EPFL

Thèse n° 8205

Exploring dynamic functional connectivity by incorporating prior knowledge of brain structure

Présentée le 18 décembre 2020

à la Faculté des sciences et techniques de l'ingénieur
Laboratoire de traitement d'images médicales
Programme doctoral en génie électrique

pour l'obtention du grade de Docteur ès Sciences

par

Anjali Bagunu TARUN

Acceptée sur proposition du jury

Prof. J.-Ph. Thiran, président du jury
Prof. D. N. A. Van De Ville, directeur de thèse
Prof. E. Tagliazucchi, rapporteur
Dr M. Bastiani, rapporteur
Prof. S. Micera, rapporteur

École polytechnique fédérale de Lausanne

2020

It is a paradoxical but profoundly true and important principle of life
that the most likely way to reach a goal is to be aiming not at
that goal itself but at some more ambitious goal beyond it.
— Arnold Toynbee

To my parents, my brothers and sisters, and my fiancé...

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Professor Dimitri Van De Ville, who has given me the opportunity to pursue a PhD in EPFL, one of the leading technical universities in the world. I am grateful for all the productive meetings we had, for all the talks about science, the continuous guidance all throughout the PhD, and for leading MIP:lab to become a positive research environment that nurtures.

Second, I am thankful to my thesis committee, Prof. Enzo Tagliazucchi, Dr. Matteo Bastiani, Prof. Silvestro Micera, and my jury president, Prof. Jean Philippe Thiran. Thank you for the time you spent reading and evaluating my work. I am grateful for all the insightful comments and the lively discussions, as well as the positive assessment of my work.

Next, I would like to thank all my collaborators, Prof. Sophie Schwartz, Prof. Nikolai Axmacher, Dr. Virginie Sterpenich, and Dr. Lampros Perogamvros, not only for providing us the data on simultaneous EEG-fMRI sleep recordings, but also for the interesting discussions, the productive research meetings, and the help in revising our manuscript. I would also like to thank my braingraph collaborators Dr. Hamid Behjat and Martin Larsson, for all their help in rereading our manuscripts and for the overall guidance in our work.

A huge thanks also to all my former and present colleagues in the MIP:lab – Younes, Giuli, Thomas, Van, Nawal, Giulia, Dani, Laura, Serafeim, Miljan, Nico, Lorena, Naghmeh, Isik, Valeria, Elvira, Silvia, Raph, Alessandra, Yasaman, Amir, Enrico, Ellie, and Karin. Life in MIP:lab has been very colorful, thanks to the happy environment that we all created together. I will always remember the lunch break games (Killer Game!) and the house/apartment parties we all enjoyed. Thank you for all the happy research discussions, the special moments we shared abroad during conferences, and for all other adventures we did in the last four years.

I am grateful to my big brother in academia, Thomas, who has helped me tremendously in rereading my manuscripts, in teaching me neuroscience concepts that I needed to know at the start of the PhD, and for giving me *academia pep talk* whenever I reach a low-point during the PhD. I'm also especially grateful to Younes, Giuli, and Van for the amazing moments we shared, and for the deep discussions on personal growth and life in general.

I would also like to thank my friends in Lausanne, Arielle, Charles, Sanjin, Miki, Dani, AK,

Acknowledgements

Martina, Soham, and Apa for all the happy moments that we shared, the weekly board games, the dinner celebrations, and the picnics along the lake. All of these activities have made my stay in Switzerland cheerful and memorable.

To my previous students, Marta, Sergio, Zhiwei, Ilaria, and Céline, for giving me the chance to mentor you. My experience supervising you has been tremendously helpful in my personal development as a mentor. Special thanks to Ilaria and Hermina, both of whom I've worked with closely in the past year. Thank you for the warm and friendly collaboration that extends beyond work.

Gusto ko ring pasalamatan ang aking mga kaibigan, magmula pa nung highschool. Ang mga pinakamatalik sa lahat – Kathleen at Elaine, sa mahabang pagkakaibigan, sa mga pang araw-araw na kwentuhan sa chat. Salamat sa pagbibigay kabuluhan sa salitang *kaibigan*. Maraming salamat sa aking pamilya sa Europa, Tor, Java, Aeriel, at Esme sa lahat ng ating pinagsamahan. Maraming salamat sa mga pagkakataong tayo ay sabay-sabay na nangarap. Maraming salamat sa pagiging inspirasyon sa bawat desisyon sa buhay. Kayo ang isa sa mga dahilan kung bakit naging mas makulay ang paglalakbay at pagsisimula ng panibagong buhay dito sa Europa.

Orly, Sevérine, and William, my immediate family in Switzerland – thank you for all the wonderful moments we shared. It feels especially good to have a family next to me ready to give a helping hand. Thank you for all the ski trips and the sumptuous dinners we all shared.

A special thanks also to the family of my fiancé, Mamina, Tatino, and Dominik, for welcoming me to your home. I spent a lot of my time with them in the last four years, and I feel very grateful to them for being open and for taking me as their own. Thank you to Mamina for being the sweetest mother-in-law one could ever ask. Dakujem velmi pekne.

Sa aking pamilya, ang pinakamatibay na haligi ng aking support group, mama, papa, Kuya Al, ate Cynthia, Millet, Kuya Ado, Balds, Wilbert, Emong, Orly, and ate Oxinne. Salamat po sa pagbibigay sa akin ng suporta. Salamat sa pagbibigay ng inspirasyon. Kayo ang dahilan kung bakit ko ipinagbubuti ang aking pag-aaral, at ang aking trabaho magmula nung ako ay bata pa lamang. Maraming salamat kay mama at papa, na parating nandiyan para makinig sa aking mga kwento at sa pagpapalaki sa amin mga magkakapatid na puno ng pangarap at malasakit sa isa't isa. Maraming salamat sa pagdadasal para sa aming lahat.

Last but not the least, I would like to give my sincerest gratitude to the love of my life, Igi. We started the PhD together, and it has been the most amazing 4 years. We complement each other so well in many different aspects of life, and I've grown so much as a person because of you. Thank you for being the most positive person I know and for loving me unconditionally. I love you, Igi.

Lausanne, December 1, 2020 A. T.

Abstract

The synchronized firing of distant neuronal populations gives rise to a wide array of functional brain networks that underlie human brain function. Given the enormous perception, learning, and cognition potential of the human brain, it is not surprising that network-level representations of brain function reveal intrinsically rich organizational structure. What remains puzzling is how the human brain maintains its vast repertoire of functional connectivity (FC) states despite being constrained by the underlying fixed anatomical substrate. Fortunately, advances in modern neuroimaging technologies, such as diffusion-weighted magnetic resonance imaging (DW-MRI), have made it possible to map the brain's anatomical scaffold, while functional MRI (fMRI) provides complementary information on neural activity. In this thesis, we develop and apply methods for combining DW-MRI and fMRI data into an integrated framework to analyze the interplay between brain structure and function.

We first explored the dynamics of brain function during wakefulness and across the different non-rapid eye movement (NREM) sleep stages. We applied the innovation-driven co-activation pattern (iCAP) analysis, an advanced approach that can characterize the spatial and temporal organization of overlapping large-scale brain networks. Our results reveal new spatial patterns covering regions that support the physiological organization of sleep and arousal. We observe that the persistence in time of individual brain maps across different sleep stages was significantly altered. Contrary to the previously observed decreasing FC that accompanies increasing sleep depth, we instead observe a surge of network activity and cross-network interactions during NREM stage 2, followed by an abrupt decrease in NREM stage 3.

We relate the observed changes in FC during sleep to structural connectivity (SC) by going from voxel-level iCAP analysis to region-wise parcellation of brain data. This downsampling of functional data to match the resolution of SC is a common practice in the field. Indeed, graph modeling of brain structure has always been limited to defining nodes as regions of a brain atlas. In this next step of the thesis, we then propose a new method that models all brain voxels as nodes of a high-resolution voxel-level brain graph. We provide two ways to construct the graph and we characterize their properties by performing a spectral analysis of the graph Laplacian operator. Our findings show that despite the huge dimensionality of

the proposed brain graphs (around 750,000 voxels), the majority of the structural information is captured by the Laplacian's lowest frequency eigenmodes. More interestingly, this lower end of the Laplacian spectra also captures about 85% of the energy content of functional MRI, suggesting that functional patterns are overall smooth over the structure, thus providing, for the first time, a direct and quantitative measure of how much brain function is shaped by the anatomy.

Going beyond a scalar measure of the SC-FC link, we introduce a new framework that interpolates gray matter signals onto the white matter using the structure embedded in the voxel-level brain grid to guide the process. This enables visualization of key white matter structures that link temporally coherent gray matter areas. By applying classical dynamic FC analysis on the interpolated volumes, we observe whole-brain structure-function networks that extend currently known spatial patterns that are limited within the gray matter only. This new approach reveals the collective mediation of white matter pathways across short and long-distance functional connections.

Finally, we assessed the observed collective structural mediation by giving a quantitative measure of the overall anatomical range between temporally coherent gray matter areas. As a first major demonstration of the powerful use of the voxel-level brain graph in the context of graph signal processing (GSP) for functional brain imaging, we utilized a canonical model of graph diffusion to extract the anatomical range of functional network interactions. We find that this measure meaningfully differentiates brain regions according to a behaviorally relevant macroscale gradient that divides the cortex between low-level primary sensory areas and high-level cognitive functions.

Keywords: Brain structure and function, functional MRI, diffusion MRI, white matter inpainting, interpolation, voxel-wise brain graph, sleep, diffusion, linear models, Laplacian, eigenmodes

Résumé

L'activation synchronisée de populations neuronales distantes donne lieu à un riche répertoire de réseaux fonctionnels qui sous-tendent la fonction cérébrale humaine. Au vu de l'énorme potentiel du cerveau en terme de perception, d'apprentissage et de cognition, il n'est pas surprenant qu'une telle description révèle une riche structure organisationnelle. Toutefois, la manière dont celui-ci maintient son vaste répertoire d'états de connectivité fonctionnelle (CF) malgré une anatomie fixe demeure mystérieuse. Heureusement, les avancées en techniques de neuroimagerie modernes, comme l'imagerie par résonance magnétique (IRM) par diffusion, ont rendu possible la cartographie de l'organisation anatomique du cerveau. Dans le même temps, l'IRM fonctionnelle apporte une information complémentaire sur l'activité neuronale. Dans cette thèse, nous développons et appliquons des méthodes combinant les données d'IRM fonctionnelle et par diffusion en un outil multimodal pour analyser la dépendance entre structure et fonction du cerveau.

Nous avons tout d'abord exploré la dynamique de la fonction cérébrale durant l'éveil et les différents stages du sommeil lent. Nous avons généré des cartes de co-activation par innovation, qui caractérisent l'organisation spatiale et temporelle de réseaux neuronaux superposés à grande échelle. Celles-ci comprennent de nouvelles configurations spatiales incluant des régions qui contribuent à l'organisation physiologique du sommeil et de l'éveil. Nous avons aussi observé que la persistence dans le temps de configurations spécifiques est significativement modifiée en fonction du stage de sommeil. Contrairement à une CF moindre préalablement observée lors d'un sommeil plus profond, nous avons mis en lumière un pic d'activation des réseaux et une augmentation de leurs interactions pendant le sommeil lent de stage 2, avant une diminution nette lors du stage 3.

Dans un deuxième temps, nous relions les changements de CF observés pendant le sommeil à la connectivité structurelle (CS) en utilisant une parcellation régionale des données du cerveau au lieu d'une description à l'échelle du voxel. Ce sous-échantillonnage des signaux fonctionnels pour mieux correspondre à la résolution de la CS est une pratique courante ; en effet, la modélisation par graphes de la structure du cerveau a toujours été limitée par la définition des noeuds représentant les différentes régions du cerveau. Nous proposons ensuite une nouvelle méthode traitant chaque voxel comme un noeud d'un graphe à haute résolution du cerveau. Nous introduisons deux façons de construire le graphe, et nous caractérisons leurs propriétés par une analyse spectrale de son Laplacien. Nos résultats montrent qu'en dépit de la très large

dimensionalité des graphes obtenus (environ 750,000 voxels), la majorité de l'information structurelle est capturée par les modes propres du Laplacien avec la fréquence la plus faible. De plus, cette partie basse du spectre du Laplacien comprend aussi environ 85% du contenu énergétique de l'IRM fonctionnelle, suggérant que les configurations fonctionnelles sont dans l'ensemble lisses par rapport à la structure sous-jacente. Ceci permet, pour la première fois, de mesurer directement et quantitativement à quel point la fonction cérébrale est façonnée par l'anatomie.

Pour aller au-delà d'une mesure scalaire du lien entre CS et CF, nous poursuivons par l'interpolation des signaux de la matière grise au sein de la matière blanche, en utilisant un graphe structurel du cerveau à l'échelle du voxel pour guider le processus. Ceci permet la visualisation de structures essentielles de la matière blanche qui relient des régions temporellement cohérentes de la matière grise. En appliquant une analyse classique de CF dynamique sur les volumes interpolés, nous observons des réseaux structuro-fonctionnels à grande échelle transcendant les configurations spatiales limitées à la matière grise connues jusqu'ici. Cette approche novatrice révèle la médiation collective des faisceaux de matière blanche sur les connections fonctionnelles de courte et de longue distance.

Finalement, nous avons considéré la médiation structurelle collective par la quantification du degré anatomique global entre les régions de la matière grise temporellement cohérentes. Comme premier jalon de l'usage d'un graphe à l'échelle du voxel dans le traitement de signal sur graphe pour l'IRM fonctionnelle, nous avons utilisé un modèle canonique de la diffusion sur graphe pour extraire l'étendue anatomique des interactions entre réseaux fonctionnels. Les régions du cerveau sont alors différenciées selon un gradient macroscopique comportementalement pertinent qui divise le cortex entre les aires sensorielles primaires de bas niveau et les fonctions cognitives de haut niveau.

Mots-clés : Structure et fonction cérébrale, IRM fonctionnelle, IRM de diffusion, interpolation de matière blanche, interpolation, graphe des connexions cérébrales (par voxel), sommeil, diffusion, modèles linéaires, Laplacien, modes propres

Contents

List of Figures

List of Tables

List of Abbreviations

Brain regions and networks

amPFC	anterior medial prefrontal cortex
CEB	cerebellum
DMN	Default Mode Network
dmPFC	dorsal medial prefrontal cortex
ECN	Executive Control Network
Fmajor	Forceps major
Fminor	Forceps minor
Fornix	Fornix
GM	gray matter
HF	Hippocampal formation
L/Rcing	Left / right cingulum
L/Rcing2	Left / right hippocampal cingulum
L/Rcst	Left / right cortico-spinal tract
L/Rifo	Left / right inferior fronto-occipital fasciculus
L/Rilf	Left / right inferior longitudinal fasciculus
L/Rslf	Left / right superior longitudinal fasciculus
L/Runc	Left / right uncinate fasciculus
LTC	lateral temporal cortex
mPFC	medial prefrontal cortex
OFC	Orbitofrontal Cortex
PCC	posterior cingulate cortex
PHC	Parahippocampal cortex
pIPL	Posterior inferior parietal lobule
Rsp	Retrosplenial cortex
TempP	temporal pole
TPJ	temporal parietal junction
vmPFC	ventral medial prefrontal cortex
WM	white matter

General acronyms

AgCC	agenesis of the corpus callosum
BIC	Bayesian information criterion
BOLD	blood oxygenation level-dependent
CAP	co-activation pattern
CDF	cumulative distribution function
CSF	cerebrospinal fluid
DARTEL	diffeomorphic anatomical registration
dFC	dynamic functional connectivity
DTI	diffusion tensor imaging
DW-MRI	diffusion-weighted magnetic resonance imaging
EEG	electroencephalography
FC	functional connectivity
FD	framewise displacement
FDR	false discovery rate
fMRI	functional magnetic resonance imaging
FWHM	full width half maximum
GFT	graph fourier transform
GLM	general linear model
GLMM	Graph Laplacian Mixture Model
GSP	graph signal processing
GSR	graph signal recovery
HARDI	high angular resolution diffusion imaging
HCP	human connectome project
HRF	hemodynamic response function
ICA	independent component analysis
iCAP	innovation-driven co-activation pattern
inCAP	interpolated co-activation pattern
MNI	Montreal Neurological Institute
MRI	magnetic resonance imaging
N1	non-rapid eye movement sleep stage 1
N2	non-rapid eye movement sleep stage 2
N3	non-rapid eye movement sleep stage 3
NREM	non-rapid eye movement
NMR	nuclear magnetic resonance
ODF	orientation distribution functions
PCA	principal component analysis
PPA	point process analysis
RCD	relative cumulated durations
REM	rapid eye movement
RMSE	root mean squared error
ROI	region of interest

rs-fMRI	resting-state functional magnetic resonance imaging
RSN	resting-state network
SC	structural connectivity
SPM	statistical parametric mapping
SVD	singular value decomposition
SWS	slow wave sleep
TA	total activation
TDC	typically developing controls
TV	total variation
TW-dFC	track-weighted dynamic functional connectivity
ZC	zero crossing

1 Introduction

1.1 Motivation

The brain is composed of a myriad of neurons that are grouped into subpopulations according to their functional specialization. Different regions in the brain are engaged depending on what we do, perceive, feel, and interact with our surroundings and other human beings. For example, visual inputs are encoded in the occipital region of the brain following an orderly arrangement known as retinotopy [Tootell et al., 1998], which emerges from the spatial specificity of neuronal connections in different parts of the visual system. Similar organizing principles also govern other sensory inputs: auditory cortex encode tonotopic maps which reflect the representation of temporal frequency in the cochlea [Romani et al., 1982], while the somatosensory cortex incorporates maps corresponding to the surface of the body [Baumgartner et al., 1991]. Although topographic mapping of sensory inputs has been well-studied in animals for centuries, studies that aim to differentiate human cortical regions according to their functional roles have only bloomed significantly with the introduction of functional magnetic resonance imaging (fMRI). By incorporating the timing of the tasks when analyzing fMRI signals through regression analyses, it is possible to map distinct regions that respond to specific stimuli [Friston et al., 1994, Friston et al., 1999]. Executive processing, for instance, typically elicit a response in the prefrontal cortex, while introspection is associated with activations in the posterior part of the brain [Alvarez and Emory, 2006]. Meanwhile, other higher-order cognitive functions such as mind-wandering and episodic memory evoke coordinated responses in different areas in the brain, forming a network or activity maps in general. Altogether, this coordination of distant and neighboring neurons is measured within the framework of **functional connectivity (FC)**.

Along with the understanding that coordinated inter-areal communication underlies varied functionalities, also comes the fundamental question of how these neuronal coordinations take place. Communication between brain areas relies on specific routes or high-ways that relay messages from each region of the cortex. White matter (WM) nerve fibers, which are made up of nerve cell axons that extend from the neuron cell bodies of the gray matter (GM),

serve to connect the cerebrum. Due to the controversial interpretation of the biophysical origins of blood oxygenation level-dependent (BOLD) contrast in WM, most research in fMRI is only limited to GM BOLD. As such, an inevitable discontinuity arises when relating GM and WM particularly through the landscape of fMRI analysis. A popular avenue for investigating WM connectivity is through scalar-valued anisotropy measures derived from reconstructed models of water diffusion at every brain voxel using diffusion-weighted MRI (DW-MRI). In contrast to neuronal activity that is indirectly measured by the BOLD contrast, DW-MRI captures the strength of physical local connections between brain voxels, *i.e.,* the likelihood that two brain voxels are physically connected by a nerve fiber. Although very useful in identifying structural anomalies in patients suffering from structural brain deficits, these anisotropy values are merely summarizing measures that reflect maximal diffusivity and WM morphology. As such, this discontinuity in GM and WM connectivity remains and the question of how coordinated brain activity is mediated by the anatomy remains unanswered.

Alternative to voxel-level measures of anisotropy values, tractography approximates WM nerve fibers directly [Catani et al., 2002]. With the understanding that high level of FC suggests the existence of direct anatomical connections between these brain areas to support the ongoing information transfer between them, two pioneering studies [Greicius et al., 2009, van den Heuvel et al., 2009] have directly extracted the anatomical connection between FC by initializing fiber streamlines at predefined seed regions obtained via prior computation of FC. This approach allowed them to visualize specific WM tracks that connect a few GM regions which display temporally coherent activity. The main disadvantage of this method, however, is the direct selection of regions of interest (ROIs) which makes the analysis observer-dependent. Moreover, through directly correlating regional measures of FC and averaged fiber streamlines through **structural connectivity (SC)**, it has been shown that FC commonly exists between regions with no direct anatomical connectivity, suggesting that brain function is constrained and shaped by the anatomy, but not entirely [Honey et al., 2009, Andrews-Hanna et al., 2007, Hagmann et al., 2008]. Thus, taking instead the anatomy as a constraint, other studies have attempted to reproduce brain activity from predefined SC through numerical simulations [Galán, 2008, Deco et al., 2011, Deco et al., 2012].

Approximating FC from SC, or *vice versa* are hypothesis-driven and entail many explicit assumptions. To understand how distributed patterns of FC arise from a fixed underlying anatomy, a need for research methodologies that are observer-independent and purely data-driven is of utmost importance. Moreover, we need to find a solution that jointly integrates fMRI and DW-MRI into a single integrated framework that puts contributions from GM and WM at equal footing. Although fiber tracking can be done using each voxel in the GM as a seed, the complexity of the approach makes it difficult to accurately map the fibers at such a high spatial resolution. Thus, a more practical solution is to average all approximated fiber bundles coming from specific regions that are defined by a brain atlas. In doing so, WM contribution is only incorporated through the strength of SC connecting the GM regions, further amplifying the discontinuity between observations in GM and WM.

Over the last decade, region-wise parcellations of the brain which summarize measures that characterize brain structure and function have been the foundation of network neuroscience [Sporns, 2018] which has propelled our understanding of the brain. With the increasing interest to understand brain function with a dynamical point of view [Preti et al., 2017] and along with which lies the fundamental question of how the underlying brain anatomy supports FC, it is only timely that we advance network neuroscience to a more unifying analysis approach that takes into account the interplay between brain structure and function. Borrowing concepts from the emerging field of graph signal processing (GSP), functional data are defined atop the nodes of a structural graph and treated as time-dependent graph signals, on which connectome-informed signal processing operations can be performed [Huang et al., 2018a]. Recent findings on GSP have started to elucidate the role of structure-function relationship in supporting cognition and behavior [Medaglia et al., 2018, Preti and Van De Ville, 2019].

Nevertheless, despite these major advances, the understanding of how the overall brain organization works in its entirety, considering the interplay between SC and FC remains only partly answered. In particular, the understanding of GM neuronal coordination through WM structural mediation remains elusive. *To what extent is brain function shaped by the structure? Which WM tracts are recruited whenever functionally coupled GM areas communicate?* In this thesis, we attempt to answer these questions by developing a new set of methodologies that jointly integrate information from three imaging modalities based on MRI, namely the structural MRI, DW-MRI, and fMRI. We advance the GSP concept by extending region-wise analyses to a voxel-level perspective. In this approach, GM and WM are both defined as nodes and local structural connectivity is encoded onto voxel-to-voxel edges. This allows defining all fMRI BOLD signals from each voxel atop the nodes of a high-resolution brain grid (around 750,000 voxels), thus removing the discrepancy in GM and WM contributions that is inherent when defining nodes only according to GM regions. Using methods from GSP that are generalized from classical signal processing, we interpolate functional signals, measured on the GM onto the WM, using the voxel-level structural brain grid to guide the process. This allows observing how particular patterns of brain activity jointly rely on an ensemble of WM pathways. As the interpolated brain signals comprise the whole-brain at the voxel level, classical dynamic FC approaches can thus be applied to the data, just as how it is applied to typical fMRI data, enabling the visualization of key WM structures that mediate the interactions of temporally coherent GM areas.

In the following section, we provide a more detailed overview of the content of the thesis. A visual summary of the thesis organization and contributions is also supplied in Figure 1.1.

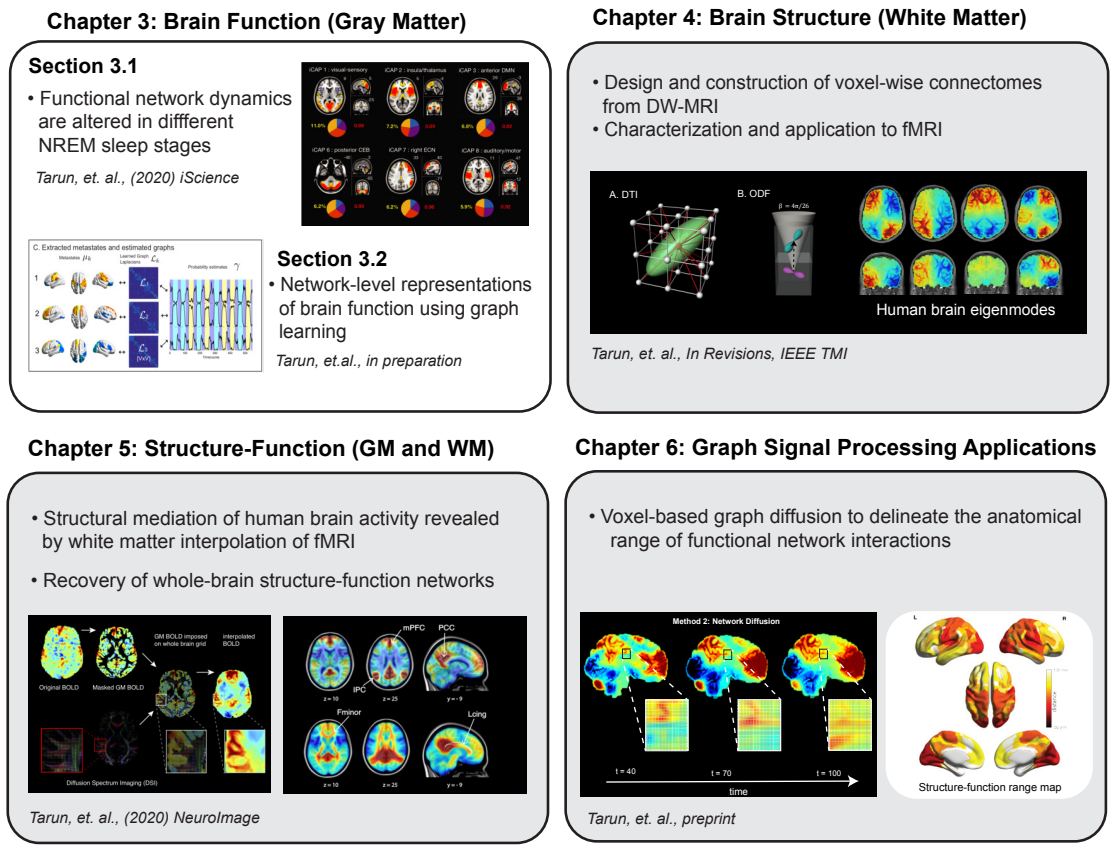

Figure 1.1 – **Thesis Overview and Contributions.** In **Chapter 3**, we applied two recently developed methodologies – one through (Section 3.1) innovation-driven co-activation patterns (iCAPs) and another (Section 3.2) through graph learning – to estimate network-level representations of brain function during tasks, rest, and across different NREM sleep stages. We tackle the similarity of these networks with region-wise structural connectomes. **Chapter 4** extends region-wise analyses to a higher resolution, voxel-level representation of brain structure. We discuss the design, construction, and spectral characterization of voxel-level brain graphs that encode structural information from DW-MRI. We also provide, as the first demonstration of its potential applications to functional MRI, a direct measure of how much brain function is shaped by the underlying anatomy. This brain graph will be used in **Chapter 5** to guide the interpolation of human brain activity from GM (GM) into the WM (WM) to extract full brain structure-function networks. These give the overall picture of how temporally coherent GM areas are mediated by specific combinations of WM bundles. In **Chapter 6**, we apply these newly developed methods in the context of graph signal processing applied to functional brain imaging. In particular, we propose a new approach that leverages voxel-based graph diffusion to extract the anatomical range of functional network interactions. The plate in white denotes contributions where we applied advanced methodologies to functional brain data, while those in gray correspond to contributions where new methodologies aiming to integrate diffusion and functional MRI data were designed and developed during the PhD.

