional Connectivity is More Consistent Within than Between Individuals

After establishing that frontoparietal FC is consistent over time within individuals, we assessed consistency of FC between individuals. Figure 3.2 visualizes FC on the cortical surface for two example participants for each of the two working memory task runs. These visualizations suggest that FC within the FPN is more consistent within than between individuals. These observations were similar for the other tasks, for resting-state data, and for other participant

pairs.

We quantified these observations by assessing the consistency of FC matrices between individuals for the same resting or task state. This was assessed by computing pairwise Pearson correlation coefficients between FC matrices for runs of the same resting or task state acquired in different participants, shown schematically in Figure 3.3a. The light grey bars in Figure 3.3b show the average replicability of FC matrices between individuals for each task and for each day of resting-state data acquisitions. We performed inference on the difference in within- and between-subject consistency using subject-label randomization tests (5,000 randomizations, one-sided test). For both resting-state days and all tasks, the difference was significant ($p<0.001$ for each comparison). The motor and WM tasks appear to show higher consistency within individuals than the other tasks. Overall, these results indicate that frontoparietal FC is more consistent within than between individuals. In other words, the internal functional organization of the FPN is partially individually unique. This is an important observation because it raises the possibility that individual variability in the network's internal organization may contribute to individual variability in cognitive function.

3.3 Functional Connectivity is More Consistent Within than Between Tasks

We next investigated whether the internal functional organization of the FPN is affected by task. If the FC matrices are fully consistent across tasks within an individual, this would suggest that the FPN does not internally reconfigure itself to accommodate changing task demands. Instead, any task effects would be predominantly driven by changes in input to the FPN (Cole et al., 2016). Alternatively, if the FC matrices are uncorrelated across tasks within an individual, this would suggest that the FPN completely reconfigures itself from one task to the next. To get an impression of the consistency across tasks, we visualized frontoparietal FC on the cortical surface for the seven tasks within individuals. Figure 3.4 demonstrates results for an example

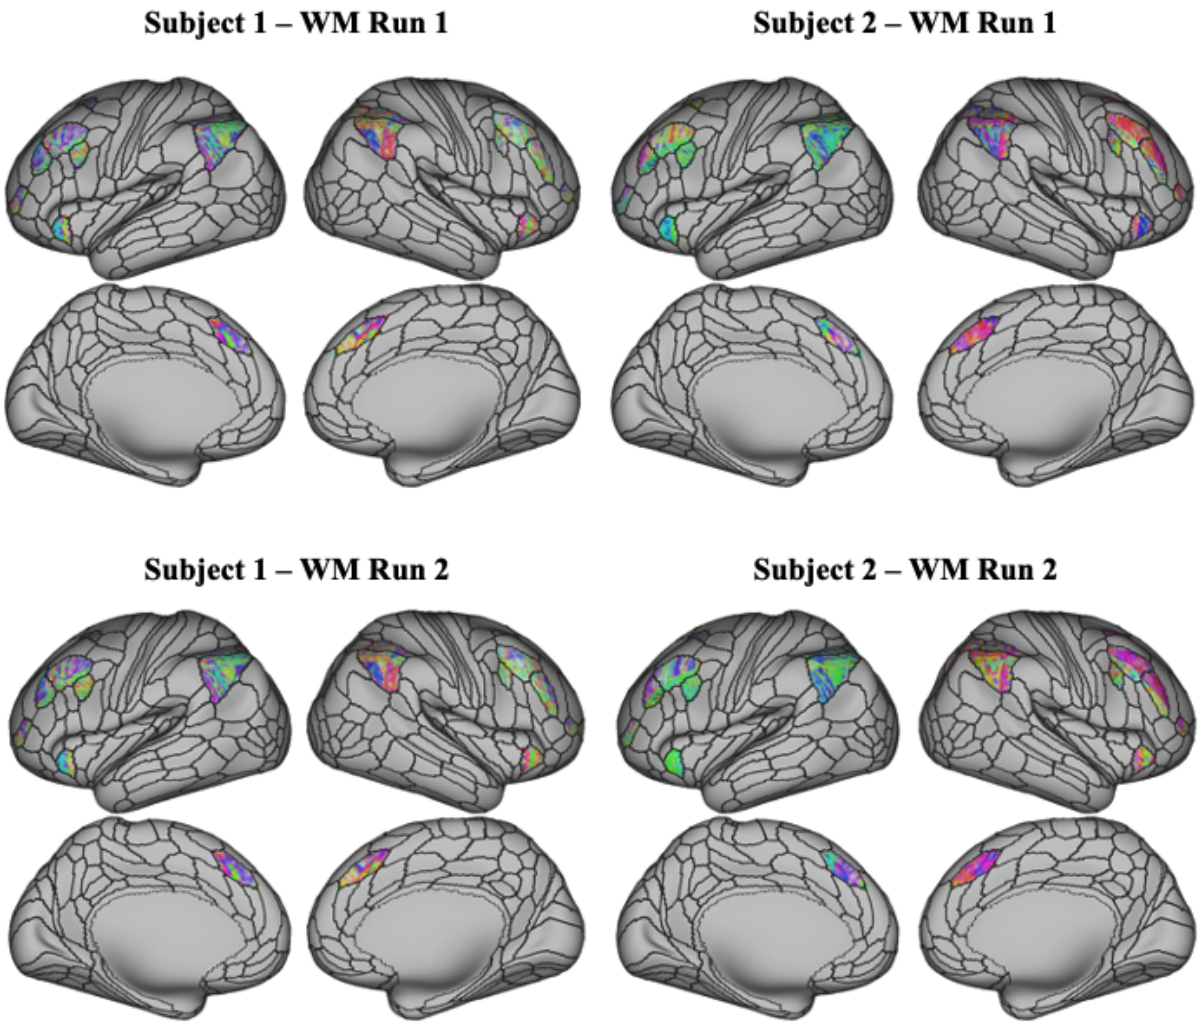

Figure 3.2: FC within the FPN visualized on the cortical surface. Vertices with similar colours exhibit similar time courses. Working memory task state data for run 1 and run 2 of two example subjects are shown with a lateral and medial view of each hemisphere. FC appears more consistent within than between individuals.