1.2 Thesis Organization and Contributions

This thesis is organized as a compilation of 5 research articles, some of which are published while others are currently under review/revisions. Chapter 2 presents the background necessary for each of the studies included in this work. It introduces the reader to the basics of neuroimaging and the current tools and techniques for analyzing MRI data. We present existing dynamic FC methods for characterizing human brain function, followed by a discussion on current methods for probing brain structure. We then introduce different techniques in linking SC and FC, and give emphasis on GSP which is the basis of all analyses performed in the thesis. Chapters 3 to 6 contain reproduced research articles. Finally, Chapter 7 summarizes the results and provides a discussion for future research directions.

In the following section, we give a summary of the content of each chapter and present the goal and the research questions that we aim to answer in each chapter of the thesis.

Chapter 3: Large-scale functional brain network dynamics in tasks, rest, and sleep

This chapter presents two research articles [Tarun et al., 2020d, Tarun et al., 2020c] that mainly tackle the extraction and characterization of network-level representations of brain function using two advanced methods, one based on innovation-driven co-activation patterns (iCAPs) [Karahanoglu and Van De Ville, 2015], and another based on graph learning [Petric Maretic and Frossard, 2020]. We explore the similarity of the observed networks with the underlying structure using a direct statistical metric, similar to the classical comparison of SC and FC [Honey et al., 2009, Liégeois et al., 2016]. This comparison is done at the region-wise resolution by defining each region following a brain atlas.

Journal Article: NREM sleep stages specifically alter dynamical integration of large-scale brain networks

How do the dynamics of brain function change across the different sleep stages?

The main hypothesis underlying most FC studies is that large-scale brain networks reflect ongoing cognitive processes. It was therefore hypothesized that if conscious awareness dissipates as the brain transitions from wakefulness to deep sleep, then the activity of networks associated with high-level cognitive processes should also decrease. Surprisingly, it has been shown that these networks persist even in the deepest non-rapid eye movement (NREM) sleep stage [Larson-Prior et al., 2009]. Nevertheless, significant modifications have also been observed in terms of their spatial patterns, such as decreases in connection strengths [Larson-Prior et al., 2009], changes in the hierarchical organization into smaller modules [Boly et al., 2008], and reduced long-range temporal dependence in the BOLD signal [Tagliazucchi et al., 2013].

In this section, we highlight the changes in the temporal dynamics of large-scale functional net-

works as the brain transitions across NREM sleep stages, as well as their interactions/couplings. Using a novel method that relies on transient brain activity applied to two complementary studies of simultaneous EEG-fMRI recordings of sleep, we recover a wide array of spatially and temporally overlapping large-scale brain networks, thus revealing the organization principles of human brain function from wakefulness to deep sleep. Our results reveal new spatial patterns covering regions that support the physiological organization of sleep and arousal. Unexpectedly, contrary to previously observed strong functional dissociations that accompany the deepening sleep, we instead observe a surge of network activity and cross-network interactions during NREM stage 2, followed by an abrupt decrease of activity in NREM stage 3 [Tarun et al., 2020d]. To relate these findings with the underlying anatomy, we also provide a dynamic and region-wise comparison of FC and SC across the different NREM sleep stages. We observe that as FC breaks down during the deep sleep, we find an increasing alignment between correlation measures of brain structure and function.

Journal Article: Dynamics of Functional Network Organization Revealed by Graph Laplacian Mixture Models

Data-driven estimation of large-scale functional brain networks using graph learning.

Different areas in the brain are activated in a temporally coherent manner according to ongoing cognition. Classical methodologies in extracting brain networks typically employ regression techniques that incorporate prior information from the timing of the task paradigms. Due to the increasing interest in understanding the intrinsic functional organization of the brain in the absence of any stimulus, several methodologies have been proposed to allow the detection of prevalent spatial patterns in a purely data-driven manner.

In this section, we introduce a new method based on graph learning to extract meaningful repeating network patterns from averaged regional time-courses. We used a generative model that treats functional data as a collection of signals that live on multiple graphs. Using this approach, we are able to extract functional networks whose activity profiles accurately capture the timings of the task paradigms. We also present the corresponding connectivity matrices (*i.e.*, learned graphs) that describe the direct interactions between regions which comprise these metastates. We then estimate the degree of similarity between the estimated graphs and the underlying structural connectome at different time-points of task and rest.

Chapter 4: High-resolution voxel-level brain graphs

This chapter presents the first method developed in the thesis which characterizes brain structure at the voxel-level resolution [Tarun et al., 2019, Tarun et al., 2020a]. Both FC and SC graph models are typically defined based on region-wise parcellations of brain data. Although a number of dynamic FC approaches are available at the voxel-level resolution, graph models of SC are currently limited to region-wise definitions, for which WM contribution is only

represented through the strength of physical connection encoded in the edge weights between GM regions. In order to address the research questions raised in this thesis, we propose a new framework that addresses the discrepancy between GM and WM contributions by removing the limitation of doing all analyses in region-wise resolution and explicitly defining all brain voxels as nodes of a high-resolution brain graph. In this chapter of the thesis, we provide two methods to build a large brain graph onto which voxel-to-voxel connectivity from DW-MRI is encoded. We then provide, for the first time, a direct measure and a quantitative description of how much brain function is shaped by the underlying structure.

Journal Article: Voxel-Wise Brain Graphs from Diffusion-Weighted MRI: Spectral Analysis and Application to Functional MRI

How can we extend region-wise graph models of brain structure to the voxel-level perspective? To what extent is function shaped by brain structure?

In the previous chapter, we introduce two different approaches for extracting network-level representations of brain function. In each work, we also provide a comparison between brain structure and function, but only through direct correlation of SC and FC measures derived using region-wise parcellations of the brain. Indeed, graph modeling of structural connectivity has been very useful but has only been limited to defining nodes as regions from a brain atlas. Recent developments in computational neuroimaging used GSP operations to link brain structure and function. Region-wise connectomes are first constructed from DW-MRI. Then, functional data are defined atop the nodes of the graph and are interpreted as time-dependent graph signals. Signal processing operations can then be performed using methods generalized from 1D temporal signals to irregular graphs [Huang et al., 2018a]. In this design of the graph, however, nodes are only defined within GM, and WM is only incorporated in terms of the strength of the structural connections between GM.

In this chapter, we model WM explicitly as nodes of a high-resolution voxel-level brain grid (*i.e.*, 750,000 voxels). We propose two different methods for constructing voxel-level brain graphs that encode structural information from DW-MRI. The properties of the brain graph are characterized using spectral analysis of the graph Laplacian operator. The intrinsic dimensionality of the eigenspace is explored through a Procrustes validation scheme that characterizes inherent inter-subject variability. By decomposing task and resting-state fMRI data in terms of low-frequency eigenmodes, we find that about 85% of the energy content of fMRI signals are concentrated in the lowest 0.125% of the structural eigenspectrum. This provides, for the first time, a direct measure and a quantitative description of how brain function is shaped by the underlying anatomy. We conclude the chapter by providing research avenues for directly integrating brain structure and function at the voxel-level resolution.

Chapter 5: Recovery of whole-brain structure-function networks

This chapter presents a novel framework that takes advantage of the proposed voxel-level brain grid to address the main research problem of the thesis. We developed a novel approach that enables the observation of key white matter structures that mediate the interactions of temporally coherent GM areas [Tarun et al., 2020b]. This approach is purely data-driven and combines measures from fMRI and DW-MRI into a single integrated framework.

Journal Article: Structural mediation of human brain activity revealed by WM interpolation of fMRI

How is human brain activity mediated by structure?

So far, most of the works that linked brain function and structure can be categorized into four different types: (1) employing a direct statistical inter-dependence between two separately defined metrics for SC and FC [Honey et al., 2009] (*e.g.*, analyses we conducted in Chapter 3), (2) initializing FC and approximating fiber streamlines that connect predefined ROIs that are implicated [Greicius et al., 2009, van den Heuvel et al., 2009, Chamberland et al., 2015], (3) approximating brain activity from SC using numerical simulations [Galán, 2008, Honey et al., 2007, Deco et al., 2011, Deco et al., 2012], or (4) employing advanced graph theoretical approaches to examine how FC arises from, is controlled by, and is constrained by measures of SC [Gu et al., 2015, Becker et al., 2018, Goni et al., 2014]. Extracting SC from FC, or *vice versa*, are hypothesis-driven and require explicit assumptions. Moreover, none of these methods jointly and simultaneously integrates functional and diffusion MRI data into a *single* integrated framework.

In this chapter, leveraging on the proposed brain grid introduced in Chapter 4, we present a new method for interpolating functional signals measured on the GM, into the WM, using the structure encoded in the voxel-wise connectome to guide the process. This gives a visualization of how temporally coherent GM regions are mediated by combinations of different WM tracks. The output is a set of 4D volumes, similar to original functional volumes, but also includes interpolated signals in the WM. Using classical static and dynamic FC analyses, we present, for the first time, full brain structure-function networks, which extend currently known spatial patterns that are limited within the GM only. The functional relevance of these structure-function patterns is then further validated through investigations of their relation to task fMRI. Our findings reveal the structural mechanisms underlying the dynamic switching between task-positive and -negative subsystems of the default mode network (DMN) in different phases of working memory and relational task paradigms

Chapter 6: Graph signal processing applications on functional brain imaging: a voxel level perspective

This chapter demonstrates the powerful use of voxel-level brain graphs in the context of GSP applications for functional brain imaging. We present a novel framework to quantify the anatomical range of functional network interactions [Tarun and Van De Ville, 2020].

Journal Article: Anatomical range of functional interactions meaningfully differentiates brain regions

How do we obtain a measure of the anatomical range of functional network interactions?

Functionally coupled brain regions recruit different combinations of white matter structures to accommodate efficient inter-areal communication. Distance between functionally connected areas is usually defined using either euclidean or geodesic distance between two ROIs [Achard et al., 2006, Sepulcre et al., 2010, Oligschläger et al., 2017]. The brain topology, as observed in region-wise graphs, has shown small-world properties characterized by short path lengths and high clustering [Bassett and Bullmore, 2017]. This maintains the efficiency of inter-areal communication in the brain and minimizes the use of more costly long-distance functional connections.

In this chapter, we consider a canonical example of signal diffusion as a tool to extract the anatomical range of functional network interactions. We used a voxel-based graph diffusion framework that leverages methods borrowed from spectral graph theory [Kondor and Lafferty, 2002] as a canonical example to model the time it takes for a signal from the GM to be diffused to the WM. Starting from masked GM signals, taken in this application as a source of heat, we diffuse these signals into the WM using the structure encoded in a voxel-wise brain grid to guide the process. The length of time it takes for the signals to be completely diffused, equivalent to the obtained interpolated volumes proposed in Chapter 5, gives the optimal diffusion time necessary for WM bundles to connect temporally coherent GM areas. We found that this characteristic time, which is then converted to distance (in millimeters), is different for each region of the brain, and when projected to a cortical surface, follows an overarching organization that differentiates low-level primary sensorimotor areas from the more complex high-order association and limbic areas.

2 Background

Magnetic resonance imaging (MRI) has made it possible to map the brain's anatomical scaffold from diffusion-weighted MRI (DW-MRI), and in parallel, extract proxy of neural activity using functional MRI (fMRI). The main focus of this thesis is understanding the link between brain structure and function using simultaneous analysis of functional, structural, and diffusion MRI data. As such, we begin the chapter by providing an overview of the physical principles that underlie the acquisition of MRI. We delve into the use of functional MRI to capture blood-oxygnenation levels as a proxy of neural activity, structural MRI to image brain morphology, and diffusion MRI to capture white matter connectivity. We demonstrate how these modalities are exploited in different aspects of brain research. Specifically, we discuss the current methods and findings in the field of dynamic functional connectivity, followed by a discussion on region-wise structural connectomes. Finally, we enumerate recent developments in integrating complementary information from structural, diffusion and functional MRI to probe how brain function is shaped by the underlying structure.

2.1 Magnetic Resonance Imaging (MRI)

MRI has remarkably shaped the field of clinical and computational neuroscience over the past decades. It offers extensive possibilities to probe the brain's organization in different contexts. Although it is a relatively young field, the number of publications utilizing this modality in brain research is rapidly growing, thanks to its numerous advantages, such as its non-invasive feature and high spatial resolution.

Nuclear Magnetic Resonance (NMR)

Long before MRI was invented, there was nuclear magnetic resonace (NMR) which is applied to study the variation in molecular electron distribution of different chemicals. This technique is founded upon the idea that charged nuclei of atoms (*e.g.*, hydrogen atom) under the influence of a static magnetic field resonate when induced with a weak and secondary oscillating

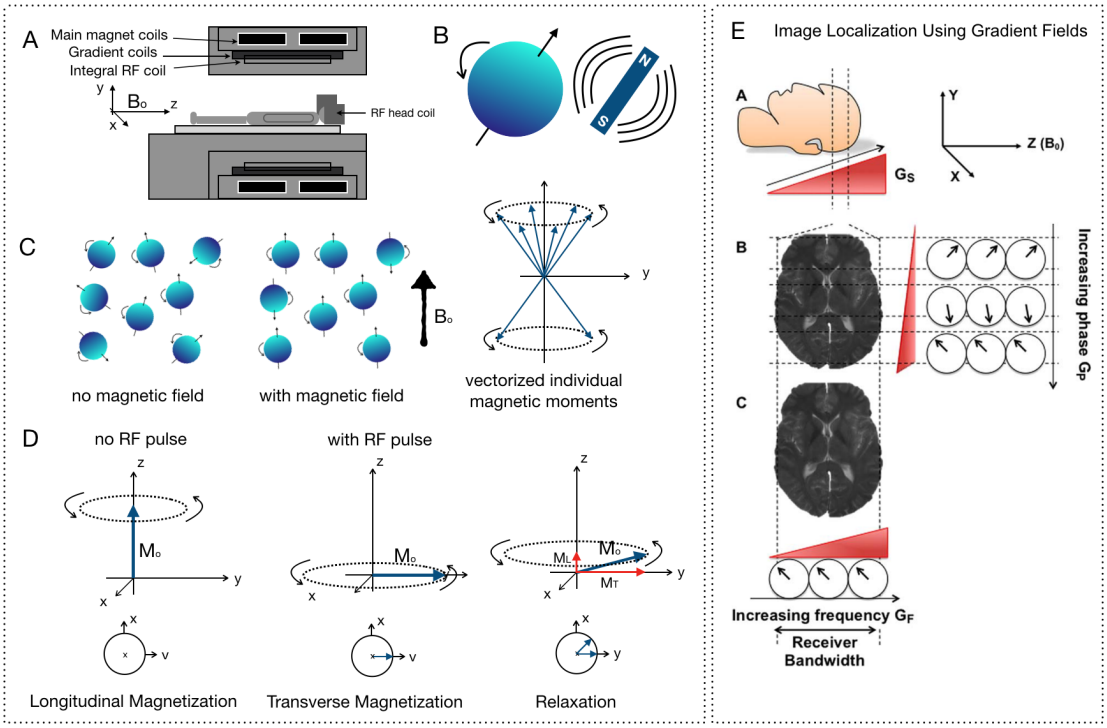

Figure 2.1 – **Basic principles of MRI.** (A) Major parts of an MRI scanner. (B) Hydrogen proton spins and acts like a bar of magnet. (C) Without magnetic field, protons spin in random orientations; with magnetic field B_0, they spin and precess around the direction of B_0. All individual protons can be visualized as vectors – the sum of their magnetic moments is equal to the (D) Longitudinal Magnetization in the B_0-direction. Under the influence of an RF pulse, this net magnetization becomes transverse to B_0. When the RF pulse is removed, the protons go back to their alignment with B_0 and relaxation occurs. Finally, (E) the NMR signal is localized using gradient fields in $x-$, $y-$, and $z-$ directions. Figure in (A) and (E) were adapted and reproduced with permission from [Currie et al., 2013].

magnetic field [Block et al., 1946]. Many years later in the 1970s, it was observed that there is a stark difference in magnetic resonance of cancerous and non-cancerous tissues due to the fact that cancerous cells hold more water, and consequently more hydrogen atoms than normal cells [Damadian et al., 1974]. This idea motivated Paul Lauterbur to use magnetic resonance to image living tissues [Lauterbur, 1973], and together with Peter Mansfield, introduced magnetic field gradients in three different directions to idenfity localized sources of NMR signal, enabling them to produce high resolution 3D images of bodily tissues.

To put simply, MRI relies on magnetic fields, radio-frequency (RF) pulse, and their interactions with protons found in water that makes up living tissues – for a comprehensive review, see [Currie et al., 2013]. We can therefore think of the hydrogen nuclei as the main driving element in producing signals in MR-brain imaging. It consists of a proton which spins in random directions, and whose motion can be likened to tiny toy tops that wobble as they spin. This movement causes it to generate its own magnetic field, known as its magnetic moment, and

consequently behave like little bars of magnets (see Figure 2.1(B)).

When a subject lies down in the bore of the MRI scanner, a static magnetic field (B_0) is introduced along the antero-posterior direction or the z-axis as shown in Fig 2.1(A). The presence of this field triggers the realignment of individual hydrogen nuclei to either anti-parallel or parallel to B_0 (Figure 2.1(C)) causing them to precess around this axis, but at different phases. When all magnetic moments are summed up (Figure 2.1(D)), the overall magnetization M_L of these group of protons, known as longitudinal magnetization, point towards the direction of B_0. Since these protons precess, there is no net polarization in the $x-$ or $y-$axes, however, one particular proton at any moment in time will be pointing in some direction in the xy plane. The number of times that these protons spin within a second is known as the Lamor frequency, which is proportional to the magnetic field according to $\omega_0 = \gamma B_0$, where γ is the gyromagnetic ratio unique to each nuclei (*i.e.*, for hydrogen, $\gamma = 42.57$).

Now, when a radio-frequency (RF) pulse oscillating at the transverse plane is applied, the orientation of these protons are disturbed and they fall off from their alignment with B_0. This only happens when the frequency of the RF pulse matches the precessional frequency of the proton, a phenomennon called *resonance*. The transfer of energy from the RF pulse to the protons causes some of them to (1) gain energy and flip their alignment anti-parallel to B_0 and to (2) spin in phase so that all protons are in the same orientation at all times. The first effect increases the number of opposing protons and consequently reduces the amplitude of the longitudinal magnetization. The second effect induces a moving net polarization in the transverse direction to B_0, known as the transverse magnetization (M_T). When the RF pulse is turned off, the protons relax and go back to their alignment with B_0. The process of relaxation occurs in terms of two aspects: (1) M_L starts to go back to its original value after the application of the RF pulse, and (2) M_T disappears as protons fall out of phase with each other due to magnetic field variations (inhomogeneity) within local tissue. Another reason for the loss of phase coherence among group of photons along the x–y plane is due to inhomogeneity within B_0. Magnetic field variations cause slightly different Larmor frequencies for protons at different locations within the field. This process describing the effects that result from the combination of T2 relaxation and the de-phasing due to inhomogeneity in B0 is characterized by the T2* (T2 star) relaxation. The evolution of these relaxations follow an exponential relationship of $M_L(t) = M_0(1 - e^{-t/T1})$ and $M_T(t) = M_0 e^{-t/T2}$, with time constants T_1 and T_2, respectively [Block et al., 1946].

Image Localization of MRI using Gradient Fields

The MR signal comes from the indirect detection of moving transverse magnetization rotating at the Larmor frequency. This induces an alternating voltage in a conductive receiver coil which in turn generates electrical current. The goal is to convert this signal into a resolved three-dimensional picture of the body. Just as how pixels comprise each square of a 2D picture, MRI obtains a 3D image that is made up of tiny *voxels* of different signal intensities. To do this,

MRI, in addition to the basic principles of NMR discussed above, introduces spatially varying gradient fields which allow for the detection of localized signals in living tissues. These are applied along the $x-$, $y-$, and $z-$ directions. Recall that protons only exchange energy with one another efficiently if the frequency of RF pulse matches their precession frequency. When a magnetic field gradient is introduced within a specific location or slice along the $z-$direction, the protons in this area experience a slightly different magnetic field, resulting to a gradient of precession frequencies that differ. The scanner can therefore select a particular slice to image by altering the frequency of the excitation pulses in order to match the frequency at the desired slice position. Meanwhile, protons next to the considered slice which precess at a different frequency will remain unaffected. Next, to identify the intensity values along the xy-plane, gradients along the $x-$ and $y-$ directions, called *frequency-encoding* and *phase-encoding gradients*, respectively, are used [Kumar et al., 1975]. The gradients in the fields along these directions induce different spin frequencies and phases across the protons, respectively. By modifying the temporal sequence of these gradients, the resulting signals encoded in the frequency domain, the so-called the *k-space*, can be converted to a two-dimensional image following an application of the inverse Fourier transform. This process is done for each slice along the z-direction until a full 3D brain volume is obtained.

2.1.1 Acquisition of functional MRI to monitor neurovascular coupling: a proxy for neural activity

The addition of the three-step gradient pulses causes de-phasing of the transverse magnetisation. In order to off-set this result and to ensure maximum signal intensity, a process known as *gradient echo* is performed. That is, two additional gradient pulses are applied: one immediately after the slice-selection and the other before the frequency-encoding [Winkler et al., 1975]. Next, to obtain a series of 3D brain volumes, the whole process discussed above, from the excitation to localization, is repeated at every constant and short temporal window. The time it takes between two successive acquisition is called the Repetition Time (TR), which is typically around 2 seconds for functional MRI acquired in regular clinical research studies. Functional data used for most part of this thesis are acquired using a much faster technique (TR < 1 second) and widely applied approach called *echo-planar imaging* [Mansfield, 1977].

The role of physiology

Functional brain imaging relies on the change in the oxygen state of the haemoglobin in the blood [Ogawa et al., 1990], which is based on the metabolical processes underlying brain activation, the so-called *neurovascular coupling*. Neural signaling in the brain requires a local increase in energy consumption, which in turn results to an up-regulated cerebral metabolic rate of oxygen at the affected brain area. By luck of nature, it turned out that oxygenated haemoglobin is diamagnetic, similar to most tissues in the brain, while deoxygenated haemoglobin is strongly paramagnetic due to the presence of unpaired electrons. This induces local magnetic field gradients detected by the MRI whose strength depends on the deoxy-

genated haemoglobin concentration. Recall from the discussion above that T2* relaxation is caused by the magnetic field inhomogeneities. These depend upon the physiological state of the brain, and in particular the composition of the local blood supply. Functional MRI, through the detection of the T2*, therefore relies on the increased local cerebral blood flow (CBF) and changes in oxygenation concentration (Blood Oxygen Level Dependent, or BOLD contrast), following neuronal activation. Thus, the signal we capture is not a direct measure of neural activity, but rather only a *proxy* of it. Because fMRI relies on haemodynamic changes, the signal that is measured is not instantaneous, and instead follow a 5-second delay described by the haemodynamic response function (HRF) [Logothetis et al., 2001].

2.1.2 Acquisition of structural and diffusion MRI to capture brain morphology and white matter connectivity

Gray matter (GM) contains cell bodies, dendrites, and axon terminals of neurons, while the white matter (WM) is made of long myelinated axons or nerve fibers. These two tissues can be distinguished by manipulating the imaging parameters that is set before a MRI scan is launched. In structural and diffusion MRI, the TR and TE are the primary parameters that the MRI technician chooses in order to influence the tissue weighting of the image to be generated. T2 relaxation characterized by the time constant T2 varies across different tissue types. Bone and lungs, for example, have short T2 values, while water and cerebrospinal fluids (CSF) have longer relaxation times and therefore long T2. Meanwhile, WM and GM have relatively shorter T2 values than CSF, but longer than bones and lungs. This difference in T2 values allows neighboring tissue types in the brain to be differentiated from one another resulting to a contrast image. Complementary, different tissue types also have different T1 relaxation times. An image contrast that is predominantly due to differences in tissue T1 relaxation time is called a *T1-weighted image*. One of the most widely used imaging setting in acquiring high-resolution anatomical data uses a T1-based weighting is called the Magnetization Prepared Rapid Gradient Echo (MPRAGE) [Mugler III and Brookeman, 1991]. In most part of this thesis, we simply use T1 when we refer to the anatomical/structural brain data.

WM can be thought as a highway that acts as a relay to coordinate communication between different GM areas. It's physically made of bundles or *tracts*. When MRI images the brain, it captures the diffusion of water molecules as they traverse long narrow tissues. In this scenario, the diffusion is anisotropic and the water molecules are expected to follow a specific fiber direction. Diffusion MRI utilizes a type of imaging sequence called the *spin-echo*, which is the refocusing of spin magnetisation by introducing a 180 degrees RF pulse, opposite to B_0, after an initial 90 degrees excitation [Hahn, 1950]. This causes the protons to precess in the opposite direction, such that those who spin the slowest are suddenly in the front. At some point in time, the ones who spin fast catch up and all the protons begin to spin in phase, resulting to a maximal signal intensity.

2.2 Probing brain function in task and rest

After the acquisition of fMRI data, we are usually left to deal with noisy 4D volumes (temporal sequence of 3D volumes). Basic preprocessing steps are as follows: (1) Motion correction is performed by realigning all frames to their mean functional image; (2) These volumes are then co-registered to a higher resolution T1-weigthed image, which is (3) subsequently segmented to obtain tissue probability maps for GM, WM, and CSF; (4) Co-registered volumes undergo spatial smoothing using a Gaussian filter with a full-width-half-max (FWHM) within the range of 4-6 mm; (5) Because each brain has different size and shape, all functional volumes are normalized to a common template or coordinates so that group-level analysis can be performed. There are many variants of pre-processing steps, for a review see [Van Dijk et al., 2010]. In this thesis, we used these basic preprocessing steps, and are followed by nuissance regression (6 motion parameters, signal drifts, and WM and CSF mean signals), except when stated differently.

First studies of fMRI relied on task-based paradigms, where functional brain changes induced by an applied stimulus were mapped directly using regression analysis, the most widely used of which is the General Linear Model (GLM) [Friston et al., 1994]. To describe simply, subjects are subjected to an external stimulus or task during specified blocks (task blocks), and are asked not to do anything during the control epochs. Statistical contrast of BOLD signal during the task block is then generated with respect to the control blocks. The resulting maps show specific brain regions that are attributed to the specific stimulus/task. Because of direct mappings between neural and cognitive states, many open-source meta-analytical tools are available, such as the brainmap.org [Fox and Lancaster, 2002] and the Neurosynth database [Yarkoni et al., 2011], wherein an automated brain-mapping framework is performed in order to support a more generalized decoding of broad cognitive states from brain activity.

2.2.1 Time-averaging quantification of brain function

Regression analyses are classic examples of static quantification of brain function and are particularly useful for mapping spatial maps related to different epochs of a task paradigm. However, without the predefined timings of the task that guide post-experiment analyses, the GLM approach cannot be used. The discovery that BOLD signal obtained from a *resting* brain is not just noise, and rather a manifestation of meaningful functional connectivity of the brain [Biswal et al., 1995] has led to a number of studies that attempted to obtain large-scale functional brain networks derived from spontaneous brain recordings – for a review of static FC approaches, see [van den Heuvel and Hulshoff Pol, 2010] and [Smitha et al., 2017].

One of the simplest and classical ways to extract static spatial maps is through the **seed-based correlation** analysis. This was method adopted by [Biswal et al., 1995] and colleagues when they identified regions of the sensorimotor cortex to be activated secondary to hand movement. Seed-based analysis is a model-based method in which users are able to select a seed or region of interest (ROI), and of which mean BOLD signals are linearly correlated to all

Figure 2.2 – **Intrinsic FC networks during rest** reproduced from a previous study [Thomas Yeo et al., 2011] using a large cohort of subjects (N=1000). Local networks are confined to sensory and motor cortices while association networks are composed of spatially distributed regions.

.

other voxels in the entire brain, resulting to a seed-based FC map. Due to the need to specify a seed, it is not possible to extract multiple brain patterns at once. Nevertheless, such approach has led to the discovery of a wide array of resting-state networks (RSNs), such as primary and secondary visual, temporal/ACC network, left and right fronto-parietal networks, to name a few. If instead of brain voxels, we consider each region as a seed in an atlased data consisting of N regions, then taking the Pearson correlation of each seed to the rest of all other regions will generate $N \times (N-1)/2$ values describing the functional interplay of each pair of regions. This is the classic definition of a *functional connectome* or *FC*.

Another method that is widely used in understanding the functional architecture of the brain is **graph theory**. A graph, mathetically represented as $\mathcal{G} := (\mathcal{N}, \mathbf{A})$, models brain regions as nodes ($\mathcal{N}$), where $|\mathcal{N}| = N$, and the degree of their functional interplay as edges. The

connectivity information is described in an adjacency matrix **A**, which in fact, constitutes the classic definition of FC (*i.e.*,, statistical interdependence of brain areas [Friston, 1994]). Using different graph metrics (*e.g.*, shortest path, degree, clustering, modularity), the overall functional organization of the brain can be examined and insights in how the human brain operates can be derived [Bullmore and Sporns, 2009, Achard and Bullmore, 2007].

2.2.2 Time-resolved quantification of brain function

For a while, FC analyses with static assumptions dominated the field of computational neuroscience. However, since the discovery that fluctuations in the BOLD signal can be meaningful [Chang and Glover, 2010], most methodological developments for extracting FC has taken the dynamic perspective into account – for a comprehensive review of dynamic FC methods, see [Preti et al., 2017] and for the dFC's relation to cognition and behavior, see [Bolton et al., 2020].

Most of the early dynamic FC methods were extensions of the static approaches. One perfect example is the **sliding-window analysis**, which computes FC within each temporal window that is cascaded at a constant time step over the whole duration of a resting-state scan. Due to its simplicity, many studies utilized it to uncover intrinsic functional brain dynamics [Kucyi and Davis, 2014, Elton and Gao, 2015]. An important parameter in this technique is the user-defined temporal window length. Too short windows would increase the risk of introducing spurious fluctuations, while too long windows would not detect the meaningful temporal variations that are being investigated [Leonardi and Van De Ville, 2015].

As an extension of sliding-window, which can also be viewed as an extension of a static graph-based analysis, is **dynamic graph analysis**, in which graph measures are computed for each temporal window of FC modeled as a graph, so that a temporal evolution of various graph metrics (*e.g.*, modularity [Betzel et al., 2016b] can be obtained [Tagliazucchi et al., 2012b]. Further extensions include decomposition of dynamic FC matrices into distinct connectivity states using principal component analysis (PCA) [Leonardi et al., 2013] or singular value decomposition (SVD), k-means clustering [Allen et al., 2014], or hierarchical clustering [Yang et al., 2014]. Meanwhile, a direct extension of ICA is the approximation of the time-course for each individual independent components by back-reconstruction [Erhardt et al., 2011].

In the discussed methods above, information from all time-points are used, however, it is not for certain that the brain undergoes ever changing temporal dynamics. In the same time that sliding-window analysis was rising in popularity, Tagliazucchi and colleagues proposed the idea that equivalent results extracted using static FC analyses can be obtained by investigating only timepoints that show a relatively large amplitude [Tagliazucchi et al., 2012a, Tagliazucchi et al., 2016]. Known as the **point-process analysis (PPA)**, only relevant timepoints above a certain threshold contain the same amount of information as a regular full timecourse. This is advantageous because it decreases the amount of data needed to extract meaningful patterns, and it preserves the temporal information, unlike the averaging-effect of classic seed-analyses.

In fact, it was shown that if we take the average of all frames corresponding to the top 15% of mean BOLD activation within a chosen seed region, the result is equivalent of that of seed-correlation map [Liu and Duyn, 2013]. Moreover, if we run a clustering algorithm in the original frames whose mean BOLD signal within the seed were above the chosen threshold, we generate a set of **co-activation patterns (CAPs)** corresponding to the chosen ROI. Furthermore, if we consider the *transients* or moments of activity changes in the BOLD signal instead of the peaks as is done in regular CAPs, we can extract whole-brain FC networks without the need of an ROI, and at the same time, be able to capture the temporally overlapping nature of intrinsic brain organization [Karahanoglu and Van De Ville, 2015]. Known as the **innovation-driven co-activation patterns (iCAPs)**, this method takes into account the the effects of the HRF in the BOLD signal, and apply a deconvolution step followed by a derivative to obtain the transients [Karahanoğlu et al., 2013], which are then fed to a clustering algorithm.

From frame-resolved family of approaches, one may also consider methods of extracting **dFC states** using **temporal modeling**. Using *k-means clustering*, Allen and colleagues were able to extract recurring connectivity patterns (cluster centroids) which are temporally non-overlapping [Allen et al., 2014]. Alternatively, there is also the *hidden Markov Model (HMM)* approach which considers a temporal sequence of hidden states parameterized by a covariance matrix [Eavani et al., 2013]. Similarly, *principal component analysis* can be used to maximize the variance in resting-state data [Leonardi et al., 2013].

2.2.3 Dynamics of brain function in the descent to sleep

The study of sleep, as well as studies related to consciousness and different levels of arousal (e.g., anesthesia, and drugs), are notable areas of research that largely benefit from continuous development of data-driven methodological tools that allow to detect prevalent spatial brain patterns. Many seminal works, especially those that initially sought to understand the neurophysiological relevance of the default mode network (DMN), a set of regions known to be engaged when there is no extraneous task [Greicius et al., 2003], assumed that these RSNs reflect spontaneous ongoing cognition [Andrews-Hanna, 2012]. It was therefore hypothesized that if conscious awareness decreases as the brain transitions from wakefulness to deep sleep, then it should rightly follow that the regions of the brain associated with high cognition also decrease in activity. This hypothesis was tested by [Larson-Prior et al., 2009] and they observed that FC was maintained in each network throughout the descent from wakefulness to light NREM sleep, as can be seen in a side-by-side comparison of the networks in Figure 2.3(A). This was verified by a later work performed by Tagliazucchi et. al, which reported that even in deep sleep (N3), almost all of the observed RSNs during wakefulness can also be retrieved [Tagliazucchi et al., 2013]. These findings therefore oppose the general notion that RSNs merely reflect spontaneous ongoing cognition. Moreover, stronger evidence against this notion is provided by separate studies that demonstrate maintained RSNs even in cases where conscious awareness is presumed to be completely abolished, such as in subjects under anesthesia and even in coma [Boly et al., 2008].

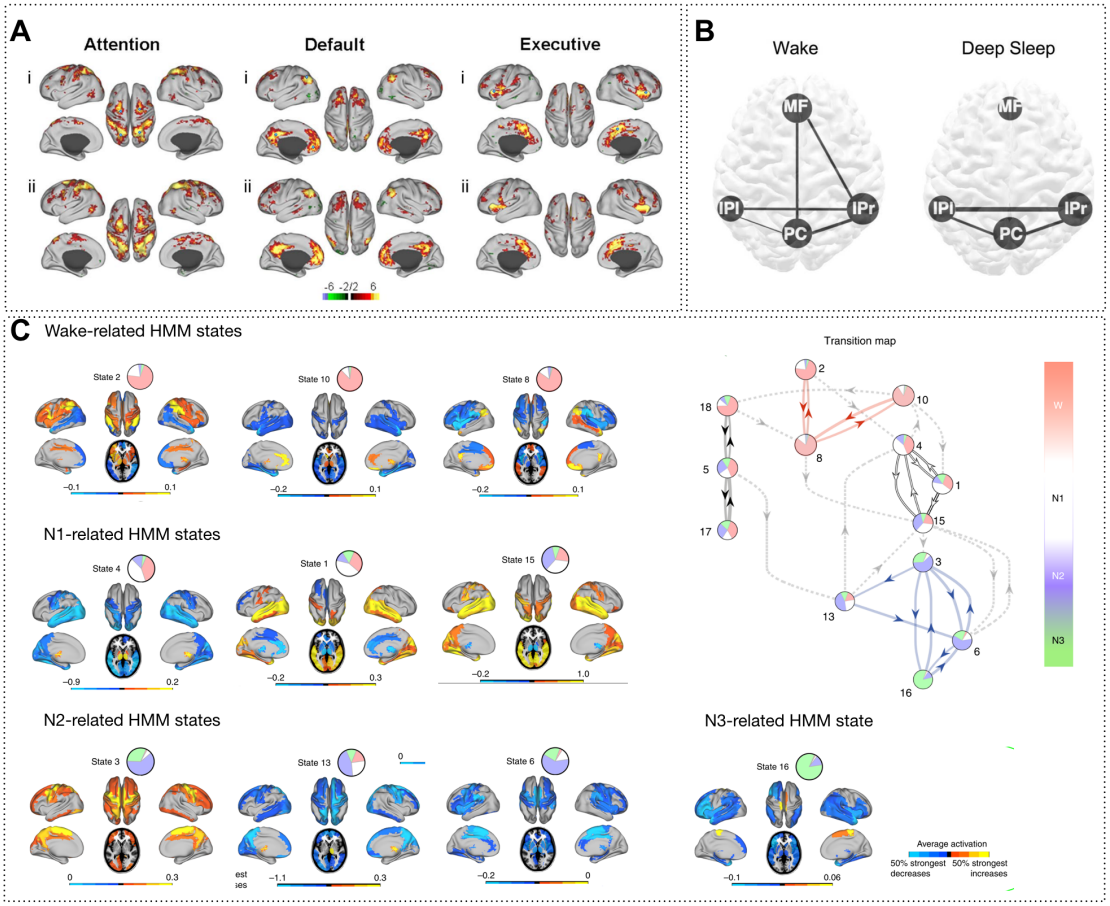

Figure 2.3 – **FC states in the descent to sleep**. (A) Using static seed correlation analysis, Larson-Prior and colleagues showed maintained RSNs in (i) wake and (ii) N1 sleep. (B) Meanwhile, it was also revealed that a dissociation between the posterior DMN and the anterior cingulate occurs when the brain transitions to deep sleep [Horovitz et al., 2009]. (C) Using a dynamic temporal modeling approach, the likelihood of different brain states to occur at different sleep stages and the key trajectories that the brain undergoes from wakefulness to deep sleep are presented [Stevner et al., 2019]. All figures are adapted with permission from the publishers.

Nevertheless, the maintained presence of these networks does not consequently mean that they remain the same all throughout different brain states. This is specially not surprising since the brain undergoes physiologically distinct changes which are particularly well-studied using electroencephalography (EEG). Our natural sleep can be broadly divided into rapid eye movement (REM) and non-rapid eye movement (NREM) sleep. The latter can further be subdivided into different sleep stages characterized by relaxed wakefulness (N1) to light sleep (N2), up to slow-wave sleep (SWS) or deep sleep (N3). Using static FC analysis, several works have attempted to compare the difference in FC patterns across NREM sleep stages. One very prominent finding is the dissociation of the DMN into posterior and anterior parts [Horovitz et al., 2009] as illustrated in Figure 2.3(B).

While static FC analyses has been very insightful, there are a lot of information yet to be explored in terms of the temporal dynamics of FC in the descent to sleep. Lately, an elegant framework that utilizes Bayesian-based temporal modelling has been introduced, which uses a HMM [Vidaurre et al., 2017] to estimate repeating spatial brain paterns in NREM sleep [Stevner et al., 2019]. They showed that there are specific FC networks that have higher propensity to occur at different sleep stages – see Figure 2.3(C). For instance, wake-related HMM states comprise regions from the thalamus, and some key regions of the DMN. Meanwhile, at the other extreme, deep sleep is characterized by an increased activation in somatosensory areas and a strong decrease of activation in the parietal and temporal cortices. These states were then elegantly laid out in terms of the tendency to shift from one state to another as the brain transitions from wakefulness and across all NREM sleep. One specific drawback of this method, however, is the strong assumption of temporal segregation between states, *i.e.*, only one state can occur per time instance. This is particularly questionable, especially in light of recent findings that FC networks are spatially and temporally overlapping [Karahanoglu and Van De Ville, 2015], and that the brain exhibits unstable synchronization in N2 sleep [Kung et al., 2019]. In the latter work, they investigated changes in variance of FC across temporal windows in a sliding-window analysis. They showed that FCs computed across all temporal windows within N2 sleep are highly fluctuating, suggesting instability in brain integration. In contrast, FC networks during N3 exhibit low variance thereby reflecting a low but stable inter-regional coordination [Kung et al., 2019].

In this thesis, we tackle changes in the dynamics of FC in wakefulness and across different NREM sleep stages using the iCAPs approach, which was introduced in the previous section. This method is able to extract spatially and temporally overlapping networks [Karahanoglu and Van De Ville, 2015], which gives a more accurate network-level representation of brain function. We will explore these changes and give interpretations of their physiological relevance, especially in terms of their relation to consciousness [Tononi, 2004, Tagliazucchi and van Someren, 2017].

2.3 Characterization of brain structure

2.3.1 Voxel-wise characterization of brain structure using scalar anisotropy measures

DTI models the diffusion of water molecules in nerve fibers using diffusion tensors represented by a $3 \times 3 \times 3$ matrix. A popular method to characterize brain structure is the use of **diffusion anisotropy measures** [Basser and Pierpaoli, 1996], such as fractional anisotropy (FA), apparent mean diffusivity (MD), axial diffusivity (AD), and radial diffusivity (RD). All of these measures are scalar valued and rotation-invariant. We can imagine diffusion tensors as ellipsoids of varying size and sphericity, which is determined by the comparative length of their three major cross-sectional axes (also known as the tensor's eigenvalues). AD is equivalent to the largest eigenvalue which is proportional to the length of the major axis of the ellipsoid. The

average between the two other eigenvalues is given by the RD. MD is the average of all three. Meanwhile, FA, the most widely used among all four, describes the degree of anisotropy in that particular voxel. This is proportional to the sum of squared differences of all eigenvalues to their mean. A value of zero would mean that the diffusion is isotropic and that the tensor is perfectly spherical in shape. On the other end, a value of one would mean that water diffusion occurs only along one axis and is fully restricted along all other directions. Due to a direct and easy interpreation of FA as a measure of white matter integrity in the area being investigated, several clinical research studies have adopted its use [Lerner et al., 2014], such as Alzheimer's disease [Nir et al., 2013], brain tumor [Jellison et al., 2004], and multiple sclerosis [Cercignani et al., 2000], to name a few.

2.3.2 Region-wise characterization of brain structure

Diffusion anisotropy measures, as described above, capture white matter integrity at the voxel level. A more advanced approach in characterizing white matter structures is the approximation of fiber streamlines using **tractography**. Signal reconstruction models of DW-MRI, such as the ODF and the tensor dictates the primary direction of water diffusion. In brief, tractography uses this local voxel-to-voxel information to trace the direction of fibers from one seed-position to another [Jbabdi et al., 2015]. There are many variants of tractography algorithms, and improving its accuracy remains to be one of the mostly sought research area in neuroscience [Maier-Hein et al., 2017].

The number of fiber streamlines connecting two regions of the brain gives an approximation of the strength of structural connectivity between them. It is then useful to summarize the results from tractography into a *structural connectome* or $N \times N$ structural connectivity (SC) matrix, where each entry a_{ij} describes the number of fiber streamlines connecting regions i and j. This is typically normalized by dividing the number of fibers with the average of the sizes of the two ROIs in question. From here, various analyses can be done to characterize brain structure. Just as how **graph theory** is useful in characterizing FC, it is also used in understanding the topology of the brain's anatomy. In particular, the human brain is described to have a *small-world* property, which combines high clustering and low path lengths [Bassett and Bullmore, 2006, Bassett and Bullmore, 2017] to maintain efficiency in inter-areal communication. This transfer of information is done, thanks to network hubs [van den Heuvel and Sporns, 2011] that are highly interconnected among themselves, forming a so-called *rich club*.

2.4 Structural basis of brain function

2.4.1 Current methods for linking brain structure and function

Both FC and SC describe the relationships between brain regions that are chosen using a predefined atlas. If this is chosen to be the same when computing FC and SC, then it is straightforward to compare the two by simply taking a direct **SC-FC correlation** for the whole

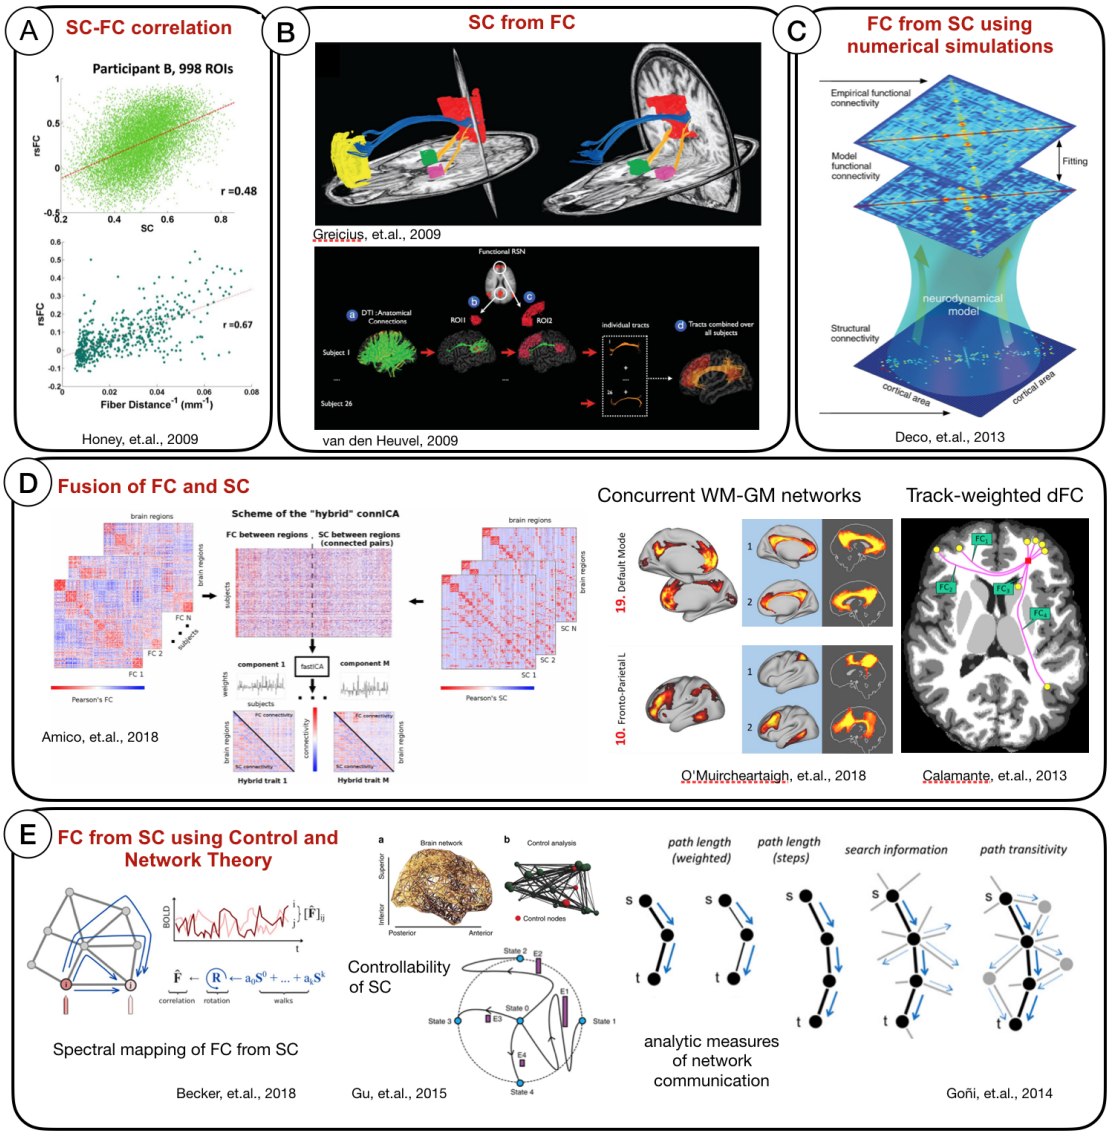

Figure 2.4 – **Overview of existing methods for linking SC and FC.** (A) By vectorizing the SC and FC matrices, one can directly compute the correlation between resting-state FC and (i) SC or (ii) the inverse of fiber distance [Honey et al., 2009]. (B) SC can be derived from predefined ROIs computed using classical FC [Greicius et al., 2009, van den Heuvel et al., 2009]. (C) In contrast, FC can be approximated from SC using numerical simulations [Deco et al., 2013]. (D) Alternatively, vectorized SC and FC measures can be combined together forming a hybrid datamatrix [Amico and Goñi, 2018] on which ICA can be applied to obtain concurrent WM-GM networks [O'Muircheartaigh and Jbabdi, 2018]. Meanwhile, a more integrated approach that maintains the dynamic view is to compute the average FC of voxel endpoints corresponding to all fibers that passed every single voxel in the WM [Calamante et al., 2017]. (E) Finally, more advanced methods borrowed from other research fields such as control theory [Gu et al., 2015] and network or graph theory [Becker et al., 2018, Goni et al., 2014] examine how FC arises from, is controlled by, and is constrained by measures of SC.

duration of the RS scan [Honey et al., 2009]. Using this simple measure, it was found that the persistence and strength of FC are constrained by the underlying anatomical bedrock, even in between regions that have no direct structural linkage. When this analysis is applied in different sliding temporal windows [Liégeois et al., 2016], FC patterns can be observed to oscillate between different phases of high and low modularity, which are primarily shaped by the anatomy and the inter-network connectivity, respectively.

Due to the bivariate nature of the correlation analysis, such measures are merely reflecting the general trend of the relationship between SC and FC. Another direct approach is to **approximate SC from FC** by initially specifying regions that are temporally coherent, then initialize one of these regions as a seed where fiber streamlines can be estimated using tractography [Greicius et al., 2009, van den Heuvel et al., 2009, Chamberland et al., 2015]. Using this approach, it was confirmed, for the first time, that major regions of the DMN, such as the medial pre-frontal cortex (MPFC) and the posterior cingulate cortext (PCC) are structurally connected via the cingulum, while the medial temporal lobe (MTL) is directly linked to the PCC via the hippocampal cingulum.

In contrast to extracting SC from predefined FC, one can also **approximate FC based on prior measures of SC** using numerical simulations [Galán, 2008, Honey et al., 2007, Deco et al., 2011]. Estimating FC from SC or vice versa are hypothesis-driven and requires many explicit assumptions at the beginning of the analyses. To overcome this limitation, several studies proposed methods that are data-driven, such as *ICA* applied to a hybrid data matrix composed of vectorized FC and SC measures [Amico and Goñi, 2018]. The output of this are joint functional-structural patterns that are task-dependent. Also with ICA, concurrent GM and WM network patterns were extracted from a whole-brain tractography data [O'Muircheartaigh and Jbabdi, 2018], where paired GM "nodes" are extracted together with associated anatomically meaningful white matter "edges". As in any ICA analysis, the choice of number of ICs remain a crucial parameter, which can be complicated to choose on top of the caveats that come with its principal assumption of spatial independence.

ICA applied to a hybrid data-matrix is one of the very few examples of a family of methods that **jointly integrate SC and FC**. Another example is that of Calamante and colleagues, the so-called **track-weighted dynamic FC (TW-dFC)**, where each voxel in the WM is assigned a certain value based on the average of all FC measures between fiber endpoints [Calamante et al., 2017]. Thus, as a first step, tractography is launched to approximate fiber bundles that connect one GM voxel to another. Some pairs of GM voxels may not have fiber streamlines connecting them. Then, for each voxel in WM, TW-dFC computes the FC between the voxel endpoints in GM of each fiber that passed through this WM voxel. The average of the computed FC from all fibers is recorded in each WM voxel. This process is repeated for each temporal sequence of a sliding window, forming a temporal sequence of 3D FC volumes.

An alternative family of techniques in relating FC and SC consists of advanced methods borrowed from other research fields to approximate SC's influence in formation of FC. Using

control and network theory, Gu and colleagues studied how the anatomy constrains the brain as it moves between cognitive states [Gu et al., 2015]. They found that structural network differences between functional subsystems determine their distinct role in brain function. Highly modular networks, as such as that of the default mode system facilitate movements of the brain to easily reachable states. Contrary, weakly connected areas such as in cognitive control systems support movement of the brain to difficult-to-reach states. Meanwhile, averaged modular systems such as attentional control systems dictate facilitate the integration or segration of cognitive systems. In line with these findings, using **spectral graph theory**, it has also been demonstrated that the amount of structural walks between two regions influence the formation of functional correlations [Becker et al., 2018], just as how analytical measures of network communication (*e.g.*, path transitivity) was found to predict resting brain activity [Goni et al., 2014].

2.4.2 Graph signal processing for functional brain imaging

Graph theory applied to brain studies remain a popular approach in characterizing both brain structure and function. In fact, it has delineated itself as a separate field of interdisciplinary research called *network neuroscience* [Bassett and Bullmore, 2017]. As discussed, structural or functional-based connectomes are typically defined separately, where vertices are associated to brain regions, and the links are weighted according to the strength of the physical connection, or the coherence of neural co-activation, e.g., pairwise Pearson correlation of brain activity between vertices, respectively. These connectomes are then analyzed using various graph measures in order to get a better understanding of the organization principles of brain networks, and determine its influence in cognitive function [Bullmore and Sporns, 2009].

Throughout the years, more interest has been given to addressing neuroscientific questions with a dynamical point of view, and along with it, lies the fundamental question of how the brain's fixed anatomy supports the emergence of spatially and temporally varying distributed patterns of brain activity. With this fundamental research question at hand, it is only fitting that network neuroscience advances to a more sophisticated approach of integrating brain structure and function. Borrowing from the emerging field of graph signal processing (GSP), a new and elegant framework has been recently introduced, where functional data are defined atop a fixed anatomical structure. These datapoints are interpreted as time-dependent graph signals on which structurally-informed signal processing methods can be performed [Huang et al., 2018a].

The GSP framework generalizes principles from classical discrete signal processing for time series to graph signal processing of signals defined on irregular structures. The most important entity to establish is the SC, from which the graph $\mathcal{G} := (\mathcal{N}, \mathbf{A})$ is defined, *i.e.*, SC is equivalent to $\mathbf{A}$. The graph Laplacian matrix is then computed as $\mathbf{L} = \mathbf{D} - \mathbf{A}$, where $\mathbf{D}$ is the degree matrix which contains the degree of each node on its diagonal: $D_{i,i} = \sum_{j \in \mathcal{N}} a_{i,j}$. This is essentially a discrete representation of a continuous second order derivative Δ that is used in many areas

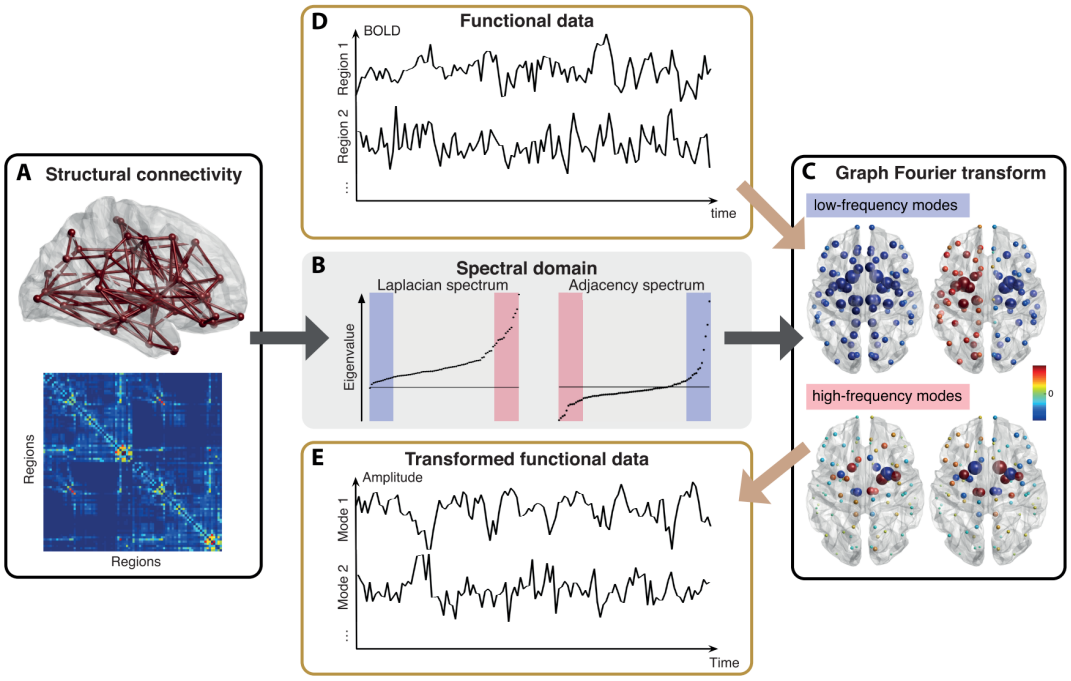

Figure 2.5 – **Graph signal processing overview for fMRI.** SC, constructed from DW-MRI, contains region-wise description of the brain's anatomy, which can be rewritten via the Laplacian. This can be decomposed into orthogonal basis functions called the eigenmodes together with their associated frequency/eigenvalues. Using low-frequency eigenmodes, slowly varying spatial patterns can be decomposed; meanwhile, highly varying spatial maps are made up of higher frequency eigenmodes. Using these modes that form the basis of Graph Fourier Transform (GFT), any functional data can be decomposed, which can then be analyzed using graph signal processing (GSP) tools. Figure is adapted with permission from [Huang et al., 2018a].

of physics and engineering.