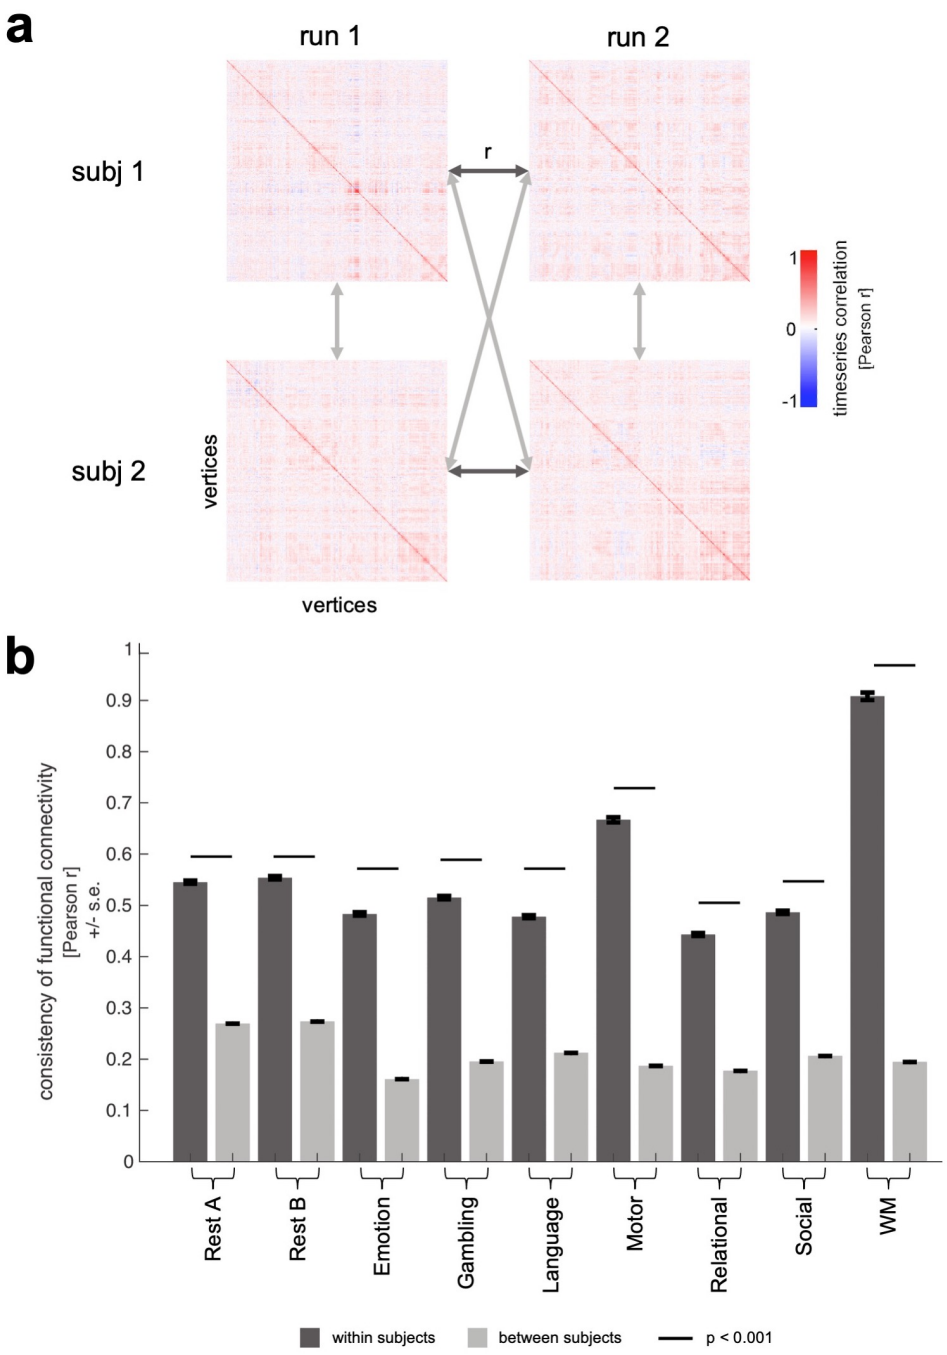

Figure 3.3: FC within the FPN is more consistent within than between individuals. Results are based on the 100 unrelated individuals in the HCP data set. Panel a schematically shows how we computed consistency of FC within and between individuals. Panel b shows results averaged across all within- and between-subject run pairs for each task and for each of the two resting-state acquisition days. The x-axis represents the resting state runs (rest A=run 1-2; rest B=run 3-4) and the seven different tasks. Error bars reflect standard error based on subject-label randomization. Statistical inference was performed using subject-label randomization (5,000 randomizations).

participant. Residual time courses were concatenated across the two runs for each task before computing the FC matrices that were used for the visualization. These visualizations suggest that FC within the FPN is moderately consistent across tasks. However, between-task consistency seems lower than within-task consistency (compare to Figure 3.2). These observations were similar for the other participants.

To quantify these observations, we estimated the consistency of FC matrices within and between tasks. This was estimated by computing pairwise Pearson correlation coefficients between FC matrices for runs of the same task and for runs of different tasks within individuals. Figure 3.5a shows the resulting correlation matrix for an example participant. The correlation matrix suggests that frontoparietal FC is more consistent within than between tasks. Figure 3.5b shows within- and between-task FC correlations, averaged across run pairs and individuals. Results are shown for all individuals, for related individuals, and for unrelated individuals. We performed inference on the difference in within- and between-task consistency using paired t-tests (one-sided test). The difference was significant for all three groups ($t(923) = 236.5$, $t(823) = 222.0$, and $t(99) = 82.7$; $p<0.001$ for each), confirming that frontoparietal FC is more consistent within than between tasks. This provides evidence for task-dependent internal reconfigurations of the FPN. These outcomes suggest that flexible reorganization may be a principle of frontoparietal function, and provide a starting point for further investigations into the network's internal functional organization.

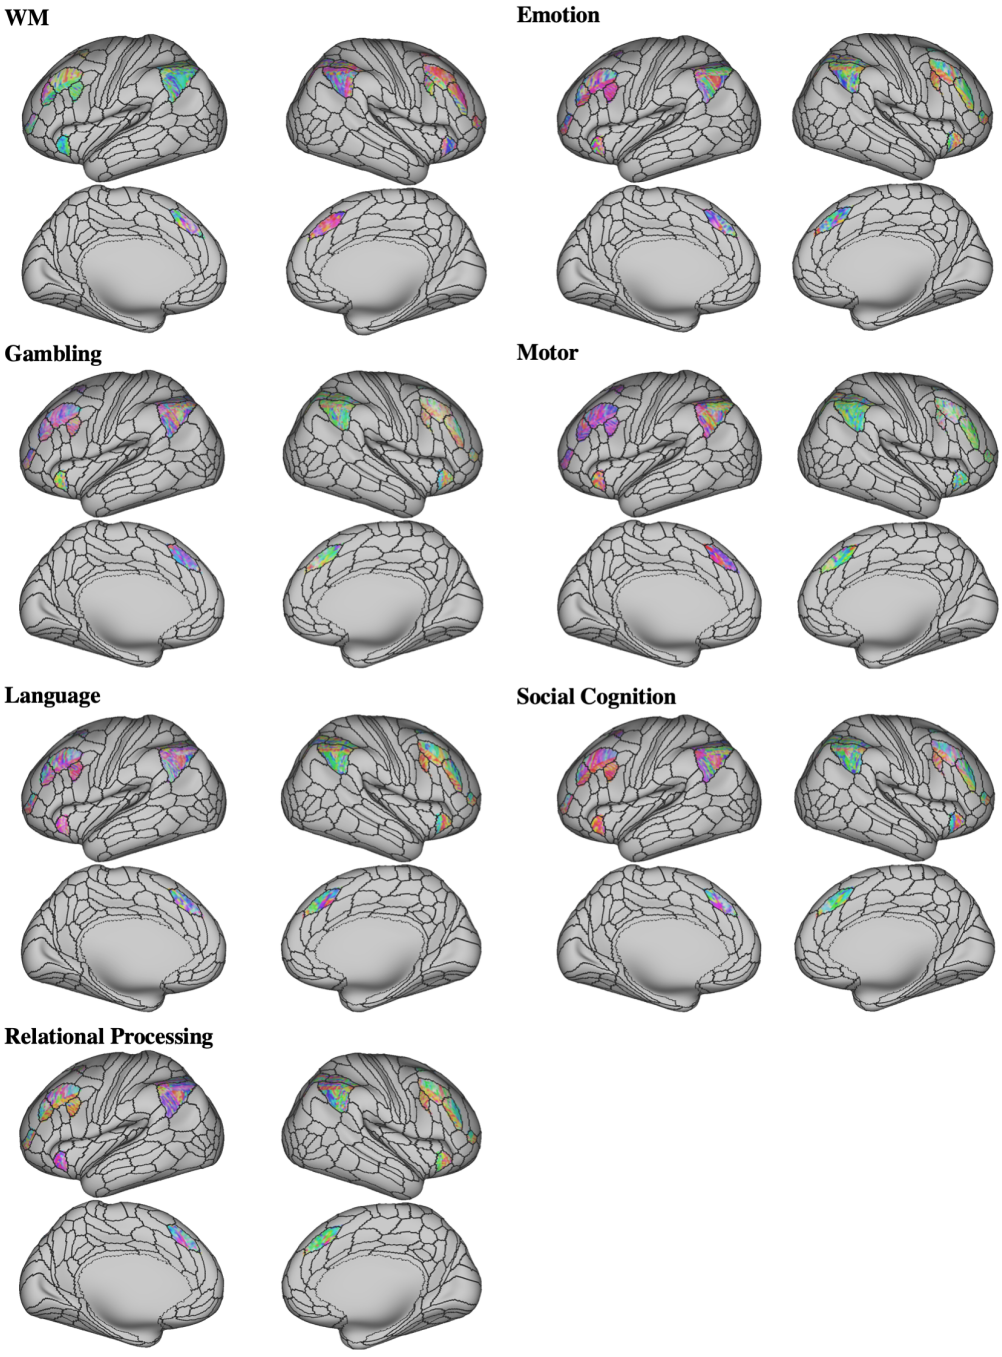

Figure 3.4: FC within the FPN visualized on the cortical surface. Vertices with similar colours exhibit similar time courses. Data from all seven tasks (time series were concatenated across runs for each task) for one example subject are shown with a lateral and medial view of each hemisphere. While the visualizations suggest some consistency in FC across tasks, there appear to be differences as well. These observations suggest that FC may be more replicable within (see Figure 3.2) than between tasks.