The fundamental GSP concepts that are employed in neuroscience are **Graph Fourier Transforms (GFT)** and **graph filters**, both of which are rooted from spectral graph theory and build on the decomposition of the graph Laplacian matrix. Because it is a positive, semi-definite matrix, $\mathbf{L}$ can be diagonalized to satisfy $\mathbf{L} = \mathbf{UVU}^\top$, where $\mathbf{U}$ are the eigenvectors, or better known in neuroimaging as *eigenmodes*, and $\mathbf{V}$ contains the associated eigenvalues. Previous works have demonstrated that these eigenmodes are reminiscent of well-known functional networks [Atasoy et al., 2016]. In the same manner that sinusoids form a basis for Fourier transforms, these eigenvectors represent a building block for GFT, and the corresponding eigenvalues carry the notion of frequency, or spatial variability. Functional signals in both healthy and clinical populations can be decomposed through these connectome-specific components [Atasoy et al., 2017], and are related to behavior and cognition [Medaglia et al., 2018, Preti and Van De Ville, 2019]. The theory of GSP has been rapidly growing throughout the years [Shuman et al., 2013], and with these developments come new avenues for better

integration of brain structure and function. For a detailed review of GSP applied to functional brain imaging, see [Huang et al., 2018a].

3 Large-scale functional brain network dynamics in task, rest, and sleep

3.1 General Motivation

Before we go directly to methods development for integrating brain structure and function, we first introduce two novel approaches for time-resolved characterization of brain function. Since the inception of the idea that the brain is always active even at rest, several methods that attempt to characterize spontaneous brain activity have been proposed [van den Heuvel and Hulshoff Pol, 2010, Beckmann et al., 2005]. This goal has become even more exciting with the discovery that brain activity exhibits meaningful fluctuations [Chang and Glover, 2010]. As such, succeeding methodological developments were geared towards data-driven methods that also take into account the brain's dynamic nature [Preti et al., 2017, Tagliazucchi et al., 2012a, Liu and Duyn, 2013].

This chapter is dedicated to contributions that largely focus on the extraction of large-scale functional brain networks and the analysis of their temporal properties. In the first study, we investigate how the dynamics of functional brain networks are altered across the different NREM sleep stages. Many neuroimaging studies have used the brain's functional connectivity (FC) to investigate alterations in network integrity along transitions across various vigilance states [Larson-Prior et al., 2009, Picchioni et al., 2014, Horovitz et al., 2009]. These alterations are mostly observed in terms of *spatial* modifications of whole-brain connectivity patterns. In this work, we focus on the changes in the *temporal dynamics* of large-scale functional networks as the brain transitions from wakefulness to deep sleep. We use an advanced approach called innovation-driven co-activation pattern (iCAP) analysis that relies on transient brain activity or *moments of activity changes* in the BOLD signal. This is done by undoing the effect of the hemodynamics through spatio-temporal deconvolution [Karahanoğlu et al., 2013, Farouj et al., 2017] combined with a regular derivative. This allows us to observe a wide-array of functional networks that reveal important contributions to sleep and arousal. Unlike previous methodologies that assume temporally segregated networks [Stevner et al., 2019], the iCAP approach warrants the recovery of spatial maps that are both spatially and temporally overlapping, thus giving a more accurate description of functional brain organization during

sleep. We derive the activity time-courses of each network at the frame-wise resolution, and investigate how long they usually last. We also identify which networks are active at the same time. We investigate these temporal properties across the different NREM sleep stages and identify alterations in individual networks and in all iCAPs as a whole. We then relate the overall dynamics of the iCAPs to one of the most prominent theories regarding consciousness and its association with the brain's capacity to integrate information [Tononi, 2004]. Due to the unique nature of the iCAP approach to recover transient brain activity, we observe new features of information integration of consciousness during sleep, and provide a concrete evidence for the presence of unstable yet distributed global synchronization in NREM stage 2. Additional analyses probing the relationship of FC with structural connectivity (SC) is also discussed at the end of this study to investigate how the SC-FC link changes from wakefulness to deep sleep.

In the second study, we explore the effectiveness of a new methodology based on graph learning to identify functional brain networks from task and resting-state data. This new approach treats functional data as a collection of signals that live on multiple graphs. We layout the details of the Graph Laplacian Mixture Model (GLMM) and demonstrate its potential to extract network patterns that correspond well to the timings of task paradigms. We demonstrate that the extracted networks also consist of brain areas that are consistent with previously observed regions using classical regression analyses. We then apply the approach to resting-state data and extract well-known intrinsic networks. One notable advantage of the GLMM approach is its ability to recover learned graph matrices corresponding to each functional network. These graph matrices describe the direct interactions of the regions that comprise the associated networks, akin to precision matrices which have been shown to reveal stronger correlation to structural connectivity (SC) [Liégeois et al., 2020]. We then take advantage of the learned graph matrices by investigating their general link to SC [Honey et al., 2009, Liégeois et al., 2016]. We implement the classical method of directly comparing SC and FC using Pearson correlation.

3.2 Journal Article: NREM sleep stages specifically alter dynamical integration of large-scale brain networks

A. Tarun, D. Wainstein-Andriano, V. Sterpenich, M. Solms, L. Bayer, L. Perogamvros, M. Solms, N. Axmacher, S. Schwartz, D. Van De Ville.

Full Research Article Accepted for Publication in iScience (2020)

Author Contributions

A. Tarun performed all stages of the analysis, from conceptualization, methodology, validation, formal analysis, and writing and editing of the manuscript, **D. Wainstein-Andriano** provided the data for Study 2, **V. Sterpenich** provided the data and conducted the manual sleep scoring for Study 1, **L. Bayer** and **L. Perogamvros** both conducted the sleep scoring for Study 2, **M. Solms**, **N. Axmacher**, and **S. Schwartz** edited and reviewed the manuscript, and **D. Van De Ville** offered continuous guidance.

Abstract

Functional dissociations in the brain observed during non-rapid eye movement (NREM) sleep have been associated with reduced information integration and impaired consciousness that accompany increasing sleep depth. Here, we explored the dynamical properties of large-scale functional brain networks derived from transients or moments of activity changes in fMRI. Spatial brain maps that are largely related to the physiological organization of sleep and arousal display significant modifications in terms of their tendency to occur across wakefulness and NREM sleep. Unexpectedly, almost all networks predominated in activity during NREM stage 2 before an abrupt loss of activity is observed in NREM stage 3. Yet, we observed that functional connectivity and mutual dependencies between these networks progressively broke down with increasing sleep depth. Thus, the efficiency of information transfer during NREM stage 2 is low despite the high attempt to communicate. Critically, our approach provides relevant data for evaluating functional brain network integrity and our findings robustly support a significant advance in our neural models of human sleep and consciousness.

3.2.1 Introduction

Spontaneous brain activity, as assessed by resting-state functional Magnetic Resonance Imaging (fMRI), has provided key insights into the functional architecture of the brain. Resting-state networks (RSNs) identify sets of brain regions that exhibit synchronized fluctuations of activity over the whole duration of a resting-state session (typically 10-20 min of continuous scanning). The main hypothesis underlying most functional connectivity (FC) studies is that different RSNs reflect distinct ongoing cognitive/affective processes/states. For instance, the default mode network (DMN) typically shows reduced activity when subjects perform an externally oriented task [Greicius et al., 2003] and, contrastingly, the DMN becomes more engaged when self-referential processes or internal mentation predominate [Andrews-Hanna, 2012]. According to this interpretation of RSNs, we expect that if conscious awareness dissipates as the brain transitions from wakefulness to deep sleep, we should observe a parallel and net decrease in activity and/or FC across regions of the brain involved in higher-order cognitive operations, such as reasoning, monitoring, or metacognition. Yet, some early evidence also suggested that several RSNs, including those encompassing association networks, persisted or even increased their connectivity during the descent from wakefulness to light sleep [Larson-Prior et al., 2009] or during anesthesia and coma, when conscious awareness is presumed to be completely abolished [Boly et al., 2008]. Therefore, RSNs may reflect intrinsic dynamical properties of the brain's functional organization that are maintained across distinct levels of consciousness.

From a behavioral point of view, the brain in sleep undergoes marked and well-characterized physiological changes. Based on polysomnography (PSG), which is a combined use of EEG, electro-oculography, and electromyography, natural sleep can be broadly divided into rapid eye movement (REM) and non-rapid eye movement (NREM) periods. The latter is further subdivided into different sleep stages characterized by relaxed wakefulness (N1) to light sleep (N2), up to slow-wave sleep (SWS) or deep sleep (N3). It is therefore not surprising that, although RSNs may be detected across different sleep stages, their connectivity patterns [Horovitz et al., 2009, Sämann et al., 2011] and FC strengths undergo significant modifications [Tagliazucchi et al., 2013]. For instance, upon reaching N3, the DMN has been found to dissociate into subcomponents, with a decrease in the connectivity between the medial prefrontal cortex (MPFC) and the posterior cingulate cortex (PCC) [Horovitz et al., 2009]. Furthermore, several studies looked at changes in FC for regions associated with the reticular activation system (RAS) that regulate the physiological state of arousal during sleep (e.g. thalamus, hypothalamus) and reported decreased connectivity between these regions and the rest of the cortex during light and deep sleep [Hale et al., 2016, Picchioni et al., 2013, Tagliazucchi and Laufs, 2014]. These changes in brain network integrity, particularly its marked reduction from wakefulness to deep sleep, have been associated to diminished level of information integration [Tononi, 2004]. That is, when the brain switches to more local cortical processing, this would lead to a global loss of information integration and a concomitant reduction in consciousness [Tononi and Koch, 2015].

Methodological tools allowing the detection of prevalent brain spatial patterns, such as FC

analyses, are particularly useful when assessing imaging data collected outside of a predefined experimental paradigm (or in the absence of any temporally-defined independent variable), namely when regression techniques cannot be applied. This is typically the case for continuous data collected during varying levels of arousal or states of consciousness (e.g. resting-state, sleep, anesthesia, and drugs). The simplest approach to investigating changes in FC is to use a sliding-window technique, where time-courses from sets of brain regions (e.g. from atlas-based parcellation) are segmented into successive temporal windows so that various assessments of FC (e.g. bivariate Pearson correlations) can be applied to obtain time-evolving connectivity matrices. Another approach is to derive analogous information on resting-state FC based on time-points where the regional BOLD signal exceeds a particular threshold [Tagliazucchi et al., 2012a]. Temporal clustering can also be applied to activity patterns occurring at these active fMRI time-points to obtain patterns of co-activity among regions, also known as co-activation patterns (CAPs) [Liu and Duyn, 2013].

Furthermore, to account for the fundamental dynamic nature of the changes in neural FC, the latest developments on non-stationary FC approaches have started to successfully incorporate methods of temporal modeling. Using a dynamic Bayesian approach (i.e., Hidden Markov Model; HMM), Stevner and colleagues [Stevner et al., 2019] could recently demonstrate that some specific whole-brain functional connections are associated to each of the different stages of NREM. They found that networks with high specificity to occur in stages N2 and N3 generally expressed longer mean lifetimes, with each HMM state lasting from a few seconds to tens of seconds. They also examined the transition probabilities between the networks by extracting modules of HMM states that transitioned more often between each other than to other states [Vidaurre et al., 2017], allowing them to identify key trajectories of network activity from wake to NREM sleep.

Despite these major methodological advancements, it remains unsettling that coordinated network activity is assumed to be strictly temporally segregated, and that whole-brain states are not overlapping in time (i.e., only one RSN can occur per time instance). To overcome this limitation, we recently proposed to consider innovation-driven co-activation patterns (iCAPS), which capture transient brain activity, i.e., physiologically significant moments of regional activation and deactivation [Karahanoğlu et al., 2013, Karahanoglu and Van De Ville, 2015], instead of the actual activity time-points. This framework allows for the recovery of RSNs that are both spatially and temporally overlapping, providing a more plausible and thus putatively more accurate description of functional brain organization. Furthermore, unlike standard CAPs and other data-driven approaches (e.g., independent component analysis [Smith et al., 2012]), the iCAP approach explicitly accounts for temporal blurring by the hemodynamic response function (HRF) by incorporating a deconvolution step in the preprocessing of the fMRI signal.

Here, we adopt an iCAP approach to recover functional brain networks from fMRI data recorded during wakefulness and NREM sleep stages. We obtain temporally overlapping activity time-courses of each network for each participant and evaluate their temporal prop-

erties to capture important features of cortical organization as the brain goes from a fully awake state to deep sleep. We observe that resting-state networks such as control, salience, visual and sensory networks [Shirer et al., 2012] persist until the deep sleep, and that some of these networks dissociate into subcomponents (e.g., DMN). Moreover, since the activity time-courses are not constrained to be temporally segregated, many networks tend to overlap in time, and that these numbers are altered across the different levels of sleep. To test these hypotheses, temporal measures, such as their overall accumulated durations and mean lifetimes, were computed, as well as their co-occurrences and temporal overlaps. Using a more accurate model to extract network-level representations of brain organization, we put the focal point into detecting the temporal alterations of large-scale brain activity across the distinct vigilance states.

3.2.2 Materials and Methods

Functional MRI Datasets

We used two complementary sets of simultaneously recorded EEG-fMRI data acquired in the context of two different studies. Ethics approval was obtained for each study and all participants signed written informed consent forms. The first dataset came from a study on sustained sleep (denoted from here onwards as Study 1), where participants were allowed to sleep for as long as they could manage, while EEG and fMRI data were simultaneously and continuously recorded. Twenty-six right-handed healthy participants participated in this study. All of the participants filled out questionnaires and underwent a semi-structured interview that established the absence of neurological, psychiatric, or sleep disorders. They were non-smokers, moderate caffeine drinkers, and were not taking any medication. The experiment was conducted at around 10 pm and the total recording time lasted between 51 min and 2 hrs 40 min (mean: 1 hr 43 min). We excluded a total of 5 participants who were not able to sleep inside the scanner (resulting in 21 subjects, mean age ± SD: 22 ±2.4 years; 15 women).

The second dataset focused on the sleep onset period (denoted below as Study 2). Ten healthy subjects (4 female) underwent simultaneous EEG-fMRI recording. All participants had no history of neurological or psychiatric disease, no previous or current use of psychoactive drugs, and were non-smokers. The experiment was conducted at around 9:30 pm and the awakening recording lasted 1 hr 30 mins. Participants were placed in the MRI scanner and asked to fall asleep. Successive awakenings of up to 10 rounds per session were performed during this period to record possible dreaming that occurred. The whole experiment included around 4 to 5 sessions (5 nights) per participant. As expected from this experimental design, in the majority of the recordings, participants were in the awake state to light sleep (N1). Figure 1 displays the two experimental paradigms, together with the distribution of sleep states reached by participants. We excluded 3 subjects due to difficulty in scoring the EEG data (7 subjects remaining, mean age ± SD: 20.5 ± 1 year; 5 women). In the present work, we did not investigate

the dream reports collected in Study 2 because our main goal was to characterize dynamic neurophysiological changes occurring from wakefulness to deep sleep, and also because we did not collect dream data in Study 1 (waking up participants would have prevented them from reaching deeper sleep stages).

EEG preprocessing and sleep scoring

For the data in Study 1, the EEG setup included a 64-channel MRI-compatible EEG cap, two pairs of electrocardiogram (ECG), horizontal and vertical electrooculography (EOG), and 1 pair of chin electromyographic (EMG) electrodes (BrainAmp MR plus, Brain Products GmbH, Gilching, Germany). For the data in Study 2, a 14-channel MR-compatible system (Brain Products GmbH, Gilching, Germany) was used, plus a total of ten cortical (EEG) electrodes. Two of the ten cortical EEG electrodes served to measure EOG, three EMG chin electrodes, and one ECG electrode on the back were used to monitor participants' sleep patterns. The preprocessing of both datasets was done using the Brain Analyzer software (Brain Products GmbH, Gilching, Germany). Specifically, gradient artifacts were removed offline using a sliding average of 21 averages and subsequently, the EEG data, initially sampled at 5000 Hz, was down-sampled to 500 Hz and low-pass filtered with a finite-impulse response filter with a bandwidth of 70 Hz. For the 64-channel recordings in Study 1, ICA was the most reliable method to remove ballistocardiogram and oculo-motor artifact, while a template subtraction method was used for the 14-electrode recordings in Study 2.

The preprocessed EEG data for Study 1 and Study 2 were then scored by three experts each, according to standardized American Academy of Sleep Medicine (AASM) polysomnographic criteria for sleep scoring (Iber 2007).

Functional MRI acquisition and preprocessing

For Study 1, MRI data were acquired on a 3 Tesla whole-body MR scanner (Tim Trio, Siemens, Erlangen, Germany) using a 12-channel head coil. Functional images were acquired with a gradient-echo EPI sequence (repetition time [TR]/ echo time [TE]/flip angle = 2100 ms/40 ms/90) and parallel imaging (GRAPPA; acceleration factor = 2). Each functional image comprised 32 axial slices (thickness = 3.2 mm without gap, FOV = 235 × 235 mm, matrix size = 128 × 84, voxel size: 3.2 × 3.2 × 3.84 mm,) oriented parallel to the inferior edge of the occipital and temporal lobes. On average 2789 functional images were recorded during one continuous scanning session (between 1459 and 3589 images). Structural images were acquired with a T1-weighted 3D sequence (MPRAGE, TR/inversion time [TI]/TE/flip angle = 1900 ms/900 ms/2.32 ms/ 9, FOV = 230 × 230 × 173 mm3, matrix size = 256 × 246 × 192 voxels, voxel size = 0.9 mm isotropic).

For Study 2, MRI recording was performed on a 3 Tesla Philips Achieva MRI scanner. Functional images were acquired using a gradient-echo EPI sequence (repetition time [TR]/ echo time

[TE]/ flip angle = 3000 ms/ 30 ms/ 83). Each functional image comprised 50 axial slices (thickness = 2.5 mm without gap, FOV = 96 mm x 96 mm, voxel size: 2.5 mm isotropic) oriented parallel to the inferior edge of the occipital and temporal lobes. Structural images were acquired with a T1-weighted 3D sequence (TR/inversion time [TI]/TE/flip angle = 1570 ms/8.4 ms/3.42 ms/ 8, FOV = 256 × 256 × 220 mm3, matrix size = 256 × 256 × 220 voxels, voxel size = 0.929 mm x 0.929 mm x 1mm).

Both functional datasets were first preprocessed following standard procedures [Van Dijk et al., 2010]. Functional volumes were realigned to their mean images using SPM12[1]. The realigned functional images underwent nuisance regression (mean white matter and cerebrospinal fluid, constant, linear, and quadratic drifts). We then applied spatial smoothing using an isotropic Gaussian kernel of 6 mm full width at half maximum. For the final analyses, the first 10 volumes were discarded to achieve steady-state magnetization of the fMRI data. To correct for large motion, we applied scrubbing [Power et al., 2012] in the functional data by marking time-points with framewise displacements of more than 0.5 mm. The iCAPs framework require a constant sampling rate, therefore, marked frames were replaced by interpolating the disconnected frames using spline interpolation (Karahanoğlu et al. 2013). Finally, structural scans (T1) were co-registered to the mean functional image, and the transformed volume was segmented using SPM 12 Segmentation to obtain probabilistic gray matter masks.

Total Activation and iCAPs framework

We applied the Total Activation (TA) procedure to the subject-space fMRI recordings for each participant of Study 1 and Study 2 datasets. The TA procedure started by a deconvolution of fMRI time-series at each voxel to remove hemodynamic effects [Karahanoğlu et al., 2013,Farouj et al., 2017]. Temporal transitions or transients (changes in activity) were then derived for each participant in Study 1 and each session of each night in Study 2. Highly significant transient frames (functional images) were identified and normalized to MNI coordinate space. The normalized functional images of all subjects were then concatenated, forming one single data matrix with all significant transients. The data matrix with a dimension of number of voxels times number of significant transients was fed into a clustering procedure to obtain temporally co-activating brain patterns. Because the functional networks were obtained by clustering significant transient brain activity, also known herewith as innovations, the resulting co-activation patterns are referred to as innovation-driven co-activation patterns, hence the name iCAPs. The optimal number of clusters used in the clustering procedure was determined using consensus clustering [Monti et al., 2003], which implements a multiple subsampling procedure. We evaluated multiple numbers of cluster values in the set K = {10, 11, 12, . . . , 25}. The optimal K was selected based on different stability metrics. One measure is the cluster consensus which reflects the consistency of a transient frame to be clustered in the same iCAP. After obtaining stable iCAPs, we used a spatio-temporal transient-informed regression [Zoller et al., 2019] to extract the activity time-courses of iCAPs for each participant. The TA and iCAPs

[1] http://www.fil.ion.ucl.ac.uk/spm/

pipelines are publicly available[2], and additional information as well as the mathematical details of the whole method are laid out in Appendix A.1.1.

Extraction of temporal characteristics of iCAPs in each sleep stages

Comparison of the iCAPs time-courses across different sleep stages was performed by z-scoring the entire iCAP time-courses per participant's recording, and thresholding at $|z| > 1.5$. The time points that survived thresholding were considered as "active". The choice of threshold was motivated by previous works that implemented TA and iCAPs framework [Karahanoglu and Van De Ville, 2015, Zöller et al., 2019]. We then computed the temporal characteristics (e.g., cumulated durations, average durations, temporal overlaps, and network interactions or couplings) of iCAPs in each sleep stage. Scrubbed frames (i.e., those with framewise displacement more than 0.5 mm) are taken into account and are not included in each temporal measure explored. The relative cumulated durations (RCD) computed in percentage pertains to the likelihood of an iCAP to occur in wakefulness and in each sleep stage. This is computed by counting the number of time-points that an iCAP is active divided by the total length of time a participant spent in that particular sleep stage. The overall RCD can go above 100% due to the overlapping nature of the iCAP time-courses (i.e., more than one iCAP can occur at one time-point). The average duration, measured in seconds, is the length of time that an iCAP is continuously active. We also evaluated the interaction of iCAPs with other networks by computing the overall percentage of temporal overlaps in wakefulness and sleep, as well as the number of pair-wise iCAP co-activations. The latter is represented by the co-occurrence which reflects the number of time-points during which a pair of iCAPs were both active divided by the total number of time-points that at least either one of them was active. We take into account the signs of the activations (similarly signed or opposite signed). Z-scoring the time-courses allows positive and negative activation signs. Positive (negative) activation signs reflect the time-points when an iCAP is at least 1.5 standard deviations above (below) the overall mean amplitude of iCAPs.

Statistical analysis

Duration measures of individual iCAPs were compared across the different sleep stages using paired t-tests and 1000 rounds of permutation testing (randomly shuffled). The corresponding t-statistics, p-values and effect sizes are displayed in Table. A.2. The overall comparison of iCAP cumulated durations (general trend), network temporal overlap and coupling indices were done through the Analysis of Variance (ANOVA), and the p-values were obtained using successive multiple comparison test (Tukey's range test). The corresponding F-statistics (ANOVA) and the corrected p-values for the temporal overlap and coupling indices are displayed in Tables A.3 3 and A.4, respectively. For the analysis of temporal characteristics (e.g., average durations, RCD, temporal overlaps, and couplings), we only included subjects who reached

[2]Code Available on: https: //c4science.ch/source/iCAPs/

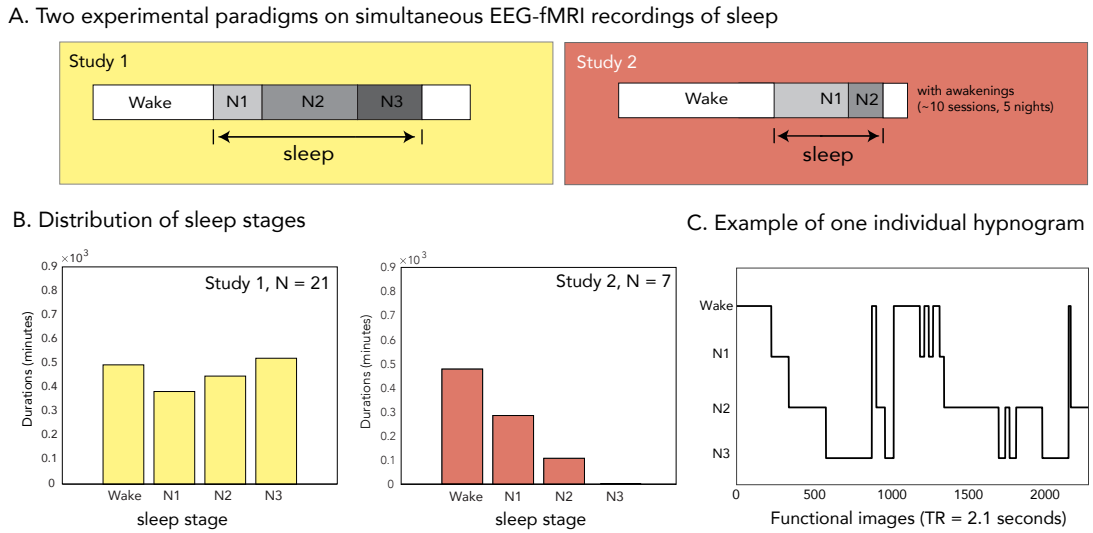

Figure 3.1 – **Experimental paradigm.** (A) The two experimental paradigms are the (i) Sustained Sleep lasting between 51 mins to 2 hrs 40 mins and the (ii) Sleep Onset experiment lasting between 5 minutes up to 20 minutes. (B) The distribution of sleep states shows how long the subjects lingered in each sleep state, with Sustained Sleep covering up to N3, while Sleep Onset covers mostly wake and N1, and a small duration in N2. (C) Example of sleep scoring (or hypnogram) for one participant in the Sustained Sleep experiment.

N3, and thus we only limit this part of the analysis to Study 1.

3.2.3 Results

Data distribution across Study 1 and Study 2

Figure 3.1 illustrates the two experimental paradigms on simultaneously acquired EEG and fMRI recordings of sleep. We used a total of 21 subjects in Study 1 (13 reached N3), and a total of 7 subjects in the Study 2 (all of them reached N1 sleep). The distribution of accumulated data from each sleep stage for each dataset is also shown in Figure 3.1, with Study 1 generating a substantial amount of N3 sleep (more than 8 hours in total), while Study 2 was predominated by wake and N1 sleep (about 8h and 4h of data, respectively).

Spatial patterns in sleep and wake

We applied a deconvolution process called total activation (TA) to all functional volumes of Study 1 and Study 2. Transient frames (i.e., moments of activity changes) are extracted through a derivative step that is incorporated in the TA framework. These frames correspond to timepoints when a change is observed in the amplitude of the fMRI BOLD signal. Significant transient frames are then selected using a two-step thresholding process (see Supplementary Methods or ‘ [Karahanoglu and Van De Ville, 2015]. Frames that survived are concatenated

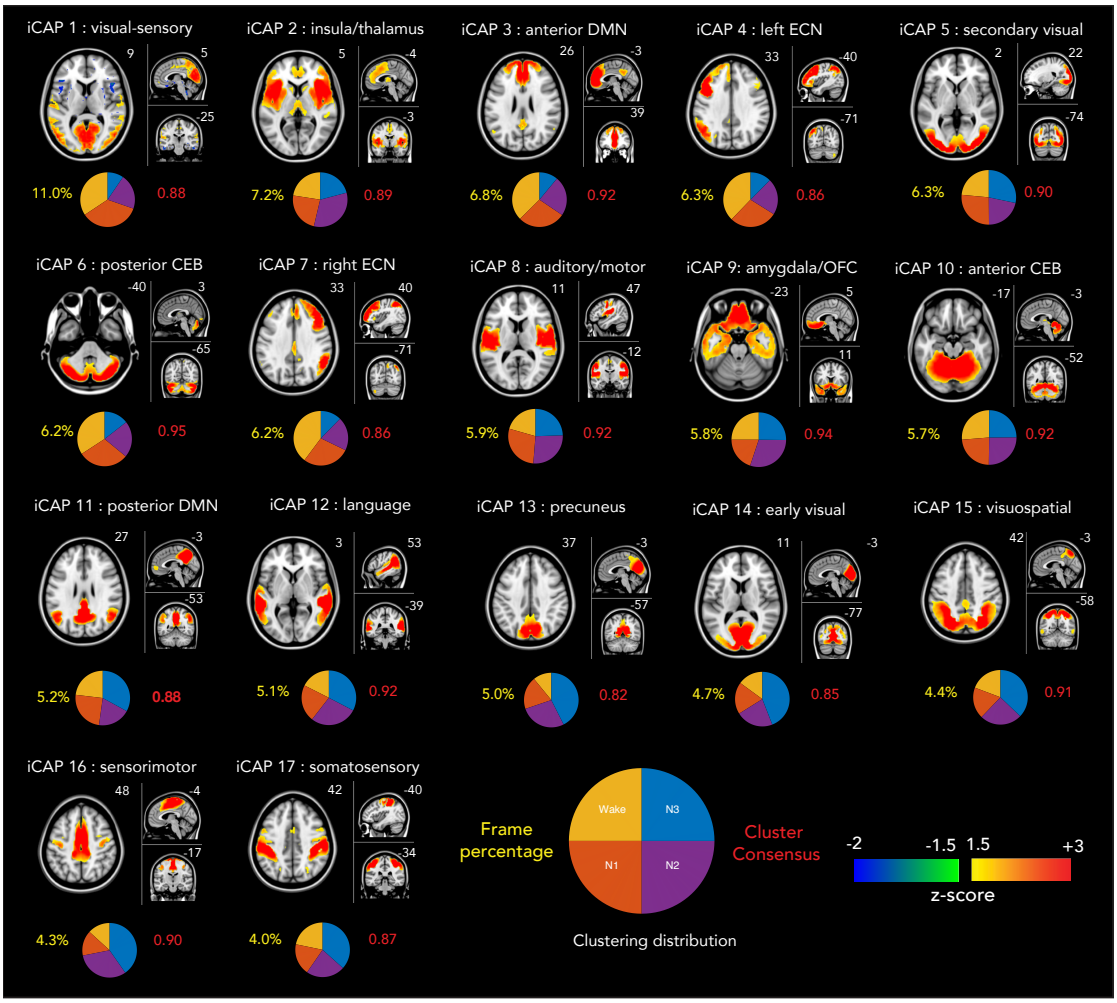

Figure 3.2 – **Spatial patterns of the 17 innovation-driven co-activation patterns (iCAPs)** derived from all recordings across both experiments. The iCAPs are numbered according to the total number of significant transients found (descending order), which are shown below the functional maps in yellow font. The cluster consensus of each iCAP is written in red. Pie charts indicate the distribution of each iCAP across sleep stage. MNI coordinates of each brain slice are indicated in white font. DMN - default mode network, ECN – executive control network, CEB - cerebellum, OFC - orbitofrontal.

together and underwent a clustering procedure known as the iCAPs framework to obtain the most prevalent brain spatial patterns. We observed 17 large-scale brain networks displayed in Figure 3.2, representing the different functional maps that dominate brain activity from wakefulness to deep sleep. The iCAPs are ordered in descending order according to the number of times that they appeared in the significant transient frames. We then looked at the wake/sleep stages where these significant transient points occurred, and identified which iCAP they corresponded to. Figure 3.2 also shows the clustering distribution for each iCAP in pie charts, revealing the proportion of transient frames that the clustered iCAP occurred in each wake/sleep stage. Spatial similarities between iCAPs generated separately using Study 1

and Study 2 are shown in Supplementary Figure A.4.

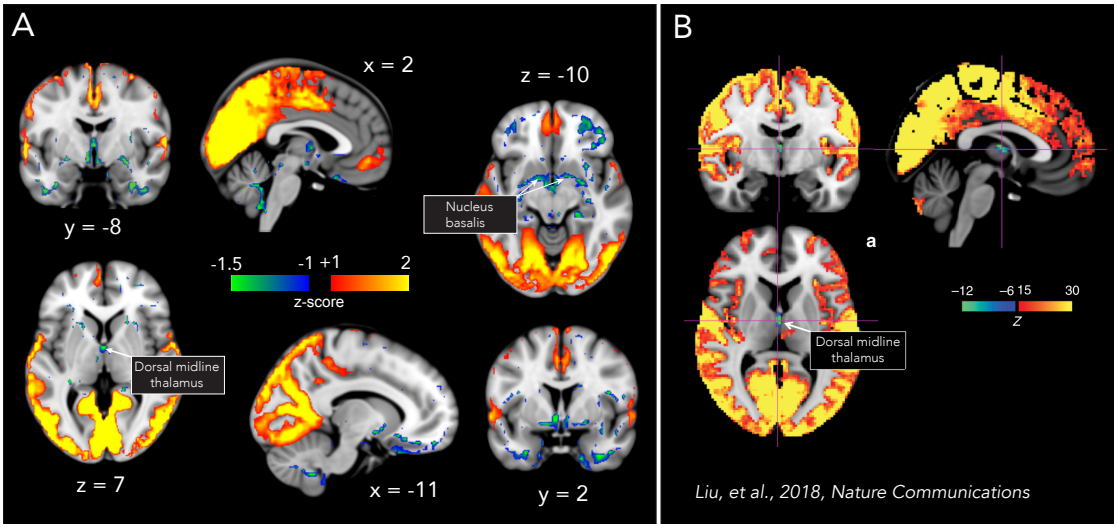

Figure 3.3 – **Visual-sensory iCAP reveals deactivation in subcortical regions.** (A) Arousal-related network showing the deactivation in dorsal midline thalamus (figure adapted from previous study ((Liu et al. 2018). (B) Visual-sensory iCAP also shows negative patterns in subcortical regions. (C) Lowering the threshold (z < -0.2) of the z-scored spatial map reveals a general deactivation in the subcortical region (orbitofrontal, and amygdala), which resembles almost similar but oppositely signed spatial pattern of iCAP 9 (amygdala/OFC). Weak and negative activation in primary sensory cortices are also observed in iCAP 9.

The first iCAP (iCAP 1) featured both visual and somatosensory regions, and resembles a previously observed functional pattern that distinguishes sleep from waking conditions. These regions were found to be more prevalent in N1 and N2 sleep here, like in previous work [Tagliazucchi et al., 2013, Tagliazucchi and Laufs, 2014]. Unlike any other iCAP extracted, this iCAP also reveals a negative activation in subcortical regions, very much similar to what was observed by Liu and colleagues (Liu et al. 2018), as is shown side-by-side in Figure 3.3.

ICAP 2 featured regions of the salience network and presented a strong and selective activation of the insula, as well as portions of the thalamus. This network showed the highest likelihood to occur during N2 sleep. ICAP 3 predominantly occurred during wakefulness, displaying the anterior portion of the default mode network (DMN), and in particular the anterior cingulate. By contrast, iCAP 11 (posterior DMN) and iCAP 13 (precuneus, ventral DMN) were predominantly active during deep sleep (N3), and displayed stronger activation in the posterior regions (e.g., PCC, precuneus, IPC). Notice that the anterior and posterior DMN were separately extracted during the clustering procedure despite their spatial similarity, as shown in Figure 3.4. This reflects a strong dissociation between the DMN sub-networks in terms of their temporal dynamics across the different sleep stages.

Meanwhile, attention-related networks such as the iCAPs 4 and 7 (left and right executive control (ECN) networks), and iCAP 15 (visuospatial) were also observed. Interestingly, both

ECN networks mostly resulted from transients occurring during wakefulness, whereas the visuospatial network predominated during N3 sleep. ICAP 9 contained the amygdala and the orbitofrontal cortex, with limbic-emotional iCAPs predominating during N2. We also found various networks corresponding to sensory areas corresponding to transients from the deep sleep (N3). These included iCAP 8 (auditory/motor), iCAP 14 (early visual), iCAP 16 (sensorimotor), and iCAP 17 (somatosensory). Finally, similar to the DMN, we also observed a

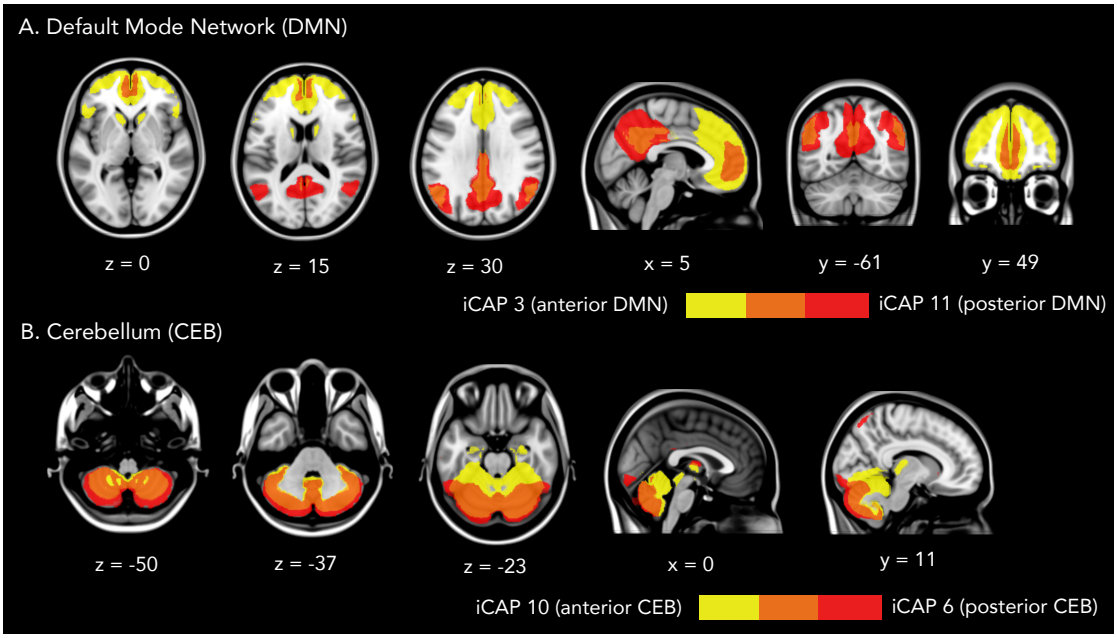

Figure 3.4 – **Dissociation of DMN and cerebellum into posterior and anterior parts.** (A) Overlay of iCAPs corresponding to anterior DMN (yellow) and posterior DMN (red). (B) Overlay of anterior cerebellum (yellow) and posterior cerebellum (red). We also see some co-activation in the thalamus in the anterior cerebellum.

dissociation of the cerebellum in the form of iCAPs 6 and 10, both featuring the anterior and posterior regions of the cerebellum, respectively. These iCAPs are overlaid in Figure 3.4B. A detailed description of the regions in each iCAP (AAL regions [Tzourio-Mazoyer et al., 2002]) and their similarity to Greicius networks [Shirer et al., 2012] are displayed in the Table A.1.

iCAP relative cumulated durations are consistent with the cluster distribution of significant transients, and both reveal stage-dependent network activity

The iCAPs were obtained by clustering the significant transients from the data (i.e., instances when functional maps markedly increased or decreased activity). The pie charts in Fig 3.2 display how the significant transients contributing to each iCAP were distributed across wake/sleep stages. In order to obtain the activity time-courses of iCAPs per participant in fMRI TR resolution (i.e., frame-wise), and be able to compute their durations (in contrast to only transients or activity changes), we performed a spatio-temporal transient-informed

back-projection of the iCAPs onto the HRF-deconvolved frames.

In Figure 3.5A, we display the RCD or the likelihood of each iCAP to appear in a specific sleep stage. This is a measure describing the cumulated time that a particular iCAP was active divided by the total time that the participant spent in a particular sleep stage. This normalization thus ensured that the reported persistence of an iCAP within one sleep stage was independent of the duration of that stage. Unsurprisingly, the relative cumulated iCAP durations appeared mostly consistent with the clustering distribution observed in Figure 3.2. Attention-related iCAPs such as the left and right ECN showed predominant sustained activity in wakefulness. Conversely, sensory-related iCAPs (e.g., early visual, somatosensory, sensorimotor) were more persistent in N2 and N3 sleep compared to wakefulness and N1. The insula/thalamus iCAP also displayed predominant activity in N2 sleep. In addition, we found that the relative cumulated durations of iCAPs related to anterior DMN is higher during wakefulness, while the posterior DMN iCAP showed higher likelihood to appear in N3. These findings are well in line with previous observations of DMN dissociating into posterior and anterior parts upon reaching deep sleep [Larson-Prior et al., 2009, Sämann et al., 2011]. Interestingly, we also observed the posterior (but not the anterior) cerebellum to be preferentially activated during wakefulness compared to N3. The corresponding test statistics (e.g., p-values, t-statistic, and effect sizes) for individual networks are reported in Table A.2.

We then looked at the general trend of network activity across the different sleep stages (inset of Figure 3.5A), and found that the RCD of iCAPs significantly increased during N2 with respect to wakefulness and N1, followed by a steep decrease in N3. Please note that values above 100% are due to the temporal overlapping nature of iCAPs, with two or more iCAPs occurring at the same time. We then computed a network-based normalization of RCD, displayed in Figure 3.5B. This is equivalent to normalizing the number of active time-points in each sleep stage by successive division to two factors: (1) overall number of time instances that an iCAP is present over the whole duration of the recording and (2) the total time that the participant spent in a particular sleep stage. The bar plots in Figure 5B show that almost all iCAPs displayed proportions above 25% in N2 sleep (broken horizontal line). We observed a marked disparity in iCAP proportions during wakefulness and N3, in contrast to the more uniform distribution in N1 and N2 sleep. Moreover, iCAPs that were more represented in wakefulness were less represented in N3, which is consistent with the RCD measures of each iCAP across wake/sleep stages in Figure 3.5A (all $p<0.05$). Specifically, iCAPs that reached more than 25% activity during specific sleep/wake stages are: (1) wakefulness – visual-sensory, left and right ECN and posterior CEB, (2) N1 – posterior CEB, visual-sensory, secondary visual, and (3) N3 – posterior DMN, somatosensory, precuneus, early visual, sensorimotor.

Contrasting with the RCD above, the average duration or mean lifetimes of an iCAP represents the average bouts (in seconds) of continuous activity. Overall, we observed iCAPs to be active between 4.5 to 10.5 seconds (7.3 ± 1.7 seconds; Figure 3.5C). In general, we also found iCAP activity to have longer durations in N2 compared to N3 ($p < 0.01$), and compared to wakefulness ($p < 0.01$). For some of the networks, the average durations were proportional to

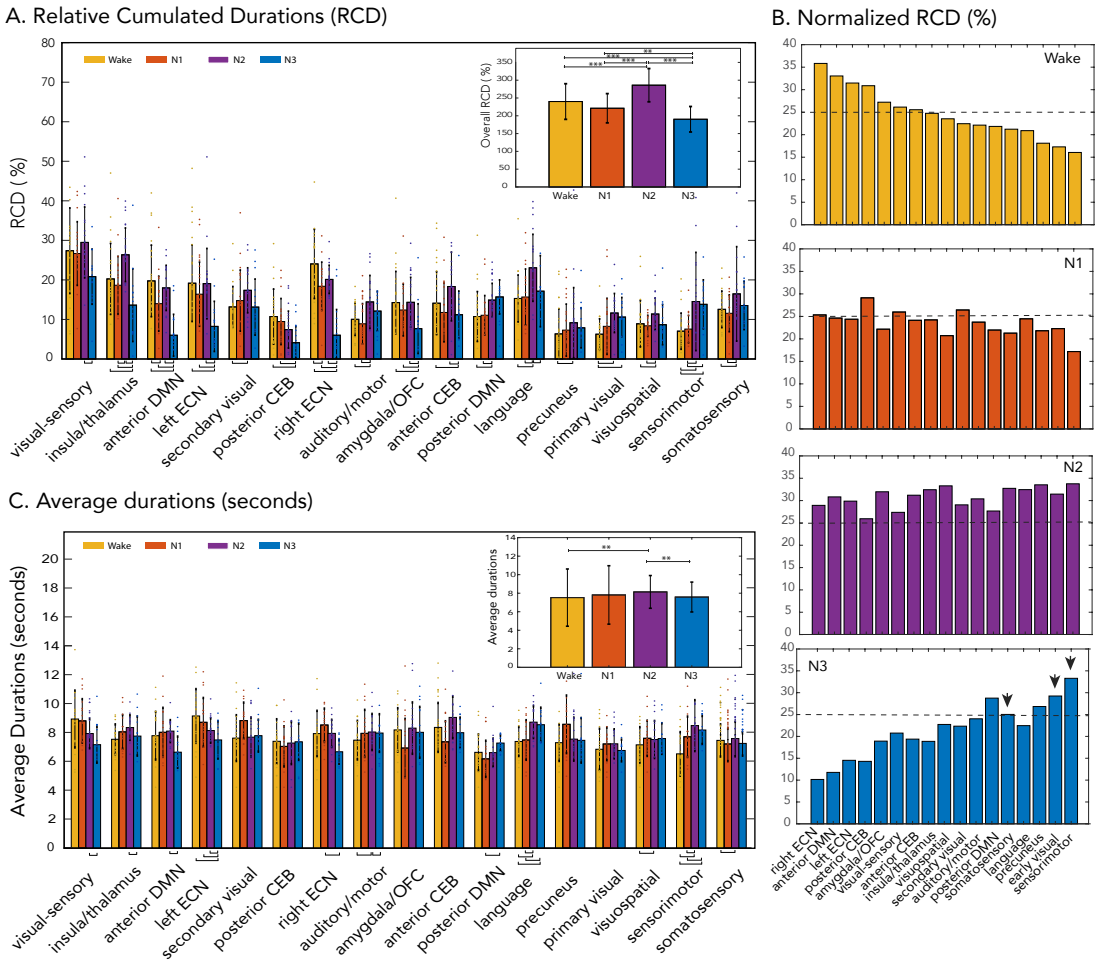

Figure 3.5 – **Duration measures of iCAPs in different sleep stages** (A) Total durations of iCAPs (in percent) showing each network's likelihood to occur in different sleep stages. The inset in shows the overall trend of the iCAP total duration in wakefulness/sleep stages, which are all above 100%, reflecting the tendency of iCAPs to overlap in time. (B) Number of time-points that an iCAP is active per sleep stage divided by the number of timepoints that an iCAP is active in the whole time-course. The normalized total durations in wake, N1, N2, and N3 for each iCAP is always equal to 100%. The broken horizontal line in each subplot indicates 25% likelihood, while the arrows in N3 subplot indicate visual and sensory-related iCAPs that are above the 25% line. (C) Average durations (in seconds) of iCAPs reflecting the length of continuous activity of iCAPs. The inset in (C) correspond to the overall trend of iCAP average durations. The height of the bars shows the mean and error bars indicate their corresponding standard deviations. The horizontal lines at the bottom of the bar plot correspond to significant differences evaluated through paired t tests and permutation testing. Horizontal lines with 3 stars and 2 stars in the insets represent significant differences with p-values less than 0.001 and 0.01, respectively.

their relative propensity to occur in each sleep stage. For instance, the left ECN displayed the least likelihood to occur in N3, and also exhibited the shortest average duration during that

stage.

Alterations in network co-occurrences in different sleep stages

To go beyond the iCAPs' individual temporal properties, we looked at the probability of having either none, one, two, or more iCAPs occurring at one time-point. Figure 3.6A shows the likelihood for different numbers of overlapping iCAPs to occur in wakefulness and in different sleep stages. We observed the overlaps to be maximally N = 10, and thus we limit the histogram for the range N = 0,1,2, . . . , 10. Among all vigilance states, N2 displayed a relatively flatter distribution especially compared to wakefulness and N3 (both $p < 0.05$), with a higher likelihood of having 2 to 5 overlapping iCAPs. For a table of test-statistics, see Supplementary Table A.3. Given that individual iCAPs had higher likelihood to occur in N2 compared to other sleep stages, it is not surprising that iCAPs are also more likely to overlap in time during N2.

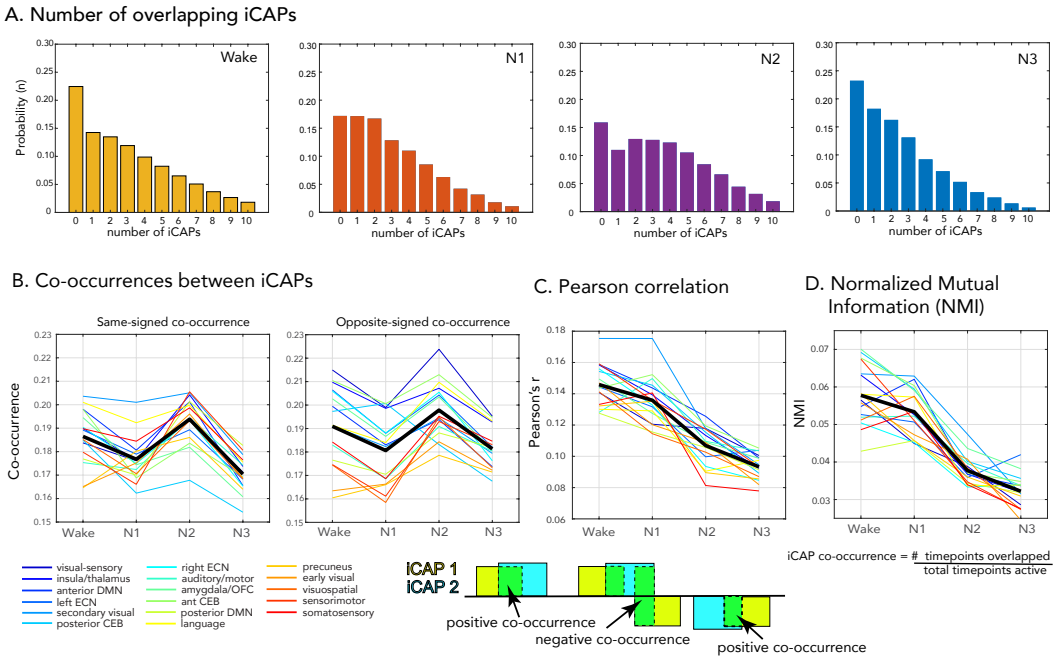

Figure 3.6 – **Network interactions across different sleep depth.** (A) The probability of different numbers of iCAPs to overlap is displayed for wakefulness and different sleep stages. (B) The iCAP co-occurrence pertains to the number of time-points during which a pair of iCAPs were both active, divided by the total number of time-points that at least one of them was active. This was computed separately for similar signed activations (same-signed co-occurrences) or opposite signed activations (opposite-signed co-occurrences). Co-occurrence was computed for each network (i.e., pair-wise co-occurrence of one network with all the other networks). (C) Mean of classical FC metric (Pearson's r) applied to iCAP time-courses. (D) Mean of normalized mutual information (NMI) of pairwise iCAP time-courses. The computation for the co-occurrence between pairs of iCAP is illustrated on the bottom of the figure.

We then evaluated the pair-wise co-occurrence between each iCAP and all other networks

(Figure 3.6B). The co-occurrence is a normalized measure of the number of time-points at which an overlap occurs between two pairs of iCAPs, divided by the total number of time-points that at least either one is active. This is different from the temporal overlap measure observed in Figure 6A because we focus on pair-wise co-occurrence between two iCAPs. This can also be referred to as the Jaccard similarity score of two iCAP time-courses. We also take into account the signs of the activations, see illustrations in Fig 6. In general, all iCAPs showed a high mean same-signed co-occurrence and opposite-signed co-occurrence in wakefulness and N2, and a decrease in N1 and N3 ($p<0.001$, see Supplementary Table A.4). Next, we compared these observations using classical FC analysis (i.e., Pearson correlation) and normalized mutual information (NMI) applied to the iCAP time-courses. We observed that the pairwise network FC across iCAPs decreased with increasing sleep depth (with no increase in N2). The same trend can be observed in Figure 3.6D which shows the overall NMI of iCAP time-courses, suggesting that the mutual dependence between iCAPs decreases from wakefulness to deep sleep. These observations are consistent with the general notion that connectivity breaks down with increasing sleep depth [Tagliazucchi and Laufs, 2014]. Both FC and the NMI also has the lowest variance (Figure A.5) during N2 and N3 sleep.

Robustness of results across datasets and across choice of parameters

In order to see whether the resulting iCAPs are robust with regards to the dataset from which the contributing frames belong to, we perform an averaging of significant transient frames corresponding to iCAP indices within a similar dataset (i.e., among those belonging Study 1 and among those belonging to Study 2). Study 1 contributes 58373 frames while the Study 2 contributes 39235, giving a total of 97608 significant frames. The number of contributing frames coming from each dataset is shown in Figure A.4 and is written in red and yellow fonts for Study 1 and Study 2, respectively. The distribution of clustering assignments is particularly telling, especially on the contribution of each dataset reflecting the tendency of the iCAP to appear in either Study 1 (reaching deep sleep) or in Study 2 (wake to N1). We observe the visual-sensory, anterior DMN, posterior cerebellum, and the left and right executive control networks to be mostly contributed by frames belonging to the Study 2 dataset. On the other hand, the Study 1 contributed more on the recovery of the remaining networks. We computed the spatial similarity of the resulting iCAPs. Figure A.4 provides an overlay of the iCAPs computed using frames belonging to Study 1 versus Study 2. Visually, all iCAPs display similar spatial patterns in both the Study 1 and Study 2, which is also well supported by their Dice coefficients.

Moreover, we also evaluated the robustness of the results in terms of the observed temporal characteristics using the data from Study 2. We computed the occurrences and co-occurrences of iCAPs using only subjects from Study 2. These are displayed in Figure A.6 and Figure A.7, respectively. In these figures, there are no values for N3 since subjects from Study 2 only reach NREM stage 2. Consistent with the findings in Study 1, we find a surge in iCAP activity during NREM stage 2. We also observe a higher likelihood for 3 or more iCAPs to overlap in this stage of sleep. Meanwhile, the iCAPs overlap the least during wakefulness.

Finally, we evaluated the robustness of the results with respect to the choice of the number of clusters, K. The choice of the number of clusters was motivated by a quantitative approach for optimal class discovery (Monti et al., 2003). To assess whether the choice of K affects the final outcome of the analysis, we also repeated the clustering procedure to the whole dataset using K = 20. Figure A.8 displays the side-by-side analysis done using K = 17 and K = 20. Here, we find that the general spatial characteristics are the same for both analyses, and that the distribution (pie charts) show similar proportions across NREM sleep stages. Moreover, we also find that increasing the number of K, beyond the optimally observed K = 17, results in the recovery of similar repeating spatial patterns that are already present for K = 17.

3.2.4 Discussion

General findings

Here we extracted large-scale brain networks or iCAPs by clustering moments of significantly changing brain activity from two studies using simultaneous EEG and fMRI recordings during sleep. The use of transient fMRI activity allowed us to obtain temporally overlapping spatial patterns together with subject-specific time-courses at a time-scale of seconds. Altogether, our findings help establish the following: (1) that coordinated activation of regional brain areas during wakefulness was largely preserved in N1, N2, and N3 sleep, yet with specific relative distributions and dynamic modulations across NREM stages, (2) that functional brain networks sustain continuous bouts of activity on an average of about 7.3 ± 1.7 seconds, and for some networks, their mean lifetimes changed accordingly to their likelihood to occur in each sleep stages, (3) that these networks exhibit maximal activity in N2 sleep and the least in N3, and (4) that the pair-wise network co-occurrences generally show marked increase in N2 sleep followed by a decrease in N3, even though (5) FC and mutual information between these networks progressively broke down from wakefulness to N3.