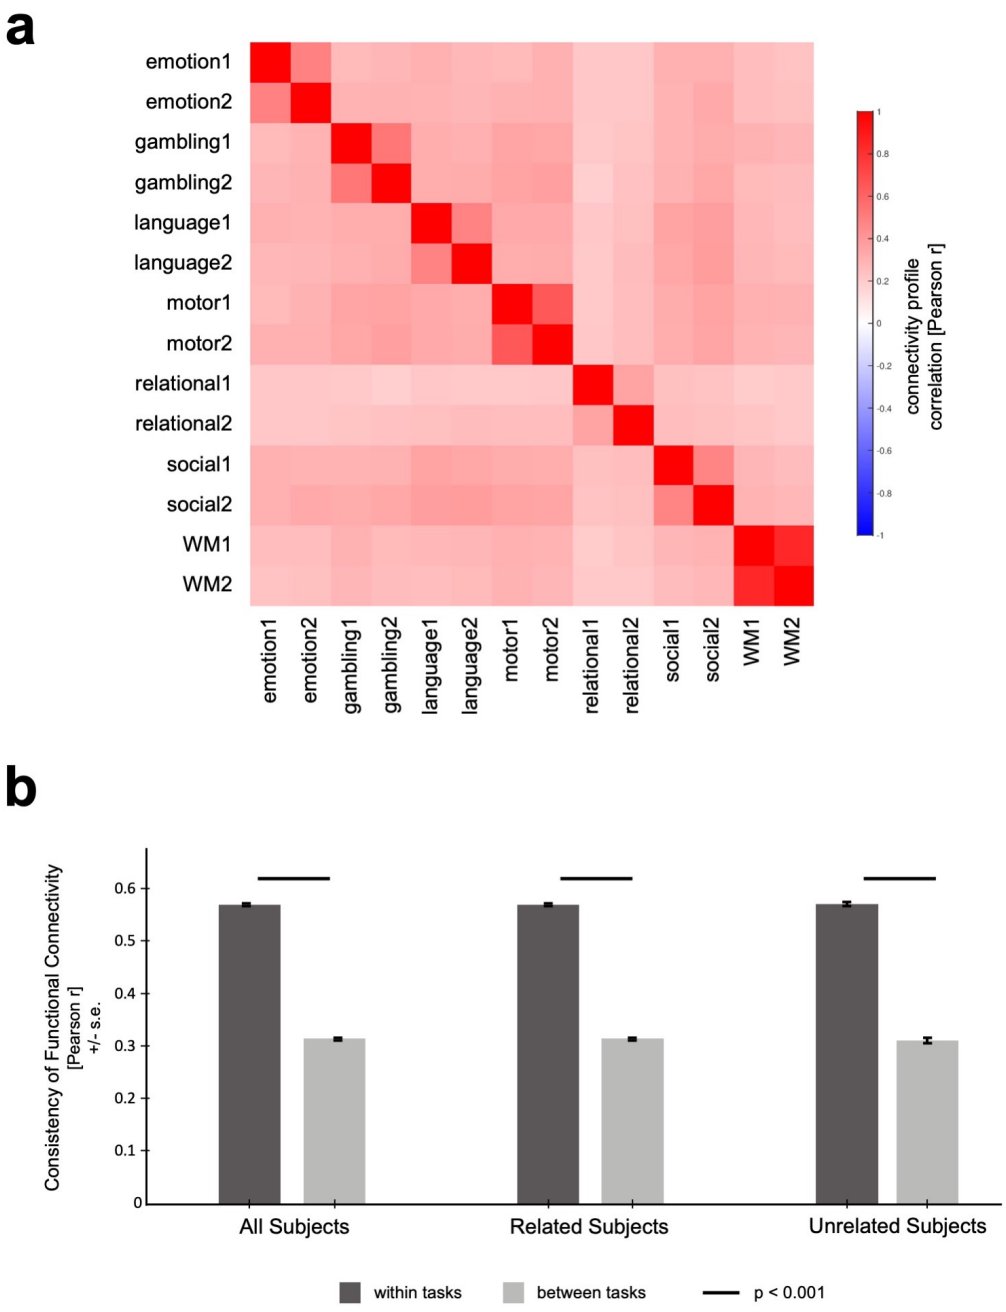

Figure 3.5: FC within the FPN is more consistent within than between tasks. Panel a: Consistency of FC within and between tasks shown for an example participant. The axes of the matrix represent the different tasks, with two runs per task. The diagonal values represent correlations of FC profiles within tasks and off-diagonal values represent correlations between different tasks. Panel b: Bars reflect consistency of FC within and between tasks, averaged across task pairs. X-axis represents groups of subjects: all subjects (n=924), related subjects (n=824), and unrelated subjects (n=100). Error bars reflect standard error of the mean across subjects. Statistical inference was performed with a paired t-test.

Chapter 4

Discussion

4.1 Summary

Throughout the current thesis, we have aimed to understand the functional organization of the human FPN on a fine-grained spatial scale as this has not been consolidated in previous literature. We did this through examining FC *within* the human FPN during the execution of cognitive tasks. We analyzed fMRI task-state data from 924 individuals from the HCP Young Adult study (Van Essen et al., 2012). Our results show that FC within the FPN is highly stable across time within individuals. Furthermore, FC within the FPN is more consistent within than between individuals. These findings provide evidence for the existence of a fine-grained functional organization within the FPN that is partially unique to the individual. Prior work on large-scale functional brain organization reported similar results (Gratton et al., 2018), suggesting that individual variability in FC is substantial at both local and global scales. Individual variability in global FC has been shown to predict variability in behaviour on higher-order cognitive tasks (Finn et al., 2015). Our findings raise the possibility that individual variability in local FC may also contribute to variability in cognitive performance across individuals.

Our results additionally show that FC within the FPN is more consistent within than

between tasks within an individual. This finding provides evidence for task-dependent reconfigurations of FC within the FPN. In other words, the FPN may adjust its input sensitivities at a local scale to best accommodate changing task demands. This is exciting because it provides additional evidence for a fine-grained functional organization within the FPN. Our findings extend prior work which showed that the FPN reconfigures its connectivity with other brain networks at a global scale to meet task demands (Cole et al., 2013). Flexible reorganization may be a general principle of frontoparietal function at both local and global scales.

This thesis also contributes tools for visualization of FC profiles on the cortical surface. We visualized FC maps of the FPN by utilizing MDS. Visualizations are important because there may be different spatial structures possible. For instance, on one side, there could have been distinct spatial clusters in the FPN where vertices with the same time courses are clustering together. However, we observed more moderate maps where there was no clearly observable spatial clustering. These different organizations are useful to visualize because they are indicative of different types of organizations. While MDS provides a useful visualization tool, the limitations that accompany MDS must be taken into consideration when interpreting results. Furthermore, the spatial structure observed in the maps needs to be quantified further in order to draw stronger conclusions, for example by performing clustering analyses before projecting results to the cortical surface.

Overall, this study shows clearly that human individuals exhibit a partially unique fine-grained functional organization within the FPN with some shared components, and that this organization contains a task-specific component.

4.2 Limitations and Future Directions

Although the Glasser parcellation (Glasser et al., 2016) provided a good framework for defining the FPN, there are limitations to its usage. For instance, since the parcellation is based on a group average, this poses a problem because in reality, individual subject brains may vary in

their parcels slightly. There may be differences in the degree of overlap with the parcellation and this may potentially be an underlying source of apparent changes in FC between subjects. This is a confound because the parcels are not defined precisely for each individual brain.

Recent studies have begun proposing improved methods for defining functional boundaries. For example, King et al. (2019) developed a method to identify functional boundaries of the cerebellum based on task activations. Haak et al. (2018) proposed a different approach: they computed connectivity gradients based on voxel-wise FC profiles. Applying these approaches would allow clustering of voxels based on how similar they are in their connectivity to the rest of the brain. This approach would be complementary to our local connectivity measurements which focused on connectivity within the FPN only. One intermediate approach could be a searchlight approach (Oosterhof et al., 2016). The conceptual motivation behind the searchlight approach is to find stable local structure in FC. It combines the stability of large-scale group parcellations with the sensitivity of local connectivity gradients to detect fine-grained functional organization. This makes the searchlight approach sensitive to the fine-grained functional organization of frontoparietal cortex. In order to develop this approach, the next steps would involve addressing how to visualize the results (e.g. colour the edges), and how to detect the functional boundaries (e.g. watershed approach, boundary search on local edges, compare global and local solutions to detect noisy vertices). In developing the searchlight approach, we should also generate the two extremes that the approach is trying to bridge: global parcellation and local connectivity gradients. Correspondence between global and local solutions would be encouraging and would suggest that the searchlight approach should give stable results.

Our results are suggestive of task-related reconfigurations of the FPN. An important next step is to further examine the network changes and investigate whether they are relevant for behaviour. For example, are some vertices more strongly reconfiguring their connectivity to accommodate task demands than others? Where are these vertices located? Which other parts of the brain are they connected to? We could potentially identify such vertices by computing a local variable connectivity measure for each vertex, reflecting the variability in its connections to

other frontoparietal vertices across task states (extending Cole et al. (2013)). If this vertex-wise variability is replicable over time within an individual, we can further examine the spatial organization of 'flexible' versus 'stable' vertices within the FPN. Furthermore, flexible vertices may show different connectivity profiles to other parts of the brain than stable vertices. The average local variable connectivity across frontoparietal vertices could serve as a summary measure of task-related frontoparietal reconfiguration. This summary measure could be used for predicting behavioural performance. Combining this 'local' connectivity measure with existing 'global' connectivity measures (Finn et al., 2015) may improve performance at predicting behaviour during task states.

The HCP has acquired a plethora of data in the young healthy adult dataset, therefore expanding our analyses to other aspects of the dataset would provide further insight into the FC of the FPN. For example, we analyzed the data from the 3T fMRI dataset, and it would be beneficial to next expand our analyses to the 7T dataset in order to assess the FC of the FPN on higher resolution data. Also, future studies should investigate the relationship between fMRI and structural MRI data. This stems from the idea that FC may be driven (partly or largely) by structural connectivity (Hermundstad et al., 2014; Deco et al., 2014). Since the HCP has acquired structural data for the subjects we analyzed, this would be a practical next step to implement. Another advantage of the HCP dataset is the large amount of siblings, including monozygotic and dizygotic twins. Therefore, differences between family members (twins and non-twin siblings) and unrelated subjects should be investigated. A limitation of the HCP dataset is that it only has seven tasks for the 3T dataset and two tasks for the 7T dataset. In future studies, it would be advantageous to extend our set of analyses to datasets with more tasks in order to engage more complex cognitive functions.

A general limitation of our work is that we are limited by the measurement technique. In the case of fMRI, this means that we have limited temporal resolution, which means that we are only sensitive to activity fluctuations in lower frequency ranges (0.01 - 1 Hz) (He et al., 2008). Low-frequency fluctuations are predominantly sensitive to network reconfigurations at a

global spatial scale and may fail to detect local reconfigurations (Matsui et al., 2016) (but see Kiani et al. (2015)). This may to some extent limit what we can infer from fMRI data about task-related reconfigurations within the FPN. Data acquired in nonhuman primates with cell array recordings (Kiani et al., 2015) may provide complementary evidence for task-related reconfigurations of FC at a local scale.

4.3 Concluding Remarks

The work presented in this thesis provides initial evidence for task-related reconfigurations of FC within the human FPN. The reported results serve as a stepping stone toward characterizing the functional organization of frontoparietal cortex and understanding its role in supporting higher-order cognitive function.

Bibliography

Agrawal, U., Brown, E. N., and Lewis, L. D. (2020). Model-based physiological noise removal in fast fMRI. *NeuroImage*, 205:116231.

Arslan, S., Ktena, S. I., Makropoulos, A., Robinson, E. C., Rueckert, D., and Parisot, S. (2018). Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage*, 170:5–30.

Assem, M., Glasser, M. F., Van Essen, D. C., and Duncan, J. (2020). A Domain-General Cognitive Core Defined in Multimodally Parcellated Human Cortex. *Cerebral Cortex*, 30(8):4361–4380.

Badre, D. and Nee, D. E. (2018). Frontal Cortex and the Hierarchical Control of Behavior. *Trends in Cognitive Sciences*, 22(2):170–188.

Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., Glasser, M. F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J. M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A. Z., and Van Essen, D. C. (2013). Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage*, 80:169–189.

Beckmann, C. and Smith, S. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 23(2):137–152. Conference Name: IEEE Transactions on Medical Imaging.

Bhandari, A., Gagne, C., and Badre, D. (2018). Just above Chance: Is It Harder to Decode Information from Prefrontal Cortex Hemodynamic Activity Patterns? *Journal of Cognitive Neuroscience*, 30(10):1473–1498.

Bijsterbosch, J. (2017). *Introduction to resting state fMRI functional connectivity*. Oxford neuroimaging primers. Oxford University Press, Oxford, United Kingdom ; New York, NY, first edition edition.

Biswal, B., Zerrin Yetkin, F., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, 34(4):537–541. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910340409.

Borg, I. and Groenen, P. J. F. (2005). *Modern Multidimensional Scaling: Theory and Applications*. Springer Science & Business Media. Google-Books-ID: duTODldZzRcC.

Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. (1996). Linear Systems Analysis of Functional Magnetic Resonance Imaging in Human V1. *The Journal of Neuroscience*, 16(13):4207–4221.

Buxton, R. B. (2012). Dynamic models of BOLD contrast. *NeuroImage*, 62(2):953–961.

Caballero-Gaudes, C. and Reynolds, R. C. (2017). Methods for cleaning the BOLD fMRI signal. *NeuroImage*, 154:128–149.

Cavada, C. and Goldman-Rakic, P. S. (1989). Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *The Journal of Comparative Neurology*, 287(4):422–445.