The first observation is largely consistent with previous work reporting the persistence of large-scale brain networks during sleep [Horovitz et al., 2009, Sämann et al., 2011] with various modifications when reaching deeper sleep stages, such as decreased connection strengths in RSNs [Larson-Prior et al., 2009], changes of hierarchical organization into smaller independent modules [Boly et al., 2008], or reduced long-range temporal dependencies in the BOLD signal [Tagliazucchi et al., 2013]. For instance, the well-known attention-related networks, such as the bilateral ECN, occurred preferentially during wakefulness and N1, whereas networks associated with primary sensory systems were most present during N2 and N3, consistent with some earlier observations [Horovitz et al., 2009, Larson-Prior et al., 2009].

In the next subsections, we discuss the most important findings of the work and give interpretations to the observed dynamical characteristics of cross-network interactions.

Visual-sensory iCAP confirms the emergence of arousal-related network

The first iCAP (iCAP 1, Figure 3.2), corresponding to the network of regions contributing to the highest number of significant transients in our dataset, was the visual-sensory iCAP, which predominated during the transition to sleep (N1). It included a large portion of the visual cortex, as well as some sensory and motor regions. This network also displayed a unique characteristic of having negative activations in subcortical areas. Using combined electrophysiological and fMRI signals, Liu and colleagues (2018) described a very similar spatial pattern, with negative activation in basal forebrain and thalamus paralleled by positive activations in sensory cortices, which arose during momentary drops in arousal, indexed by a spectral shift in local field potentials toward low frequencies. In Figure 3.3, we provide a side-by-side comparison between our visual-sensory iCAP 1 and the arousal-related network observed by Liu et al. While both networks displayed positive activations in visual, sensory, and motor regions, they also exhibited negative activation in the dorsal midline thalamus, as well as the nucleus basalis. Because the thalamus holds an important role in regulating the physiological state of arousal during sleep [Gent et al., 2018, Saper et al., 2010] and is responsible for relaying motor and sensory signals to the cerebral cortex, we can suggest that reduced thalamic activity favors sleep onset while limiting incoming external sensory stimulation. This interpretation is also supported by the fact that most of the frames that contributed to this iCAP came from wake and N1, suggesting that this iCAP might play a key role in the onset of sleep. Interestingly, however, we found that this network exhibited a persistent likelihood in terms of the cumulated duration to occur in N2, before decreasing in N3. Fittingly, it has also been found in previous studies that there is an increase in BOLD signal variance in sensory and motor cortices during N1 and N2 sleep, which has been used as diagnostic for detecting sleep in typical resting-state recordings. On the other hand, high activities in frontal, parietal, and temporal cortices are typically associated with wake conditions [Tagliazucchi and Laufs, 2014, Tagliazucchi et al., 2013]. These two brain spatial patterns bear resemblance to the visual-sensory iCAP and the bilateral ECN, respectively, whose temporal properties show consistent behavioral stage-dependence. In addition, [Stevner et al., 2019] reported the same visual-sensory network to be implicated also in N1.

DMN and cerebellum dissociates into posterior and anterior regions

The DMN has been well studied particularly in light of its connectivity changes when subjects transition from wake to light sleep, and deep sleep. The initial hypothesis was that the DMN supports a range of self-related mental processes, such as unconstrained self-referential thought and recollection. However, previous studies have found the DMN to persist during light sleep, as well as in deep sleep with reduced connectivity between the medial prefrontal cortex (MPFC) and the rest of the network [Horovitz et al., 2009, Larson-Prior et al., 2009]. In the present work, we observed a dissociation of the DMN sub-components into its posterior and anterior parts (Figure 3.4A). The pie-chart distributions of cluster frame indices displayed in Figure 3.2 and the cumulated durations in Figure 3.5A reveal that the anterior region of the

DMN mostly occurred during wake and N1 sleep whereas the posterior DMN predominated in N2 and N3 sleep. We found a similar antero-posterior dissociation in the cerebellum (overlap shown in Figure 3.4B). Studies that investigated FC in the posterior and the anterior regions of the cerebellum have remained inconclusive regarding their potential roles in sleep. The cerebellum has been found to be related to motor control and motor memory formation [Gao et al., 2012]. It has also been observed to show sleep stage-dependent activity, whose impairments disrupt the sleep-wake cycle which then leads to sleep disorders [DelRosso and Hoque, 2014]. As a general observation, signals in the cerebellum appear to be lower during N1 as compared to wakefulness [Hiroki et al., 2005, Kaufmann et al., 2006]. However, a map summarizing major cerebellar research in sleep using fMRI and PET revealed a localized change in cerebellar activity during SWS surrounding the cerebellum's larger lobules (parts IV, V, VI, and VII), as well as marked correlations with slow spindles in N2 [Canto et al., 2017]. These lobules make up the anterior region of the cerebellum [Dang-vu et al., 2010]. Moreover, it has been previously shown that the sensorimotor domain is related to the anterior cerebellar lobe, while the cognitive domain corresponds to the posterior cerebellar region [Stoodley and Schmahmann, 2010]. This particular finding is consistent with the observed co-activation of the anterior cerebellum with regions corresponding to the sensorimotor network represented by iCAP 10 (see also [Shirer et al., 2012]) and Supplementary Table A.1) Meanwhile, the inferior posterior lobe has been associated with the frontal gray matter, a well-known core area of executive function [Jung et al., 2019, Tiemeier et al., 2010]. Altogether, these studies support our current findings on the preferential tendency of posterior and anterior cerebellum to persist in wakefulness and N2, respectively.

Alterations in network temporal characteristics and interactions during sleep reveal paradoxical trend

Previous connectivity analyses of brain regions found a general decrease of FC with increasing sleep depth [Haimovici et al., 2017, Spoormaker et al., 2010, Tagliazucchi et al., 2016]. Classical FC approaches capture statistical inter-dependencies between activity from two brain regions [Friston, 2011]. Typically, if two regional time-courses exhibit simultaneous positive (or negative) values over some time-windows, and exactly opposite values on some other time-windows, the final correlation (and derived FC) between both regions would be close to zero, despite the synchronized deviations from the baseline, albeit the incoherent signs. Therefore, it follows that a decrease in FC does not mean a decrease in global network activity. This is also true in measures of mutual information.

In this work, we demonstrated that network activations increased when participants went from wakefulness to N2, while they decreased when reaching N3. This important observation was captured because we were able to extract network temporal activity at the frame-wise level for each individual iCAP. In contrast, classical dynamic FC analyses only capture the statistical relationship among ROIs comprising the networks, and not the actual activity of the networks themselves. Complementary, NMI measures the general relationship and detects both linear

and non-linear mutual dependencies between these networks. Furthermore, because the iCAP approach allows networks to overlap in time, we were able to observe that functional networks show a higher likelihood to overlap (e.g., 2 to 4 iCAPs active at the same time) in N2 compared to other stages. Congruently, we also found high co-occurrences in N2. By contrast, N1 and N3 generally showed much lower co-occurrences. What is remarkable is the simultaneous increase in both the same-signed and opposite-signed activations in N2 (Figure 3.6B). Meanwhile, classical Pearson correlation measure applied to iCAP time-courses displayed a decreasing trend from wakefulness to deep sleep, consistent with FC findings based on regional fMRI time-courses [Haimovici et al., 2017, Kung et al., 2019, Spoormaker et al., 2010, Tagliazucchi et al., 2016]. This loss of FC is supported by the observed decline of iCAP mutual dependence based on NMI measures in Figure 3.6D further strengthening the observed paradox between network activity and network interactions from wakefulness to deep sleep.

The remarkable increase in network activity during N2 is particularly telling and perhaps not coincidental, as it is to be noted that N2 sleep acts as an intermediate epoch between drowsiness and deep sleep. This is also the period when various paroxysmal events occur, i.e., spindles and K-complexes [Jahnke et al., 2012]. Functionally, the role of K-complexes is under debate whether it acts to promote deeper sleep, or whether it does the opposite and correlates with cortical arousal [Bastien et al., 2000]. As such, it is believed to often emerge at times of brain instability. Thus, it cannot be discounted that the general increase of network activity and co-occurrences in N2 may be due to a common cause or a global drive that occurs at this sleep stage. Meanwhile, the overall decrease in network activity and co-occurrences during N3 fit previous findings on global increase in functional segregation during deep sleep [Boly et al., 2012, Horovitz et al., 2009, Spoormaker et al., 2012]. Deep sleep is first and foremost characterized by high-amplitude, low-frequency brain waves. It is manifested by unresponsiveness, and the decline of the ability to react to external stimuli [Cirelli and Tononi, 2008]. Fittingly, similar findings have been shown for other unconscious brain states, such as in propofol anesthesia [Monti et al., 2013].

We therefore summarize three important findings regarding network temporal characteristics of sleep stage N2: (1) FC and NMI between networks decreaseds with sleep depth, (2) there is a higher likelihood for stereotyped networks or iCAPs to emerge in N2 compared to wakefulness and other sleep stages, and (3) both same-signed and opposite-signed co-occurrences between iCAPs increased in N2. In line with previous works, we speculate that the decreased in FC and more strongly the reduced NMI, both reflect a global reduced efficiency of information transfer between networks. These observations are possible manifestations of reduced brain network integrity [Nofzinger et al., 2013] that is commonly interpreted to reflect reduced consciousness during sleep [Cirelli and Tononi, 2008]. This interpretation is consistent with the general consensus among previous studies that view consciousness not as a persistence of functional brain networks to occur, but rather the degree of interactions among them [Horovitz et al., 2009, Nofzinger et al., 2013, Sämann et al., 2011, Spoormaker et al., 2010]. Intriguingly, despite the decrease in mutual dependence between networks in N2, a general increase in

network activity is observed, possibly reflecting the unstable dynamics among iCAPs. In other words, observation (1) likely corresponds to a decrease in global network integration, while observations (2) and (3) can be presumed to express network attempts to communicate specifically during N2 resulting in an increased inter-network co-occurrence with nonetheless unstable synchronizations. This latter finding is in agreement with a very recent observation by [Kung et al., 2019], who reported a very high variance in the dynamic FC evaluated across sliding windows during N2, compared to during wakefulness, N1, and N3. High dynamic FC variance in N2 was interpreted in their study as elevated instability of information transfer between networks, whereas low mean dFC was deduced to reflect decreased intra-network consistency.

Concerning N3, we observed (1) an overall decrease of global network occurrence, (2) a decrease in network mean lifetimes, and (3) a decrease in functional association characterized by lowest cross-network FC and mutual dependence through NMI. Altogether, our results are consistent with the interpretation of a more stable brain state [Jobst et al., 2017] yet more localized signal integration in SWS [Deco et al., 2017]. Spatially, the huge drop in anterior DMN's likelihood to occur in deep sleep compared to other stages, together with the high persistence of unimodal primary sensory networks, provide further support to the notion of a break-down of long-distance functional connections [Tagliazucchi et al., 2013] in favor of short-distance associations [Boly et al., 2012]. The global decrease of network activity, mean lifetimes in N3, and mutual dependence between networks are, altogether, potential indicators of a more stable state of the system, but also a general loss of effective information integration.

Methodological Aspects and Limitations

The iCAPs framework has already been applied to resting-state fMRI data of healthy [Karahanoglu and Van De Ville, 2015], clinical populations [Zöller et al., 2019], and spinal cord fMRI [Kinany et al., 2020]. The recovered iCAPs in the present study were highly similar to the ones observed in the first two studies, while additional spatial patterns unique in the sleep were also found (e.g., visual-sensory network with deactivations in subcortical regions, cerebellum, and insula). While dynamic FC analyses relying on direct statistical interdependencies of brain regions are effective in capturing coherent fluctuations between brain areas, it does not capture individual brain activity at a single time-point. The framework is unique in its ability to detect transients which allows for the extraction of spatially and temporally overlapping functional networks at each time-point, which is a particular advantage that is beyond the classical methodologies applied in sleep studies. Using this advantage, we were able to show, for the first time, individual network activity at precise fMRI temporal resolution across sleep stages and found a distinct stage-dependent activity for each network.

Although the method used is advantageous particularly in our goal to characterize functional brain dynamics during sleep, it also comes with some limitations. Unlike classical FC studies where one can evaluate the integrity of connectivity within a network (i.e., correlate one ROI to another ROI within the same network), the iCAP approach extracts the spatial patterns

as a whole. Thus, we are not able to evaluate the strength of connectivity within a network. Nevertheless, the clustering step itself reveals data-driven dissociations of well-known spatial patterns, such as the DMN and the cerebellum, which essentially reflects the spatial modifications that the brain undergoes as it transitions to SWS. Furthermore, while our approach does remove hemodynamic effects resulting to a much cleaner representation of the BOLD activity, the model used to deconvolve the HRF is assumed to be the same across all brain voxels, following the general practice in the fMRI field. Nevertheless, for the research questions that we aim to answer in this work, these effects can be reasonably ignored for our main findings.

Lastly, while the extracted temporal overlaps and co-occurrences between iCAPs give useful insights on the overall global brain dynamics, these measures do not reflect causal influence between these networks and therefore cannot be used to directly assess different levels of consciousness. In contrast, FC, and more strongly, NMI, capture linear and non-linear interdependencies between iCAPs. The agreeable outcome between these two measures provides further evidence on reduced information integration that accompany increasing sleep depth, and give stronger confidence on the observed paradox between the activity of large-scale brain networks and their mutual dependencies. Nevertheless, in order to more accurately quantify the level of consciousness across NREM sleep, additional analysis is suggested using previously prescribed measures of integrated information [Krohn and Ostwald, 2017, Tononi, 2008, Tononi et al., 2016].

In summary, we investigated the dynamic properties of large-scale functional networks during wakefulness and NREM sleep. We used the TA and iCAPs framework to capture precise moments of transient brain activity, giving us a quantitative view of how the brain dynamically evolves across the different stages of NREM sleep. We found new networks that are largely related to regions that support the physiological organization of sleep and arousal. We also uncovered whole brain spatial patterns that resemble currently known RSNs whose temporal profiles are consistent with previous findings. The temporal dynamics these networks exhibited alterations between different sleep stages. In particular, we observed a global dissociation/decrease of brain activity both in spatial and temporal domains, from wakefulness to deep sleep. Unexpectedly, we found an increase in network activity in N2 and a global increase in simultaneous positive and negative network co-occurrence, signaling instability of network synchronization and ineffective brain integration. Altogether, these findings support the general consensus of relating cortical integration to consciousness dissipation during sleep and provide new evidence for the presence of unstable yet distributed global inter-regional co-activation in N2 sleep.

3.2.5 Additional results

In the journal article presented in the previous section, we limited our analysis to characterizing brain dynamics across the different NREM sleep stages. In this section, we assess how the dynamics of brain function can be related to the underlying anatomy as the brain transitions

from wakefulness to deep sleep. Using the same set of subjects used in the previous study, we estimate brain dynamics using classical sliding-window analysis (*i.e.,* Pearson correlation) and relate each window to the underlying structure by directly correlating it with SC, following the method employed by [Liégeois et al., 2016].

Methods

Estimating brain dynamics using classical seed-correlation. We used AAL Anatomical atlas to parcellate the functional data into 90 regions. For the functional data, we calculated the mean signal within each region of the AAL, forming a $90 \times 90 \times N_{TC}$ datamatrix, where N_{TC} is the length of the whole fMRI recording of each subject. Because subjects were allowed to sleep for as long as they manage, N_{TC} is different for each subject. The mean BOLD signal in each region is correlated to all other regions to obtain a static connectivity matrix, FC_{static}. Next, using a temporal window of length $TR = 30$ (equivalent to 1 minute) and a step of 1TR, we performed a sliding-window correlation in the parcellated functional data. For each temporal window, we normalized the resulting FC with FC_{static}. Finally, we vectorized the FC measure in each temporal window and computed the mean of all pair-wise inter-regional correlations.

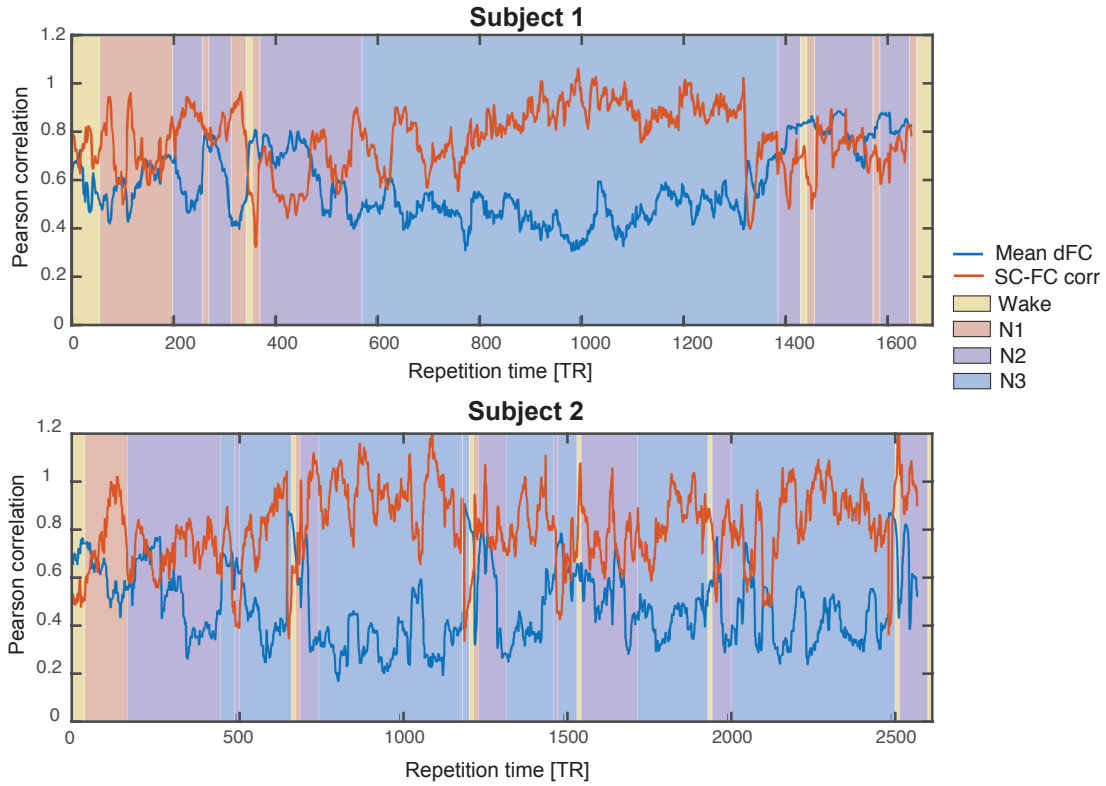

Figure 3.7 – **Example dynamic FC evolution of two subjects, and their relation to SC.** As the mean FC breaks down during NREM stage 3, the coupling between FC and SC increases, as can be observed in blue regions. Transitions between NREM sleep stages are also clearly reflected in the fluctuations of dynamic FC and SC-FC coupling.

Probing the link between brain structure and function. We used the conventional approach of extracting SC matrix from diffusion data, as discussed in Chapter 2. We used a total of 100 subjects, all of which were downloaded from the Human Connectome Project (HCP) database, WU-Minn Consortium. These data were downloaded in a preprocessed format. The details of the preprocessing are well-documented elsewhere [Glasser et al., 2013]. All volumes are then processed using MRTrix3[3] with the following operations: multi-shell multi-tissue response function estimation, constrained spherical deconvolution, and tractogram generation. We used the AAL atlas to parcellate the cortex and estimate the corresponding SC matrix. Each pair-wise fiber counts were normalized with the average of the volumes of the two regions being considered. The final SC matrix is obtained by averaging the SC of all 100 subjects. Finally, we vectorized this mean SC by taking only the upper triangular section of the connectivity matrix.

To link SC and FC, we performed a sliding-window correlation between the vectorized SC and FC matrices, using the same parameters used in approximating the dynamic FC. For each temporal window of length 30TR, we directly computed the Pearson correlation of FC to SC and normalized them with the correlation between FC_{static} and SC. Mathematically, this is given by the following formula:

$$R(t) = \frac{corr(FC(t), SC)}{corr(FC_{static}, SC)}. \tag{3.1}$$

The SC-FC coupling can have a value greater than 1.

Results and Discussion

FC breaks down as the brain transitions to deep sleep. Figure 3.7 shows (in blue) examples of averaged dynamic FC evolution from two representative subjects as they transition from wakefulness to deep sleep. From the two plots, we can clearly observe a huge drop in the overall dynamic FC during NREM stage 3, in agreement with what was observed in previous studies [Spoormaker et al., 2010, Kung et al., 2019]. This further supports our findings using the iCAPs approach – that there is a significant reduction of the overall occurrence of functional networks during the deep sleep. This is more straightforward to observe in Figure 3.8(A) which shows a decreasing mean dynamic FC across NREM sleep, averaged from all subjects.

FC increases coupling with SC during deep sleep. In contrast to the observed decrease in mean FC during the deep sleep, Figure 3.7 instead shows (in red) an increase in the SC-FC correlation, demonstrating the tendency of brain function to increase coupling with the underlying anatomy during deep sleep. This observation is also generalized across all subjects in Figure 3.8(B) which shows significant increase in SC-FC correlation upon reaching NREM stage 3.

The observation that the anatomical-functional coupling is stronger during the deep sleep

[3] http://www.mrtrix.org/

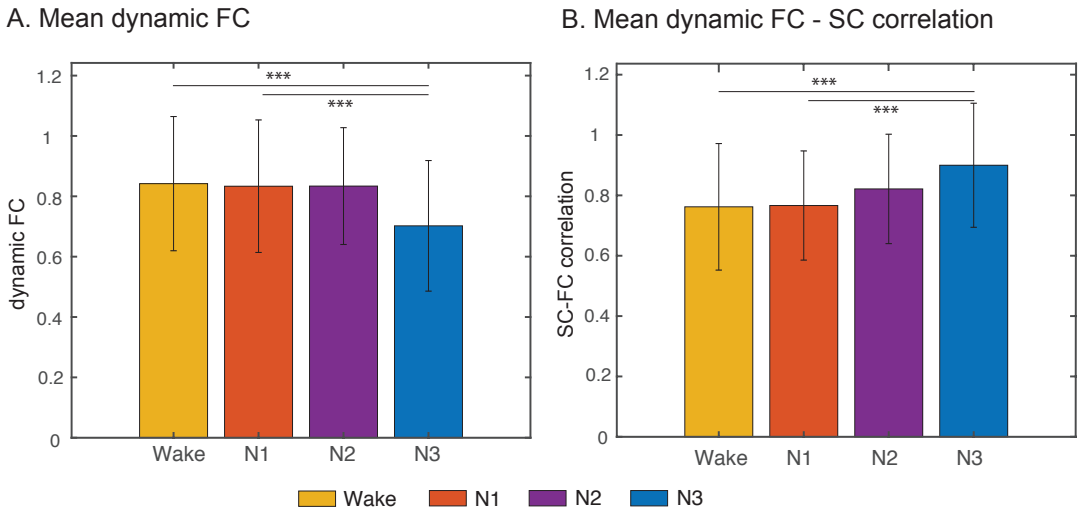

Figure 3.8 – **SC-FC relationship across NREM sleep.** (A) Average dynamic FC measure across all temporal windows corresponding to each NREM sleep stage. (B) Average FC-SC relationship computed across all temporal windows corresponding to each NREM sleep. Only temporal windows that have non-changing sleep stage within at least 30TR is considered, *i.e.,* in one sliding temporal window of 30TR, no sleep stage transition has to occur. All p-values are corrected for multiple comparison, $p < 0.001$.

is well in-line with what was previously found by Tagliazucchi and colleagues, where the similarity between the anatomical and functional connectivity networks increased during the deep sleep [Tagliazucchi et al., 2016]. In particular, they have shown that FC in primary cortices resembles that of the anatomy, in contrast to loosening structure-functional coupling between higher-order integrative cortices. Recall that in the journal article presented in the previous section, we observed that the normalized RCD during the deep sleep is dominated by the occurrence of early visual, somatosensory, and sensorimotor iCAPs. Meanwhile, at the other end, high-order association areas such as the bilateral executive control networks (ECN) and the limbic networks displayed the least RCD. This trend is of the exact opposite of that in wakefulness. We surmise that these findings reflect a previous observation that the brain retains low-order processes during deep sleep, and anchors itself on the fixed anatomy to restore its functionality upon awakening.

3.3 Journal Article: Dynamics of Functional Network Organization Revealed by Graph Laplacian Mixture Models

A. Tarun*, I. Ricchi*, H. Petric Maretic, P. Frossard, and D. Van De Ville
Preprint with preliminary results.

Author Contributions

A. Tarun conceptualized research, performed half of the analyses, supervised, and wrote the manuscript, **I. Ricchi** conducted half of the analyses and explored the optimal paramaters of the algorithm, **H. Petric Maretic** developed the method and contributed on writing part of the methods, **P. Frossard** and **D. Van De Ville** supervised and offered continuous guidance.

Abstract

Understanding the organizational principles of distributed human brain activity remains a major challenge in neuroscience. Here, we introduce a full data-driven approach based on graph learning to extract meaningful repeating network patterns from regionally-averaged time-courses. The Graph Laplacian Mixture Model (GLMM) is a generative model that treats functional data as a collection of signals that live on multiple graphs. Three types of information can be obtained using this approach: (1) meta-stable functional networks; (2) the activity profile of each network; (3) the learned graphs that describe direct interactions between regions that comprise these networks. We first validate the reliability of the proposed technique by applying it to task fMRI. The probability of each network to occur at each time-point is found to be consistent with the timing of the task paradigm, while the spatial patterns associated to each epoch of the task are in line with previously established knowledge using classical regression analyses. Second, the application to resting state leads to extracted metastates that correspond to well-known functional networks. Finally, we take advantage of the learned graph matrices to investigate the general link with anatomy. Using a dissimilarity index computed between the GLMM-based functional connectivity (FC) and structural connectivity (SC), we show that the strength of the FC-SC relationship gives rise to a behaviorally-relevant macroscale gradient that distinguishes brain regions associated to low-level and high-level cognitive functions.

3.3.1 Introduction

Functional magnetic resonance imaging (fMRI) is a unique tool to probe the functional architecture of the human brain. Specifically, spontaneous fluctuations of blood-oxygenation-level dependent (BOLD) signal have shown to be synchronised between brain regions during resting state (RS) despite the absence of task or external stimuli [Biswal et al., 1995, Smith et al., 2009]. A repertoire of functional networks have been identified in healthy [Beckmann et al., 2005, Thomas Yeo et al., 2011, Shirer et al., 2012], clinical populations [Tagliazucchi et al., 2010, Liao et al., 2014, Zöller et al., 2019], and across various brain states, for instance, different levels of sleep [Larson-Prior et al., 2009, Tarun et al., 2020d] and coma [Boly et al., 2008]. These functional networks include both sensory regions and higher-level cognitive ones, such as the default-mode network (DMN), which generally shows reduced activity during an externally-oriented task [Greicius et al., 2003], and becomes more engaged during internal mentation [Andrews-Hanna, 2012].

The interest in understanding the intrinsic functional organization of the human brain has motivated many new methods to deal with RS. Functional connectivity (FC), conventionally measured as Pearson correlation between pairs of timecourses, allows to constitute a whole-brain functional connectivity matrix [Bullmore and Sporns, 2009]. In its conventional form, FC is computed using the whole RS scan, but several dynamic extensions that acknowledge temporal fluctuations of FC have seen the light [Chang and Glover, 2010, Preti et al., 2017, Bolton et al., 2020]). One popular approach is the sliding-window technique [Kucyi and Davis, 2014, Elton and Gao, 2015, Madhyastha and Grabowski, 2014], where timecourses are segmented into temporal windows so that a time-dependent FC matrix can be obtained. Then further analysis by graph metrics [Betzel et al., 2016a], or dimensionality reduction methods such as singular value decomposition (SVD) [Leonardi et al., 2013], k-means clustering [Allen et al., 2014], or hierarchical clustering [Yang et al., 2014], can be applied.

From the analysis of windowed correlation of time-courses, one can also consider the extraction of dynamic FC states directly from averaged time-courses (*i.e.*, in contrast to taking cross-regional correlation) within specified regional atlas. The high-dimensionality of the approach, however, requires a dimensionality-reduction step using PCA to extract information that gives the highest amount of variance in the data. This then allows the use of Bayesian temporal modeling (*i.e.*,, Hidden Markov Models) to extract repeating network patterns in a whole-brain fMRI data [Vidaurre et al., 2017]. A particular advantage of this approach is the elimination of crucial parameters such as the choice of window lengths [Leonardi and Van De Ville, 2015], as well as the ability to extract temporal characteristics (*e.g.*, probability of occurrence of each metastate at each time point) for each dynamic FC state.

In this work, we propose a new framework to estimate multiple functional brain networks, or *metastates*, by building upon an emerging field of graph learning [Dong et al., 2019]. From this point onwards, we will use the terms *metastates* and *networks* interchangeably to pertain to repeating functional spatial patterns in the brain. Starting from the whole-brain time-

series of parcellated fMRI data, we apply a recently introduced Bayesian-based [Tzikas et al., 2008] generative model for mixed signals living on multiple networks [Petric Maretic and Frossard, 2020]. The Graph Laplacian Mixture Model (GLMM) infers different metastates directly from the data, and at the same time, extracts the corresponding underlying graph that gives rise to these networks. We first demonstrate the powerful use of the tool by capturing distinct networks corresponding to each task epochs of task fMRI data downloaded from the Human Connectome Project (HCP) database. The performance of the method is validated by comparing the extracted probability of occurrence of each metastate to the timing of the task paradigm. We show that the extracted metastates consist of brain areas that are consistent with previously observed regions that are implicated in the corresponding task. We then apply the GLMM to RS fMRI data and obtain the most prevalent networks of spontaneously interacting brain areas.

Furthermore, in addition to understanding the organizational principles of human brain function, recent studies in neuroimaging are also starting to elucidate the role of structure in cognition and behavior [Preti and Van De Ville, 2019, Medaglia et al., 2018]. To extract meaningful interpretation of the learned graphs, we show that the estimated graph Laplacian matrices from fMRI reveal indicative similarities with the structural connectome (SC) derived from diffusion-weighted MRI (DW-MRI). We demonstrate that the degree of their similarity captures behaviorally-relevant gradient that is consistent with the previously observed macro-scale organization of the cortex that ranks low-level sensory processing to high-level cognitive ones [Preti and Van De Ville, 2019, Margulies et al., 2016].

3.3.2 Materials and Methods

Data and Preprocessing

The functional data were downloaded from the publicly available Human Connectome Project (HCP) database, WU-Minn Consortium. MRI acquisition protocols of the HCP are extensively described and discussed in an existing work [Glasser et al., 2013]. We used 50 subjects consisting of 4 sessions of RS scans (1200 volumes each, a total of 4800 frames), and 2 sessions each of task fMRI data (working memory, relational memory, social, language, emotion, and motor tasks). Functional volumes underwent the standard preprocessing steps [**?**]. All functional images were realigned to the mean functional volume for each participant. The realigned volumes were registered to the structural T1 data using rigid-body registration (SPM12, https://www.fil.ion.ucl.ac.uk), and were de-trended (i.e., constant, linear, quadratic) to remove signal drifts. Then, the images were smoothed using a Gaussian kernel with FWHM equal to 6mm. Finally, we used the Automated Anatomical Labeling (AAL, 90 regions) atlas that was resliced to fMRI resolution to parcellate fMRI volumes and compute regionally averaged fMRI signals. Meanwhile, structural data of each subject were downloaded from the HCP and were processed using MRtrix3 (http://www.mrtrix.org/). The SC of each subject was generated based on the total number of fibers connecting two regions divided by the volumes

of connecting regions. The normalization is done to ensure that the strength of the connection is not biased towards the size of the bundles. The final SC matrix was obtained by averaging all SC matrices of all subjects.

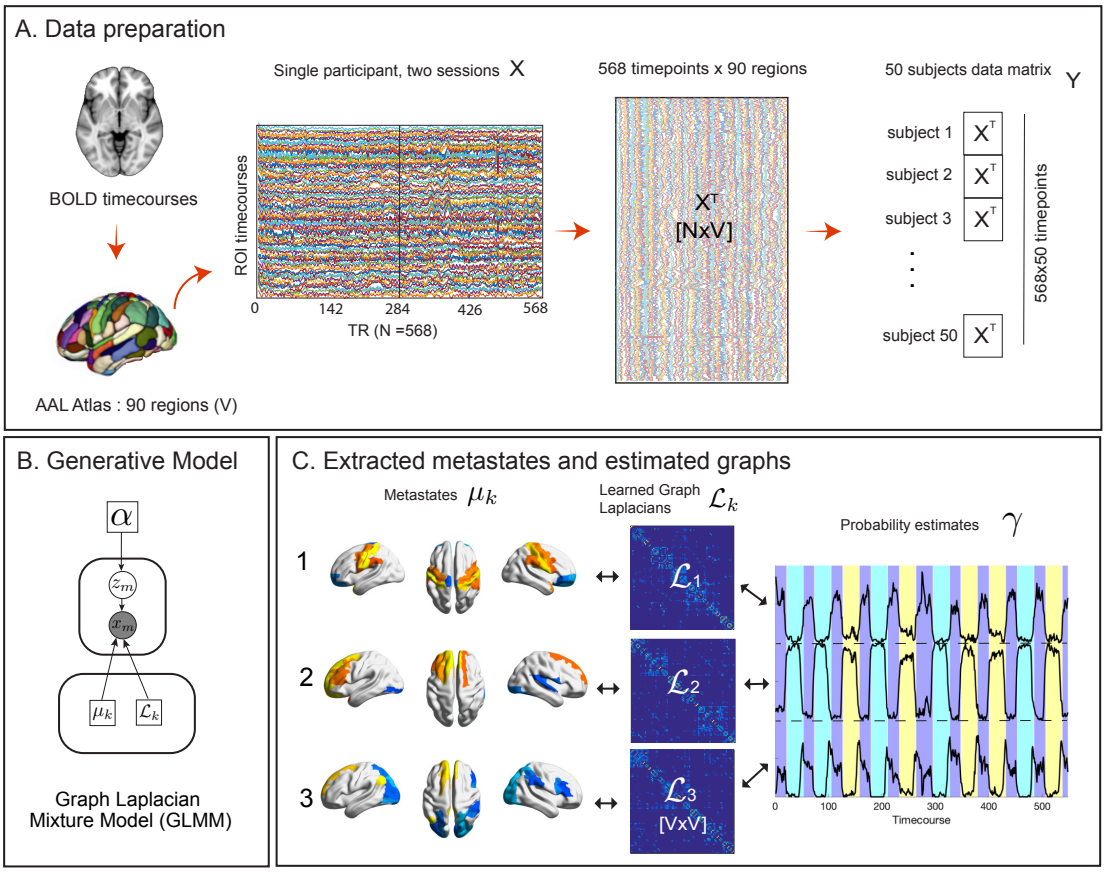

Figure 3.9 – **General workflow: from fMRI signals to the extraction of networks**. (a) Mean BOLD signals from four sessions of fMRI recordings are computed within each region of the AAL90 atlas and are concatenated together to form the subject data matrix **X**. The final data (**Y**) is obtained by concatenating 50 HCP subjects. (b) Plate notation of the generative model for extracting functional metastates and their corresponding estimated graphs represented through the Laplacians. (c) The output of the algorithm are K-number of metastates, their corresponding graph Laplacians, and the probability that each metastate would occur at a particular time-point. In the above example, we show K = 3.

Graph Laplacian Mixture Model

The Graph Laplacian Mixture Model is a generative model for a collection of signals which naturally live on different graphs [Petric Maretic and Frossard, 2020]. It assumes that there are several unknown networks, and each signal is generated from one of them. The interactions within each network are modeled with a graph Laplacian matrix. A graph Laplacian is defined as $L = D - A$, where D is a diagonal matrix of node degrees, and A is the adjacency matrix.

It is unknown a priori which signals correspond to which (unknown) network. The model fits the observed data to recover signal clusters, as well as the associated networks and graph structures. Specifically, with observed fMRI signals, it assumes they group into clusters defined by different unknown brain networks. The model recovers these clusters, together with brain networks and graph Laplacians which model interactions within each network.

Formally, each of the observed M samples $x_m \in \mathbb{R}^N$ belongs to exactly one cluster k determined by the graph Laplacian $L_k \in \mathbb{R}^{N\times N}$ and mean $\mu_k \in \mathbb{R}^N$. A binary latent variable $z_m \in \mathbb{R}^K$ has exactly one value 1, denoting the cluster k that x_m belongs to. Probability α_k models the relative size of cluster k and defines a prior distribution of x_m belonging to cluster k, namely $p(z_{m,k} = 1) = \alpha_k, \forall m$. Finally, the graph Laplacian $L_k \in \mathbb{R}^{N\times N}$ models smooth changes in signals. Large values in L_k thus capture pairs of vertices that change their values in similar ways. These connections can be seen as partial correlations between two vertices [Dempster, 1972]. Under these assumptions, signals in each cluster k follow a Gaussian distribution determined through the graph Laplacian[4] $x \sim \mathcal{N}(\mu_k, L_k^{-1})$ [Dong et al., 2016]:

$$p(x_m|z_{m,k} = 1) = p(x_m|\mu_k, L_k) = \mathcal{N}(\mu_k, L_k^{-1}) \tag{3.2}$$

Recall that the cluster of each signal is unknown. Marginalising over latent variables z denoting which cluster the signal x_m belongs to, we have:

$$\begin{aligned} p(x_m) &= \sum_{z_m} p(z_m)\, p(x_m|z_m) && (3.3)\\ &= \sum_{k=1}^{K} p(z_{m,k} = 1)\, p(x_m|z_{m,k} = 1) && (3.4)\\ &= \sum_{k=1}^{K} \alpha_k \mathcal{N}(\mu_k, L_k^{-1}), && (3.5)\\ \text{s.t.} \quad & L_k \in \mathcal{L}, \forall k && (3.6)\\ & \sum_{k=1}^{K} \alpha_k = 1, && (3.7)\\ & \alpha_k > 0, \forall k. && (3.8) \end{aligned}$$

Here (3.6) ensures that all L_ks are valid Laplacian matrices, $\mathcal{L} = \{L|L_{i,j} = L_{j,i} \leq 0, \forall i \neq j \ \& \ \sum_{j=1}^{N} L_{i,j} = 0, \forall i\}$. (3.7) and (3.8) ensure that α defines a valid probability measure.

GLMM algorithm

Given M observed N - dimensional signals, $X \in \mathbb{R}^{N\times M}$, we want to recover the parameters of our generative model (3.5). To do so, we will look at the maximum a posteriori (MAP) estimate for our parameters: probabilities $\alpha = \alpha_1 ... \alpha_K$, means $\mu = \mu_1 ... \mu_K$ and graph Laplacians $L = L_1 ... L_K$. Namely, assuming the data has been sampled independently from the distribution in

[4] Note that L_k^{-1} here denotes a pseudo-inverse of the graph Laplacian L_k

(3.5), and allowing for a prior on the graph structure, we want to maximise over the a posteriori distribution of our model:

$$\underset{\alpha,\mu,L}{\operatorname{argmax}} \ \ln p(\alpha,\mu,L|X) \tag{3.9}$$

$$\propto \underset{\alpha,\mu,L}{\operatorname{argmax}} \ \ln p(X|\alpha,\mu,L)p(L) \tag{3.10}$$

$$= \underset{\alpha,\mu,L}{\operatorname{argmax}} \ \ln \prod_{m=1}^{M} p(x_m|\alpha,\mu,L)p(L) \tag{3.11}$$

$$= \underset{\alpha,\mu,L}{\operatorname{argmax}} \ \ln \prod_{m=1}^{M} \sum_{k=1}^{K} \alpha_k \mathcal{N}(x_m|\mu_k, L_k^{-1})p(L_k) \tag{3.12}$$

$$= \underset{\alpha,\mu,L}{\operatorname{argmax}} \ \sum_{m=1}^{M} \ln \sum_{k=1}^{K} \alpha_k \mathcal{N}(x_m|\mu_k, L_k^{-1})p(L_k), \tag{3.13}$$

which does not have a closed form solution. The problem is simplified through posterior probabilities $\gamma \in \mathbb{R}^{M\times K}$, with $\gamma_{m,k}$ modelling the probability that signal x_m belongs to cluster k:

$$\gamma_{m,k} = p(z_{m,k} = 1|x_m, \mu_k, L_k) \tag{3.14}$$

$$= \frac{p(z_{m,k} = 1)p(x_m|z_{m,k} = 1, \mu_k, L_k)}{\sum_{l=1}^{K} p(z_{m,l} = 1)p(x_m|z_{m,l} = 1, \mu_l, L_l)} \tag{3.15}$$

$$= \frac{\alpha_k N(x_m|\mu_k, g^2(L_k))}{\sum_{l=1}^{K} \alpha_l N(x_m|\mu_l, g^2(L_l))} \tag{3.16}$$

The parameters can now be estimated iteratively using expectation maximisation (EM) algorithm. The graph Laplacians are estimated with a graph learning scheme, while the other parameters have closed-form solutions [Petric Maretic and Frossard, 2020].

Application to fMRI data

The general workflow of this work is summarized in Figure 3.9. We began by extracting the mean BOLD signals within regions of the AAL atlas for each session of each subject. We built a data matrix $\mathbf{X}$ that contains the timecourses of all AAL 90 regions. The timecourses of all sessions and all subjects are then concatenated together to form a huge data matrix $\mathbf{Y}$. This is then fed to the GLMM algorithm which estimates the means or *metastates* (μ_k), learned graphs represented in the output as the graph Laplacians (L_k), and signal clustering probabilities (γ).

We then validated the performance of the proposed framework using task fMRI by demonstrating that the timing of the task paradigms are captured by the averaged γ-values across all subjects. The hyper-parameters of the model (Δ, θ) are optimized using Normalized Mutual Information (NMI) using the task paradigm as the reference ground truth. Meanwhile, the number of clusters (K) is optimized using a feedback step where t cannot expect it to correspond to the number of tasks, since the metastates may be formed as a combination of tasks.

For that reason, we evaluated the results by iteratively changing K and performing a visual evaluation of the dynamics of γ with respect to the experimental task paradigm. Doing so, we observe that the meaningful clusters consistently appear in a very similar fashion, regardless of the proposed number of clusters K. Furthermore, fixing K in such a way resulted in better overall accuracy, suggesting that the number of optimal clusters might be different from the number estimated through more traditional methods.

As for RS, which is lacking a ground truth, the number of clusters K was chosen based on the optimized silhouette measure and the consensus clustering procedure [Monti et al., 2003], which is a resampling-based method for optimal class discovery. Both metrics suggested 3 as the optimal number of clusters. K has eventually been varied according to the procedure mentioned above, in order to capture multiple networks. In practice, changing the optimal number of clusters does not seem to affect the final estimation, instead, it opens the possibility of inferring more networks.

Comparison of learned functional graphs to brain structure

An important benefit of the GLMM algorithm compared to other clustering methods is the estimation of graph Laplacians (L). Not only this does help in obtaining more meaningful networks or metastates (means), but it also conveys much more information and details about our clusters. From here onwards, we use the term *graphs* to refer to the learned graph Laplacians and *networks* or *metastates* to refer to the functional means or clusters.

We explored the similarity of the learned functional graphs to the structural connectome derived from DW-MRI by using the spectral euclidean difference as a metric. This is done by first decomposing the adjacency matrices (A) of each graph derived from the Laplacians ($A = D - L$) into their constituent eigenvalues and eigenvectors. The eigenspectrum of A are transformed using the procrustes algorithm [Kendall, 1989, Goodall, 1991] to match the ordering of the eigenvectors of SC.

The output of the Procrustes transform is the rotation matrix that is used to retrieve the transformed (functional) eigenvalues. In particular, given the rotation matrix Z generated by the transformation, the rotated eigenvalues are computed as follows:

$$\hat{\lambda} = ZEW, \tag{3.17}$$

where W is the weighted adjacency matrix of the functional network taken into consideration and E is the vector of its original eigenvalues.

Considering all the functional networks inferred, regardless of the task, we computed the euclidean spectral distance between the eigenvalues of the SC and the transformed graphs. This metric reflects the degree of similarity between the learned graphs in each task and the anatomy. Thus, each network will have a certain score of distance that indicates how that function couples with structure: the smaller the metric, the closer the estimated graph to the

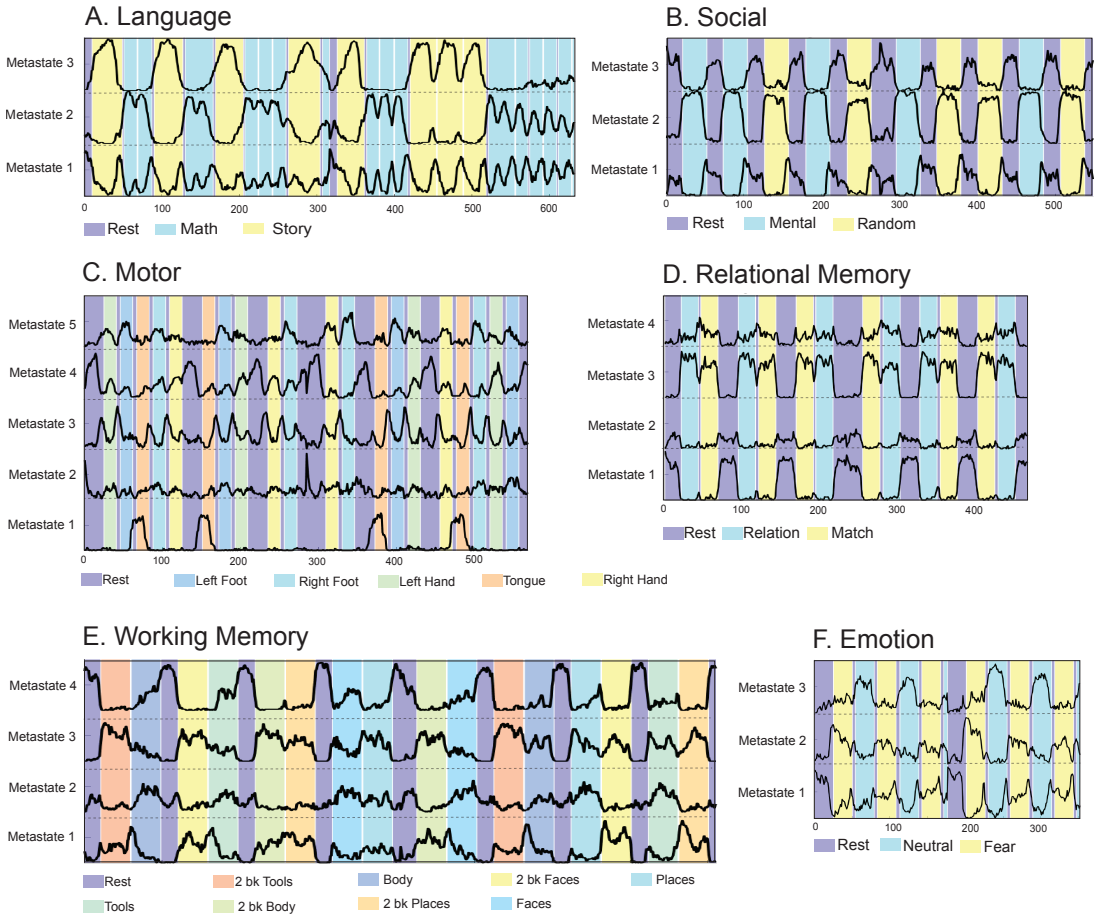

Figure 3.10 – **Estimated timecourses with respect to selected task paradigms.** The γ values are plotted over the experimental paradigm of Language (A), Emotion (B), Social (C) and Relational Memory (D) tasks. The black signal corresponds to the probability of belonging to a specific *metastate*.

structure. We perfomed bootstrapping to extract networks that show statistically different values. We then performed a Neurosynth meta-analysis similar to the one implemented by Margulies *et.al.* [Margulies et al., 2016] and Preti *et.al.* [Preti and Van De Ville, 2019], where topic terms from the database [https://neurosynth.org/] were matched with the derived euclidean spectral difference of the functional graphs and the SC. The active areas of the estimated networks were correlated with the database present in literature, thereby capturing a number of possible behavioral domains. The same 24 topics adopted these two studies were considered. This behavioral gradient has been sorted according to a weighted mean of the resulting z-statistics.

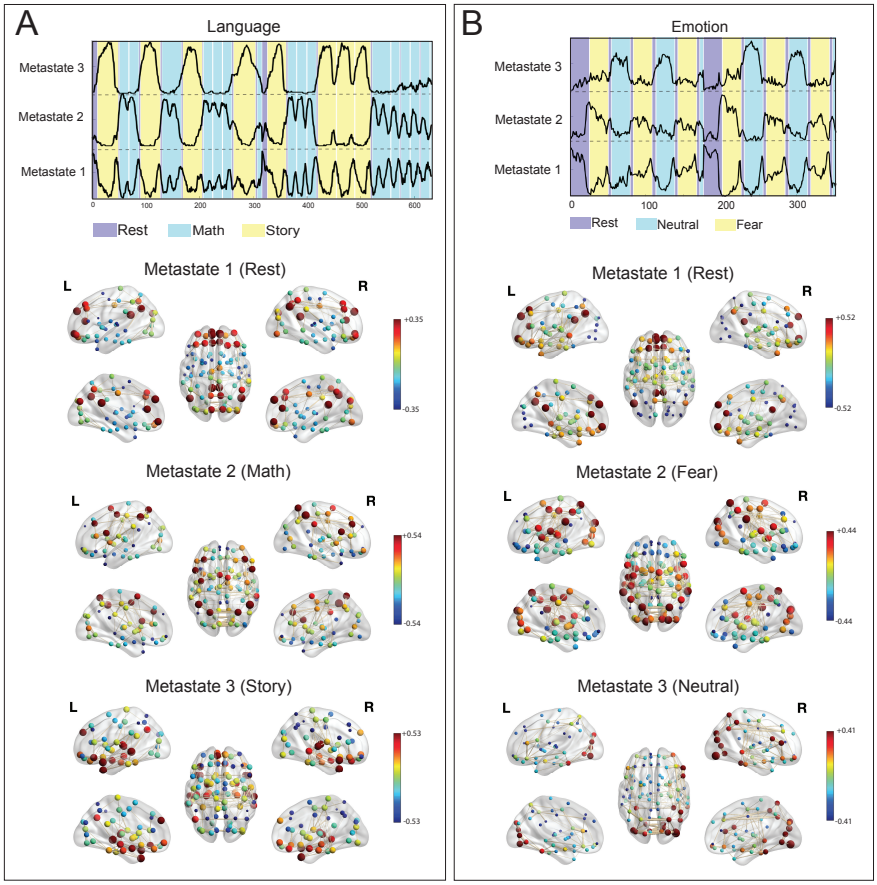

Figure 3.11 – **Spatial patterns corresponding to each epochs** in the (A) Language task and (B) Emotional task paradigms. Each node corresponds to each brain region in the AAL atlas. The colors denote the signs of the extracted clusters (means), the edges denote the connectivity derived from the estimated graph Laplacians, and the size the nodes indicate the degree of each node's connections. The rest epochs of the two tasks show regions of the default mode network.

3.3.3 Results

Estimated network time-courses are consistent with the timing of the task paradigms

The proposed framework extracts consistent patterns of brain activity in each of the considered tasks, and along with it is the estimation of their likelihood to occur at each time-point. Figure 5.2 displays the estimated time-courses for six selected tasks overlaid with the task paradigm which are distinguished in colors. Metastate 2 and 3 of the Language task capture time-points when subjects undergo the *Story* and *Math* epochs, respectively while Metastate 1 corresponds to the resting epoch of the task. Meanwhile, Metastate 3 consistently captures the RS epoch in the *Social* task. However, the algorithm does not distinguish the metastates corresponding to *Mental* and *Random* phases since the likelihood estimation of Metastate 2 corresponded to both task epochs. On the other hand, Metastate 3 captures the transition

between the rest and task epochs. The same observation can be made in the other tasks, in particular, the motor, relational, and working memory. In each of these tasks, some metastates capture transitions between epochs and uniquely identifies whether it is a transition between rest and task epoch or a transition between two different task epochs *i.e.,* Metastate 4 in Relational Memory captures transitions between *Relation* and *Match*, while Metastate 2 captures transitions between rest and any of the two tasks.

GLMM captures spatial networks corresponding to each task consistent with their known established neurophysiological descriptions

Figure 5.3(A) and (B) displays an example set of results that we obtained when running the GLMM algorithm to the Language and Emotion tasks, respectively. The RS epochs of both tasks reveal metastates that are akin to the DMN network. The regions implicated in Metastate 2 (*Math*) include the parietal areas (superior and inferior) and the frontal region (e.g., middle frontal gyrus, opercular part of the inferior frontal gyrus). These regions are well in-line with the established associated areas corresponding to numbers and calculations [Arsalidou et al., 2018]. On the other hand, active areas in Metastate 3 (*Story*) are the hippocampus, frontal, and the bilateral superior and anterior temporal cortex, consistent with previously observed regions implicated with story processing tasks [Barch et al., 2013].

Meanwhile, Metastate 1 corresponding to the *Rest* epoch of Emotion task again captures the DMN pattern. Metastates 2 and 3 corresponding to the *Fear* and *Neutral* phases, respectively both show stark differences in terms of the regions activated. In particular, *Fear* triggers activations in the bilateral central gyrus, superior occipital gyrus, and the parietal cortices. These regions cover the somatosensory cortex, which is a known region responsible not only for processing sensory information from various parts of the body but also for emotional processing, including generation of emotional states and emotion regulation [Dolan, 2002, Bufalari et al., 2007, Kropf et al., 2019]. Meanwhile, the metastate for *Neutral* shows activation of visual regions that are strongly related to shape [Anzai et al., 2007, Hegdé and Van Essen, 2000]. Unexpectedly, we found relatively low means in the amygdala, a well-known region that is typically activated in emotional matching paradigms [Hariri and Mazziotta, 2000, Preckel et al., 2019].

Extracted metastates during rest

After validating the outcome of the proposed framework with task fMRI using the known experimental task paradigm as the ground truth, we applied the method to RS fMRI. Figure 3.12 displays the estimated metastates together with the approximated likelihood to occur at each time-point for one representative subject. This likelihood of occurrence gives a proxy of the metastate's activity profiles. As expected, we observe the DMN to be highly occurring. We found a metastate that clearly covers areas of the visual cortex, another metastate that shows regions corresponding to the auditory network and some portions of the frontoparietal

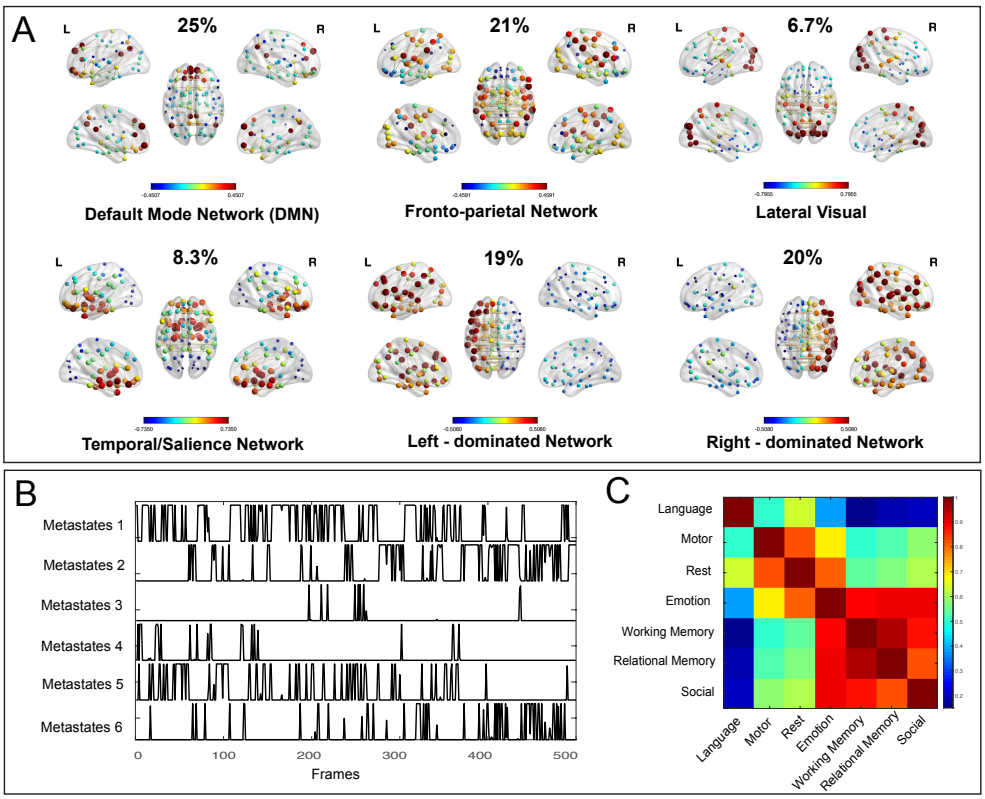

Figure 3.12 – **Extracted spatial maps during rest** (A) Six metastates corresponding to the RS data. Values in percentage show the occurrences of each metastate across all 50 subjects considered. (B) Example activtity profile of each metastate for an example subject. (C) Correlation of estimated means corresponding to DMN-networks extracted from rest and all task fMRI data.

cortex, and another one that contains the bilateral temporal cortices and the insula, which is analogous to the salience network. Additionally, we also found separately clustered activations in the left and right portion of the brain.