Chen, J. E., Lewis, L. D., Chang, C., Tian, Q., Fultz, N. E., Ohringer, N. A., Rosen, B. R., and Polimeni, J. R. (2020). Resting-state "physiological networks". *NeuroImage*, 213:116707.

Coalson, T. S., Essen, D. C. V., and Glasser, M. F. (2018). The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of*

Sciences, 115(27):E6356–E6365. Publisher: National Academy of Sciences Section: PNAS Plus.

Cole, M., Bassett, D., Power, J., Braver, T., and Petersen, S. (2014). Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron*, 83(1):238–251.

Cole, M. W., Ito, T., Bassett, D. S., and Schultz, D. H. (2016). Activity flow over resting-state networks shapes cognitive task activations. *Nature Neuroscience*, 19(12):1718–1726. Number: 12 Publisher: Nature Publishing Group.

Cole, M. W., Ito, T., Cocuzza, C., and Sanchez-Romero, R. (2021). The Functional Relevance of Task-State Functional Connectivity. *Journal of Neuroscience*, 41(12):2684–2702. Publisher: Society for Neuroscience Section: Research Articles.

Cole, M. W., Ito, T., Schultz, D., Mill, R., Chen, R., and Cocuzza, C. (2019). Task activations produce spurious but systematic inflation of task functional connectivity estimates. *NeuroImage*, 189:1–18.

Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., and Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, 16(9):1348–1355.

Cunningham, J. P. and Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 11 Primary_atype: Reviews Publisher: Nature Publishing Group Subject_term: Learning algorithms Subject_term_id: learning-algorithms.

Dagli, M. S., Ingeholm, J. E., and Haxby, J. V. (1999). Localization of Cardiac-Induced Signal Change in fMRI. *NeuroImage*, 9(4):407–415.

Deco, G., McIntosh, A. R., Shen, K., Hutchison, R. M., Menon, R. S., Everling, S., Hagmann, P., and Jirsa, V. K. (2014). Identification of Optimal Structural Connectivity Using Functional Connectivity and Neural Modeling. *Journal of Neuroscience*, 34(23):7910–7916.

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., and Fiez, J. A. (2000). Tracking the Hemodynamic Responses to Reward and Punishment in the Striatum. *Journal of Neurophysiology*, 84(6):3072–3077. Publisher: American Physiological Society.

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., and Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968–980.

Dias, R., Robbins, T. W., and Roberts, A. C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. *Nature*, 380(6569):69–72. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 6569 Primary_atype: Research Publisher: Nature Publishing Group.

Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience*, 2(11):820–829.

Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, 14(4):172–179.

Duncan, J. (2013). The Structure of Cognition: Attentional Episodes in Mind and Brain. *Neuron*, 80(1):35–50.

Duncan, J. and Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, 23(10):475–483.

Fedorenko, E., Duncan, J., and Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences*, 110(41):16616–16621.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., and Constable, R. T. (2015). Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18(11):1664–1671.

Fischl, B. (2004). Automatically Parcellating the Human Cerebral Cortex. *Cerebral Cortex*, 14(1):11–22.

Formisano, E. and Goebel, R. (2003). Tracking cognitive processes with functional MRI mental chronometry. *Current Opinion in Neurobiology*, 13(2):174–181.

Formisano, E., Kim, D.-S., Di Salle, F., van de Moortele, P.-F., Ugurbil, K., and Goebel, R. (2003). Mirror-Symmetric Tonotopic Maps in Human Primary Auditory Cortex. *Neuron*, 40(4):859–869.

Fox, M. D. and Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(9):700–711.

Friston, K. J., Frith, C. D., Liddle, P. F., and Frackowiak, R. S. J. (1993). Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets. *Journal of Cerebral Blood Flow & Metabolism*, 13(1):5–14. Publisher: SAGE Publications Ltd STM.

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., and Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178.

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., and Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124.

Goebel, R. (2007). Localization of Brain Activity using Functional Magnetic Resonance Imaging. In Stippich, C., editor, *Clinical Functional MRI*, pages 9–51. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Medical Radiology.

Goldman-Rakic, P. S. (1988). Topography of cognition: parallel distributed networks in primate association cortex. *Annu. Rev. Neurosci.*, 11:137–156.

Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., and Petersen, S. E. (2016). Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, 26(1):288–303.

Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., Nelson, S. M., Coalson, R. S., Snyder, A. Z., Schlaggar, B. L., Dosenbach, N. U., and Petersen, S. E. (2018). Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron*, 98(2):439–452.e5.

Greicius, M. D., Krasnow, B., Reiss, A. L., and Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100(1):253–258. Publisher: National Academy of Sciences Section: Biological Sciences.

Haak, K. V., Marquand, A. F., and Beckmann, C. F. (2018). Connectopic mapping with resting-state fMRI. *NeuroImage*, 170:83–94.

Hariri, A. R., Tessitore, A., Mattay, V. S., Fera, F., and Weinberger, D. R. (2002). The Amygdala Response to Emotional Stimuli: A Comparison of Faces and Scenes. *NeuroImage*, 17(1):317–323.

He, B. J., Snyder, A. Z., Zempel, J. M., Smyth, M. D., and Raichle, M. E. (2008). Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proceedings of the National Academy of Sciences*, 105(41):16039–16044. Publisher: National Academy of Sciences Section: Biological Sciences.

Hebart, M. N., Bankson, B. B., Harel, A., Baker, C. I., and Cichy, R. M. (2018). The representational dynamics of task and object processing in humans. *eLife*, 7:e32816. Publisher: eLife Sciences Publications, Ltd.

Henson, R. (2007). Efficient Experimental Design for fMRI. In *Statistical Parametric Mapping*, pages 193–210. Elsevier.

Hermundstad, A. M., Brown, K. S., Bassett, D. S., Aminoff, E. M., Frithsen, A., Johnson, A., Tipper, C. M., Miller, M. B., Grafton, S. T., and Carlson, J. M. (2014). Structurally-Constrained Relationships between Cognitive States in the Human Brain. *PLOS Computational Biology*, 10(5):e1003591. Publisher: Public Library of Science.

Kanwisher, N., McDermott, J., and Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17(11):4302–4311. Publisher: Society for Neuroscience Section: Articles.

Kaufman, M. T., Churchland, M. M., Ryu, S. I., and Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17(3):440–448. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 3 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Motor cortex;Premotor cortex Subject_term_id: motor-cortex;premotor-cortex.

Kiani, R., Cueva, C., Reppas, J., Peixoto, D., Ryu, S., and Newsome, W. (2015). Natural Grouping of Neural Responses Reveals Spatially Segregated Clusters in Prearcuate Cortex. *Neuron*, 85(6):1359–1373.

King, M., Hernandez-Castillo, C. R., Poldrack, R. A., Ivry, R. B., and Diedrichsen, J. (2019). Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nature Neuroscience*, 22(8):1371–1378.

Krienen, F. M., Yeo, B. T. T., and Buckner, R. L. (2014). Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 369(1653):20130526.

Kruskal, J. and Wish, M. (1978). *Multidimensional Scaling*. SAGE Publications, Inc., 2455 Teller Road, Thousand Oaks California 91320 United States of America.

Leavitt, M. L., Pieper, F., Sachs, A. J., and Martinez-Trujillo, J. C. (2018). A Quadrantic Bias in Prefrontal Representation of Visual-Mnemonic Space. *Cerebral Cortex (New York, N.Y.: 1991)*, 28(7):2405–2421.

Levy, I., Hasson, U., Avidan, G., Hendler, T., and Malach, R. (2001). Center-periphery organization of human object areas. *Nature Neuroscience*, 4(5):533–539.

Lewis, L. D., Setsompop, K., Rosen, B. R., and Polimeni, J. R. (2016). Fast fMRI can detect oscillatory neural activity in humans. *Proceedings of the National Academy of Sciences*, 113(43):E6679–E6685. Publisher: National Academy of Sciences Section: PNAS Plus.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. 412:8.

Lorenz, R., Violante, I. R., Monti, R. P., Montana, G., Hampshire, A., and Leech, R. (2018). Dissociating frontoparietal brain networks with neuroadaptive Bayesian optimization. *Nature Communications*, 9(1):1227. Bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Cognitive control;Cognitive neuroscience;Problem solving;Psychology Subject_term_id: cognitive-control;cognitive-neuroscience;problem-solving;psychology.

Luria, A. R. (1966). Disturbances of Higher Cortical Functions with Lesions of the Frontal Region. Pages: 218–295 Publication Title: Higher Cortical Functions in Man.

Machens, C. K., Romo, R., and Brody, C. D. (2010). Functional, But Not Anatomical, Separation of "What" and "When" in Prefrontal Cortex. *Journal of Neuroscience*, 30(1):350–360. Publisher: Society for Neuroscience Section: Articles.

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7474 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Cognitive neuroscience Subject_term_id: cognitive-neuroscience.

Matsui, T., Murakami, T., and Ohki, K. (2016). Transient neuronal coactivations embedded in globally propagating waves underlie resting-state functional connectivity. *Proceedings of the National Academy of Sciences*, 113(23):6556–6561.

Miller, E. K. and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24:167–202.

Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., and Uğurbil, K. (2010). Multiband Multislice GE-EPI at 7 Tesla, With 16-Fold Acceleration Using Partial Parallel Imaging With Application to High Spatial and Temporal Whole-Brain FMRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 63(5):1144. Publisher: NIH Public Access.

Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B., and Bandettini, P. A. (2009). The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? *NeuroImage*, 44(3):893–905.

Norman-Haignere, S. V., McCarthy, G., Chun, M. M., and Turk-Browne, N. B. (2012). Category-Selective Background Connectivity in Ventral Visual Cortex. *Cerebral Cortex*, 22(2):391–402.

Oosterhof, N. N., Connolly, A. C., and Haxby, J. V. (2016). CoSMoMVPA: Multi-Modal Multivariate Pattern Analysis of Neuroimaging Data in Matlab/GNU Octave. *Frontiers in Neuroinformatics*, 10. Publisher: Frontiers.

Oosterhof, N. N., Wiestler, T., Downing, P. E., and Diedrichsen, J. (2011). A comparison of volume-based and surface-based multi-voxel pattern analysis. *NeuroImage*, 56(2):593–600.

Parkes, L. M., Schwarzbach, J. V., Bouts, A. A., Deckers, R. h. R., Pullens, P., Kerskens, C. M., and Norris, D. G. (2005). Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magnetic Resonance in Medicine*, 54(6):1465–1472. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.20712.

Passingham, R. (1975). Delayed matching after selective prefrontal lesions in monkeys (Macaca mulatta). *Brain Research*, 92(1):89–102.

Petrides, M. and Pandya, D. N. (1999). Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *The European Journal of Neuroscience*, 11(3):1011–1036.

Poldrack, R. (2011). Inferring Mental States from Neuroimaging Data: From Reverse Inference to Large-Scale Decoding. *Neuron*, 72(5):692–697.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682. Publisher: National Academy of Sciences Section: Biological Sciences.

Rigotti, M. (2010). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Frontiers in Computational Neuroscience*, 4.

Rypma, B., Prabhakaran, V., Desmond, J. E., Glover, G. H., and Gabrieli, J. D. (1999). Load-Dependent Roles of Frontal Brain Regions in the Maintenance of Working Memory. *NeuroImage*, 9(2):216–226.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., and Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, 90:449–468.

Schultz, D. H. and Cole, M. W. (2016). Higher Intelligence Is Associated with Less Task-Related Brain Network Reconfiguration. *Journal of Neuroscience*, 36(33):8551–8561. Publisher: Society for Neuroscience Section: Articles.

Sejnowski, T. J., Churchland, P. S., and Movshon, J. A. (2014). Putting big data to good use in neuroscience. *Nature Neuroscience*, 17(11):1440–1441.

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., and Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science (New York, N.Y.)*, 268(5212):889–893.

Shepard, R. N. (1980). Multidimensional Scaling, Tree-Fitting, and Clustering. *Science*, 210(4468):390–398.

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., and Petersen, S. E. (1997). Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. *Journal of Cognitive Neuroscience*, 9(5):648–663.

Smith, A. M., Lewis, B. K., Ruttimann, U. E., Ye, F. Q., Sinnwell, T. M., Yang, Y., Duyn, J. H., and Frank, J. A. (1999). Investigation of Low Frequency Drift in fMRI Signal. *NeuroImage*, 9(5):526–533.

Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D. A., Griffanti, L., Harms, M. P., Kelly, M., Laumann, T., Miller, K. L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A. Z., Vu, A. T., Woolrich, M. W., Xu, J., Yacoub, E., Uğurbil, K., Van Essen, D. C., and Glasser, M. F. (2013). Resting-state fMRI in the Human Connectome Project. *NeuroImage*, 80:144–168.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., and Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045.

Sotiropoulos, S. N., Jbabdi, S., Xu, J., Andersson, J. L., Moeller, S., Auerbach, E. J., Glasser, M. F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D. A., Yacoub, E., Lenglet, C., Van Essen, D. C., Ugurbil, K., and Behrens, T. E. J. (2013). Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage*, 80:125–143.

Stiers, P. and Goulas, A. (2018). Functional connectivity of task context representations in

prefrontal nodes of the multiple demand network. *Brain Structure and Function*, 223(5):2455–2473.

Stuss, D. T., Levine, B., Alexander, M. P., Hong, J., Palumbo, C., Hamer, L., Murphy, K. J., and Izukawa, D. (2000). Wisconsin Card Sorting Test performance in patients with focal frontal and posterior brain damage: effects of lesion location and test structure on separable cognitive processes. *Neuropsychologia*, 38(4):388–402.

Tavor, I., Jones, O. P., Mars, R. B., Smith, S. M., Behrens, T. E., and Jbabdi, S. (2016). Task-free MRI predicts individual differences in brain activity during task performance. *Science*, 352(6282):216–220.

Thomas Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., and Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165.

Torgerson, W. S. (1958). *Theory and methods of scaling*. Theory and methods of scaling. Wiley, Oxford, England. Pages: xiii, 460.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain. *NeuroImage*, 15(1):273–289.

Uğurbil, K., Toth, L., and Kim, D. S. (2003). How accurate is magnetic resonance imaging of brain function? *Trends in Neurosciences*, 26(2):108–114.

Uğurbil, K., Xu, J., Auerbach, E. J., Moeller, S., Vu, A. T., Duarte-Carvajalino, J. M., Lenglet, C., Wu, X., Schmitter, S., Van de Moortele, P. F., Strupp, J., Sapiro, G., De Martino, F., Wang, D., Harel, N., Garwood, M., Chen, L., Feinberg, D. A., Smith, S. M., Miller, K. L., Sotiropoulos, S. N., Jbabdi, S., Andersson, J. L. R., Behrens, T. E. J., Glasser, M. F., Van Essen, D. C., and

Yacoub, E. (2013). Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *NeuroImage*, 80:80–104.

Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E. J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S. W., Della Penna, S., Feinberg, D., Glasser, M. F., Harel, N., Heath, A. C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S. E., Prior, F., Schlaggar, B. L., Smith, S. M., Snyder, A. Z., Xu, J., and Yacoub, E. (2012). The Human Connectome Project: A data acquisition perspective. *NeuroImage*, 62(4):2222–2231.

Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4):229–239. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 4 Primary_atype: Reviews Publisher: Nature Publishing Group.

Waskom, M. L. and Wagner, A. D. (2017). Distributed representation of context by intrinsic subnetworks in prefrontal cortex. *Proceedings of the National Academy of Sciences*, 114(8):2030–2035.

Williamson, R. C., Cowley, B. R., Litwin-Kumar, A., Doiron, B., Kohn, A., Smith, M. A., and Yu, B. M. (2016). Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models. *PLOS Computational Biology*, 12(12):e1005141.

Woolgar, A., Parr, A., Cusack, R., Thompson, R., Nimmo-Smith, I., Torralva, T., Roca, M., Antoun, N., Manes, F., and Duncan, J. (2010). Fluid intelligence loss linked to restricted regions of damage within frontal and parietal cortex. *Proc. Natl. Acad. Sci. U. S. A.*, 107(33):14899–14902.

Yan, L., Zhuo, Y., Ye, Y., Xie, S. X., An, J., Aguirre, G. K., and Wang, J. (2009). Physiological origin of low-frequency drift in blood oxygen level dependent (BOLD) functional magnetic resonance imaging (fMRI). *Magnetic Resonance in Medicine*, 61(4):819–827. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.21902.

Zhi, D., King, M., and Diedrichsen, J. (2021). Evaluating brain parcellations using the distance controlled boundary coefficient. preprint, Neuroscience.

Appendix A

Glasser Parcellation

Table A.1 lists the ten parcels from the Glasser parcellation that were selected to represent the FPN according to Assem et al. (2020).

Parcel Name	**Area Description**
8C	Area 8C
IFJp	Area IFJp
a9-46v	Area anterior 9-46v
p9-46v	Area posterior 9-46v
i6-8	Inferior 6-8 Transitional Area
AVI	Anterior Ventral Insular Area
8BM	Area 8BM
IP1	Area Intraparietal 1
IP2	Area Intraparietal 2
PFm	Area PFm Complex
8BM	Area 8BM

Table A.1: Parcel descriptions based on Glasser parcellation (Glasser et al., 2016)

Curriculum Vitae

Name: Geetika Gupta

Education

Master of Science, Neuroscience 2019 - 2021
University of Western Ontario
London, Ontario, Canada

Bachelor of Medical Sciences 2014 - 2019
University of Western Ontario
London, Ontario, Canada

Honours and Awards

2020-2021	NSERC Canada Graduate Scholarship (CGS-M)
2017	Dean's Undergraduate Research Opportunities Award (DUROP)
2017	USC Alumni Award
2016	Western Campaign Scholarship
2014	Western Scholarship of Excellence

Teaching Experience

2019-2020	Teaching Assistant, Western University

Publications

1. Rafeh, R., & **Gupta, G.** (2020). Information-Limiting Correlations in Neural Populations: The Devil Is in the Details. *Journal of Neuroscience, 40*(41), 7782-7784.

2. Lau, J. C., Parrent, A. G., Demarco, J., **Gupta, G.**, Kai, J., Stanley, O. W., ... & Peters, T. M. (2019). A framework for evaluating correspondence between brain images using anatomical fiducials. *Human Brain Mapping, 40*(14), 4163-4179.