Comparison of learned graphs that give rise to DMN patterns

We found that metastates corresponding to the rest epochs of the tasks always corresponded to the DMN, which is expected. To understand the nature of the recovered DMN-related metastates, we computed the correlation across all pair-wise networks corresponding to each task and RS data. Figure 3.12(C) displays the correlation between the extracted DMN-related metastates from all tasks and RS fMRI. The $x-$ and y-axes were sorted from lowest to highest correlation, showing a distinction between tasks that require low-level sensory-motor functions (*e.g.*, language, motor, rest) versus tasks that associated to high-level cognition (*e.g.*, emotion, working memory, relational memory, and social).

Learned graphs bear similarity to the underlying brain structure and their relation reveals behaviorally-relevant organization

For each of the tasks evaluated, we obtained different metastates and corresponding graph Laplacian matrices. We, therefore, have a total of 28 metastates across all tasks, and consequently, also 28 graph Laplacian matrices. We evaluated the similarity of these graph Laplacian matrices to the underlying brain structure derived from DW-MRI. Figure 3.13(A) shows the group-averaged structural connectome across all subjects considered in this work, and Figure 3.13(B) displays an example adjacency matrix computed from the estimated graph Laplacian matrix of an example task. The adjacency matrix is visually more sparse than the SC. Moreover, a direct statistical measure (*i.e.*, Pearson correlation) between the estimated GLMM adjacency matrices and SC reveals a similarity values within the range of $r = 0.48$ to $r = 0.60$ as is shown in Figure 3.13(C). Compared to the classical FC-SC relationship where FC is obtained by correlating inter-regional BOLD time-courses using either Pearson correlation or Partial correlation, the strength of FC-SC coherence is much higher using GLMM-based FC.

We then performed a meta-analysis [Yarkoni et al., 2011] of this SC-FC relationship by using spectral difference as a measure. In essence, this translates to a dissimilarity index between the SC and the GLMM-based FC obtained in all tasks and RS fMRI. Figure 3.13(D) shows the results of the meta-analysis performed using the Neurosynth database, analogous to previous works of Preti, et.al. [Preti and Van De Ville, 2019] and Marguiles, et.al., [Margulies et al., 2016]. We displayed seven metastates that showed a statistically significant spectral difference. This is shown in the figure as columns which is sorted in increasing order. The 24 behavioral topics describing the rows are sorted according to their z-score. The ordering of both axes reveals an organized behavioral gradient that naturally arises from this SC-FC relationship. Low-level processes (highlighted in blue) reveal a general low spectral difference while high-level cognitive behavior (in red) show high difference. A diagonal trend can be glimpsed.

An important aspect to note in this analysis is the use of spectral difference to define a decoupling index between SC and FC. This choice was motivated by the SC-FC decoupling measure introduced in previous work [Preti and Van De Ville, 2019]. In order to be consistent, we opted to use a dissimilarity measure instead of similarity (*e.g.*, correlation). It is to be noted, however, that results do not change when using either correlation or spectral difference, as can be seen in Figure A.1.

3.3.4 Discussion

General findings

We proposed a new framework for extracting repeating functional brain patterns based on a generative model that assumes functional data as a collection of signals that live on multiple graphs. Using this technique, we were able to obtain: (1) distinct sets of networks that the brain cycles to, (2) the likelihood of these networks to occur at each time-point, and (3) the

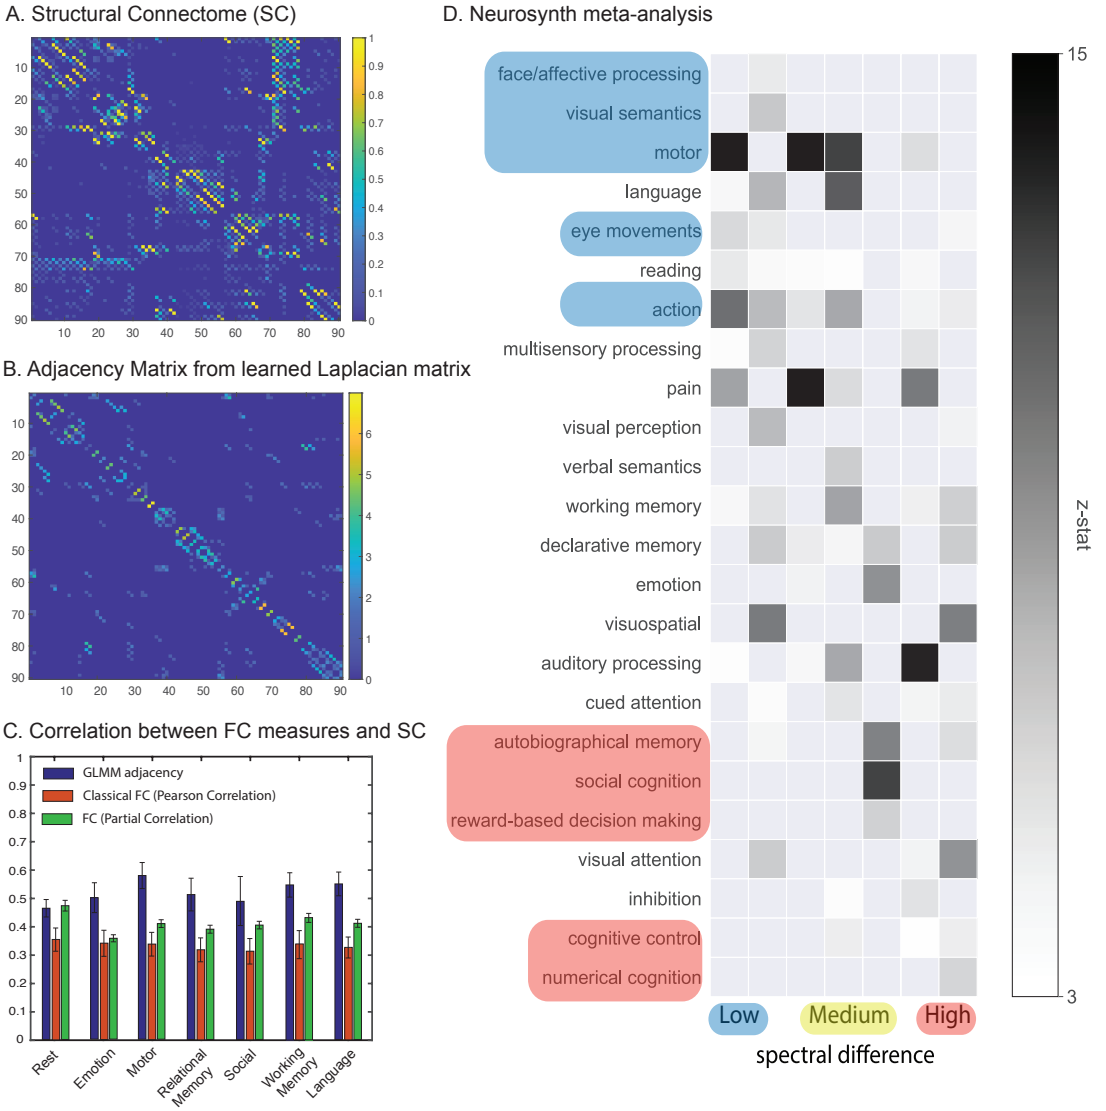

Figure 3.13 – **Learned Laplacian matrix and its relation to the structural connectome and behavior.** (A) Group-averaged SC matrix corresponding to all subjects considered in this work. (B) Adjacency matrix extracted from an example Laplacian matrix. (C) Similarity between SC and different FC measures: using (1) GLMM-based adjacency extracted from learned graphs and classical measures of FC using (2) Pearson Correlation and (3) Partial Correlation. (D) Meta-analysis of SC-FC relationship using spectral distance as a metric. The distinction between low level and high-level processing arises thanks to the sorting of the spectral difference and the z-score assigned to a behavioral topic.

underlying graphs that describe the interactions of the regions composing these networks. We validated the approach by demonstrating that the extracted metastates corresponded to spatial patterns that are consistent with their established neurophysiological descriptions. We also showed that the probability of these networks to occur at each time-point generally captured the timing of the task paradigms we evaluated. Then, we applied the approach to

RS data to reveal some of the well-known RS networks, such as the DMN, visual, auditory/attention, and salience networks. Finally, we took advantage of the estimated graph Laplacian matrices to understand the interactions of the regions composing these networks, and how they are related to the underlying brain structure obtained from DW-MRI. We showed that the adjacency matrices computed from the graph Laplacians bear closer similarity to SC than conventionally defined FC measures, *e.g.,* Pearson correlation of inter-regional BOLD signals. We then demonstrated that the degree of similarity between the learned graphs and the SC captures behaviorally-relevant gradient which distinguishes low-level and high-level cognitive processes that is consistent with previously observed macro-scale organization in the cortex [Preti and Van De Ville, 2019, Margulies et al., 2016].

Within-network regional interactions of brain function are coupled with anatomical connectivity

The relationship between brain structure and function is one of the fundamental questions in neuroscience [Goni et al., 2014, Atasoy et al., 2016, Betzel et al., 2016c, Gu et al., 2015, Tarun et al., 2020b]. One of the most straightforward method for comparing structure-function relationship is done by correlating FC with SC (*e.g.,* Pearson correlation) [Honey et al., 2009, Supekar et al., 2010, Horn et al., 2014]. The simplicity of the approach, however, bears with it indispensable limitations. Pearson correlation reflects the marginal association between regions and does not capture a direct relationship between two variables without confounding effects of other variables. To address this issue, partial correlation has been shown to improve the results by regressing out the effects of other variables [Smith et al., 2011, Wang et al., 2016]. In our approach, the GLMM algorithm extracts both an estimate of the (1) the mean signals that are activated at every time instance as well as (2) a connectivity matrix (*i.e.*, Laplacian matrix) which captures regions that are directly and coherently fluctuating together, without the confounds of the activity of other areas. As can be observed in Figure 3.13(B), the adjacency matrix derived from an example Laplacian matrix is more sparse than SC. It is therefore showing fewer connections than classical FC [Honey et al., 2009], thus highlighting only very specific regions that are implicated in the recovery of each metastate. We also show in (Figure 3.13(C)) that compared to FC's computed using classical Pearson correlation analysis and partial correlation, the GLMM adjacency matrices bear higher similarity to SC. Two reasons are believed to be the cause of this observation. First, unlike classical FC, GLMM adjacency matrices do not include spurious connections, and only captures direct interactions between regions within a network. Second, the GLMM method is a dynamic approach, whereby each GLMM FC matrix corresponds to a particular temporal window (*i.e.*, each task epoch corresponds to a specific metastate, and consequently to a specific FC matrix). It has been shown in previous work [Liégeois et al., 2016] that there are phases of synchronization and desynchronization between FC and SC matrices when sliding-window analysis is applied instead of static FC. Our findings therefore provide further evidence on the role of structure as a transitory architecture between different functional brain state. The fact that metastate-specific graphs show closer similarity to the fixed SC (Figure 3.13(C))

reflects the inherent preferential organization of the brain to be constrained by its underlying anatomical bedrock [Honey et al., 2009, Liégeois et al., 2016, Cabral et al., 2011].

Behaviorally-relevant macro-scale organization of the cortex is linked to brain functional-anatomical coupling

In addition to the findings above, there is a growing evidence for strongly linked brain structure and function that is related to cognition and behavior [Preti and Van De Ville, 2019, Medaglia et al., 2018]. Strongly coupled SC and FC are observed to be associated with low-level sensory-motor functions while strongly dissimilar SC and FC correspond to high-level processing [Preti and Van De Ville, 2019]. A similar macroscale organization in the cortex is also observed by another study using only FC measures [Margulies et al., 2016]. These findings are well in-line with what we have found in Figure 3.13(D) using the same meta-analysis performed in these two studies. Although a diagonal trend is not as apparent, we can see a clear ordering of the behavioral topics from top to bottom, from low-level processes (*e.g.*, face/affective processing, visual semantics, motor, eye movements, and action) to high-level ones (*e.g.*, autobiographical memory, social cognition, reward-based decision making, cognitive control, and numerical cognition). This organization naturally emerged together with the ordering of the spectral difference between GLMM-based FC and SC from low distance (*i.e.*, similar) to high distance (*i.e.*, dissimilar). These findings confirm that the strength of FC-SC coupling follows a behaviorally-relevant gradient that describes the topographical organization of cortex in humans [Margulies et al., 2016, Preti and Van De Ville, 2019].

Methodological perspectives

Dynamic analyses of functional imaging data during rest and tasks have been going on for a decade [Chang and Glover, 2010]. Since its inception, several methodological advances have been introduced to probe the functional organization of the brain in a dynamical point of view. A recent review of the existing methodological tools [Preti et al., 2017] has classified existing approaches into four distinct groups of methods: (i) sliding-window correlations, (ii) frame-wise analyses [Liu and Duyn, 2013, Karahanoglu and Van De Ville, 2015], (iii) state modeling [Leonardi et al., 2013, Allen et al., 2014], and (iv) temporal modeling [Vidaurre et al., 2017]. We situate our proposed technique to be an integration of state and temporal modeling, whereby the GLMM approach captures the dFC states and their corresponding activity timecourses, akin to how HMM models [Vidaurre et al., 2017, **?**] are able to capture various brain states and their likelihood to occur at each time-point. Unlike HMMs, however, GLMM extracts the metastates directly from the averaged BOLD data within parcellated brain regions without a dimensionality reduction step (*e.g.*, PCA). The reason for this is the implicit dimensionality reduction which occurs in GLMM by imposing a Laplacian structure on the inferred graphs [Petric Maretic and Frossard, 2020]. The model has been shown to outperform standard clustering methods on high-dimensional tasks, even when intuitively there is no inherent graph structure in the data, achieving better clustering accuracy. Finally, the inferred

Laplacians add a strong layer of interpretability to our findings and offers multiple possibilities for further analysis and understanding of brain networks. The extracted individual graph structures describe the interactions of regions comprising the dFC states. Apart from the enhanced visual understanding of interactions between the brain regions, these graphs enable further comparison of different brain networks, as shown in Figures 3.12 and 3.13.

Technical Limitations

While the proposed method tries to capture more information about recovered brain networks, it still suffers from numerous technical limitations, mostly due to the oversimplification of a very complex problem. Firstly, the method makes the assumption that each fMRI signal corresponds to exactly one cluster, ie. one brain network. When fitting the model, it does return probabilities of the signal belonging to each of the clusters, but that is largely different from assuming that one signal actually originated as a combination of several networks, a phenomena that could very well be present in practice. Furthermore, the method takes no temporal information into account. This fact helped us evaluate the method, showing that even with no temporal information meaningful clusters are successfully obtained. However, in practice the temporal information is valuable and incorporating it could lead to more accurate results, which this method does not support at the moment. Finally, while the additional information in terms of inferred graph Laplacians is very interesting and can lead to more accurate results, the method does not support imposing any similarities on the graph Laplacian structures. Namely, imposing any additional information known from the literature, as well as similarities or differences between networks could render the inference problem more simple, thus resulting in a higher accuracy and more reliable findings. These limitations could be addressed in future work, where the method described in this paper could serve as baseline, enabling easy comparison with more complex models and clearly underlying possible benefits obtained from the incorporation of additional information and literature knowledge.

3.3.5 Conclusion

This study presents a new framework for uncovering meaningful network-level representations of brain function, each of which is supported by a separately defined graph that describes the interactions of implicated regions. The extracted graphs are sparse and only recover links between regions that have direct interactions. These graphs also bear closer similarity to structural connectivity, more than classical definitions of FC (*e.g.*., Pearson correlation, Partial correlation). We demonstrate that the degree of dissimilarity between brain function and structure gives rise to a macro-scale gradient in the cortex which reflects the preferential tendency of brain regions involved in low-level sensory-motor processing to be similar to structure, and in contrast, those involved in high-level cognitive processing are discordant. Overall, our findings validate the potential of the proposed technique in providing additional perspective with regards to understanding the organizational principles of human brain

function and its relation to structure.

4 High-resolution voxel-level brain graphs

4.1 General Motivation

Graph modeling of structural connectivity has proven itself very useful in characterizing brain network topology [Bassett and Bullmore, 2017]. Moreover, graph measures derived from functional connectivity have also been demonstrated to present insightful neurobiological underpinnings [Sporns, 2018], which makes them effective biomarkers for delineating healthy and disease [Wang et al., 2010]. Recently, an increasing interest is given on the use of graph signal processing (GSP) to analyze the interplay between brain structure and function [Huang et al., 2018a]. This family of method is essentially a step more advanced than classical network neuroscience since it gives an integrated approach to study brain dynamics on the basis of the anatomy. Notably, the latest studies that utilized GSP tools have started to elucidate the importance of the structure-function link in supporting coordinated cognition and behavior [Preti and Van De Ville, 2019, Medaglia et al., 2018]. It is therefore very timely to survey this approach, and in particular, take a closer look on the definition of the brain graph itself, from which, all succeeding GSP operations are based upon.

Indeed, most of the GSP studies are so far limited to defining nodes as regions from a brain atlas. This Chapter is dedicated to the design and construction of high-resolution voxel-level brain graphs that are meant to extend current region-wise GSP approaches to voxel-level perspective. In particular, we propose two designs for constructing the brain graphs, one based on diffusion tensor imaging and another based on diffusion orientation distribution functions. We differentiate the designs by characterizing their respective degree distributions, followed by a comparison of their eigenvalue spectra to the weighted zero crossings of their corresponding eigenvectors. We assess the intrinsic dimensionality of the eigenspace using a Procrustes validation scheme that characterizes the inherent inter-subject variability. Finally, we demonstrate the powerful use of the proposed technique in integrating brain structure and function by decomposing resting-state and task fMRI in terms of the Laplacian's lowest frequency *eigenmodes*. We end the chapter with enumerated possibilities for future research directions, in particular, in the field of GSP applied to functional brain imaging.

4.2 Journal Article: Voxel-Wise Brain Graphs from Diffusion-Weighted MRI: Spectral Analysis and Application to Functional MRI

A. Tarun, D. Abramian, M. Larsson, H. Behjat*, and D. Van De Ville*

Part of the manuscript is published in 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). IEEE. p. 159–163.
Submitted Full Research Article. Under Revisions for publication in IEEE Transactions on Medical Imaging

Author Contributions

A. Tarun conceptualized, designed, and developed the methodology, performed all stages of the analysis, from validation, formal analysis, writing of the original draft, and editing of the manuscript, **D. Abramian** and **M. Larsson** did preliminary exploration in developing the ODF-based method and reviewed the manuscript, **H. Behjat** and **D. Van De Ville** share senior authorship. Both supervised, conceptualized research, offered continuous guidance, and equally contributed in reviewing and editing the manuscript.

Abstract

Understanding the interaction between brain structure and function is key to elucidate how human cognition and behavior arise. Recent studies in computational neuroimaging employ signal processing operations to analyze brain activity signals associated to the nodes of structural brain graphs. Current graph models for characterizing structural connectivity are mostly limited to defining nodes as gray matter regions from a brain atlas, while incorporating white matter connectivity as edge weights between nodes. Here, we model white matter explicitly as nodes of a high-resolution voxel-wise brain graph—containing around 750 000 voxels—for which local connectivity is computed from diffusion-weighted MRI data. We propose two methods for encoding structural connectivity, one based on diffusion tensor imaging and the other based on diffusion orientation distribution functions. To assess the properties of the brain graph, we apply spectral analysis to the graph Laplacian operator. The Laplacian eigenmodes are spatially highly resolved, reflecting distributed structural features that correspond to major white matter tracts. We study the intrinsic dimensionality of the eigenspace of the proposed graphs across 100 subjects through a Procrustes validation scheme that characterizes inter-subject variability. We then demonstrate the extent to which functional MRI data are shaped by the underlying anatomical structure by decomposing task and resting-state data using low-frequency eigenmodes. The proposed methods open a range of possibilities for new research avenues, especially in jointly integrating brain structure and function.

4.2.1 Introduction

Magnetic resonance imaging (MRI) has become an essential neuroimaging modality that has provided effective means to map the brain's anatomical scaffold, using diffusion-weighted MRI (DW-MRI), and, in parallel, to track brain neural activity using functional MRI (fMRI). In the age of bio-banks, extensive datasets that include diffusion and functional data on the same set of subjects, such as the Human Connectome Project (HCP) [Essen et al., 2013], have been made freely available, and with such accessibility, various methodological developments that aimed to integrate the two modalities have emerged.

Computational neuroimaging has successfully adopted graph theory and made it evolve into a new field of interdisciplinary research called *network neuroscience* [Bassett and Bullmore, 2017]. Structural and functional *connectomes* are independently defined, and are analyzed using graph theory measures to provide a better understanding of their organizing network principles and to determine their influence in cognitive function [Bassett and Bullmore, 2017, Bullmore and Sporns, 2009]. Recently, there has been increasing interest in addressing neuroscientific questions from the point of view of brain dynamics [Preti et al., 2017, Bolton et al., 2020], along with which lies the fundamental question of how the underlying brain anatomy supports the emergence of spatially and temporally varying distributed patterns of functional activity. In this perspective, it is fitting to consider advancing network neuroscience to a more unifying analysis approach that accounts for the interplay between brain structure and function. Leveraging principles from the field of graph signal processing (GSP) [Shuman et al., 2013, Ortega et al., 2018, Stankovic et al., 2020] that has emerged in the last decade, a new neuroimaging framework is taking form, in which functional data are interpreted as defined atop the structural connectome, treated as time-dependent graph signals, and, in turn, processed using methods that are informed by the structure; e.g., see [Huang et al., 2018a, Behjat et al., 2015, Atasoy et al., 2017].

The GSP framework generalizes principles from classical discrete signal processing for time series to data defined on irregular domains. It has been successfully applied to various datasets in many different fields, such as image processing [Cheung et al., 2018], sensor networks and smart grids [Jablonski, 2017, Stankovic et al., 2019], community detection in graphs [Tremblay and Borgnat, 2014], urban analysis [Jain et al., 2014], energy networks [He et al., 2018], transportation networks [Hasanzadeh et al., 2017] and food industry [Masoumi et al., 2020]. The fundamental GSP concepts that have to date been employed for neuroimaging are (i) the graph Fourier transform (GFT), e.g. through which functional signals on region-wise graphs have been characterized by connectome-specific components [Atasoy et al., 2017], and (ii) spectral graph filters, e.g. using which low- and high-level cognitive functions were associated to the relative smoothness of functional signals over the structural graph [Medaglia et al., 2018, Preti and Van De Ville, 2019]. Both of these concepts are rooted in spectral graph theory built on the graph Laplacian matrix [Chung, 1997], wherein the Laplacian eigenvectors form the basis of the GFT, in the same manner that complex exponentials form the basis of the conventional Fourier transform. As such, the Laplacian eigenvalues carry a notion

of frequency (spatial saliency) in the context of structural brain graphs. We refer the reader to [Huang et al., 2018a] for a detailed review of the application of GSP on region-wise human brain connectomes.

Research questions that can be answered through GSP-based analyses are fundamentally reliant on the definition of the graph, from which all succeeding operations are derived. The majority of existing applications of GSP in functional brain imaging are limited to region-wise analyses, where the regions are defined using a priori brain atlases containing 100 to 1000 regions. Motivated by the promising results from existing GSP studies in linking brain structure and function, we propose the construction of large-scale brain graphs at the resolution of voxels. We introduce two methods to derive the strength of connection between adjacent voxels using DW-MRI data, one based on tensors estimated from diffusion tensor imaging (DTI) [Basser et al., 1994], and the other based on diffusion orientation distribution functions (ODF) estimated from high angular resolution diffusion imaging (HARDI) data [Tuch, 2004]. The tensor and ODF models both aim to represent structural information on the intra-voxel axon fiber arrangement [Le Bihan, 2003]. The diffusion tensor model can be seen as a multivariate Gaussian describing the distribution of fiber bundle alignment, whereas the diffusion ODF model defines the radial projection of the diffusion function and estimates the empirical distribution of water diffusion. For a detailed description of the DTI and ODF signal reconstruction methods, we refer to [Alexander et al., 2007] and [Tuch, 2004], respectively.

We validate the two proposed structure-encoding graph methods, in particular, in relation to their spectral properties, using diffusion data of the 100 unrelated subjects from the HCP dataset [Essen et al., 2013]. We then investigate the intrinsic dimensionality of their eigenspace across subjects using a Procrustes validation scheme that characterizes inter-subject variability. Finally, we demonstrate the relevance of the proposed high-resolution graphs within a GSP setting, particularly through studying the energy spectral density of resting-state and task fMRI data on these graphs. We conclude the paper by discussing potential future research avenues in using the proposed graph designs in addressing the interaction between brain structure and function.

4.2.2 Preliminaries

Graphs and their spectra

Let $\mathcal{G} := (\mathcal{N}, \mathbf{A})$ denote an undirected, single-connected, weighted graph, consisting of a node set $\mathcal{N}$, where $|\mathcal{N}| = N$, and a symmetric $N \times N$ weighted adjacency matrix $\mathbf{A}$, wherein any of its nonzero elements a_{ij} represent the weight of an edge (i, j) in the graph. The graph degree matrix $\mathbf{D}$ is diagonal with elements $d_{i,i} = \sum_{j=1}^{N} a_{ij}$. One useful operator defined on $\mathcal{G}$ is the normalized graph Laplacian, which is defined as the symmetric and positive semi-definite matrix

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}, \tag{4.1}$$

where $\mathbf{I}$ denotes the identity matrix. Through eigendecomposition, $\mathbf{L}$ can be diagonalized as $\mathbf{L} = \mathbf{U}\Lambda\mathbf{U}^\top$, where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, ..., \mathbf{u}_N]$ is an orthonormal matrix stacking the eigenvectors $\mathbf{u}_i$, $i = 1, \ldots, N$, in its columns, and Λ is a diagonal matrix whose diagonal elements are the corresponding eigenvalues $\Lambda_{i,i} = \lambda_i$, $i = 1, \ldots, N$. Due to symmetry and positive semi-definiteness of $\mathbf{L}$, its eigenvalues are real and non-negative, and its eigenvectors form an orthonormal set; *i.e.,* $\mathbf{u}_i^\top \mathbf{u}_j = \delta_{i-j}$, where δ is the Kronecker-delta. Without loss of generality, we assume that the diagonal elements in Λ, and the corresponding columns in $\mathbf{U}$, are sorted based on the magnitude of the eigenvalues, *i.e.,* $\forall\, i, j$ if $i < j$ then $\Lambda_{i,i} \leq \Lambda_{j,j}$. As such, the graph Laplacian eigenvalue set satisfies $\{0 = \lambda_1 \leq \lambda_2 \cdots \leq \lambda_N := \lambda_{\max} \leq 2\}$, where the upper bound is guaranteed due to the use of the normalized Laplacian matrix [Chung, 1997]. This set of values define the Laplacian *spectrum* of the graph, and the eigenvector set $\{\mathbf{u}_i\}_{i=1,...,N}$ defines an orthonormal basis that spans the $\mathbb{R}^N$ space of vectors defined on the nodes of the graph. Hereon, we refer to the Laplacian eigenvectors as the *eigenmodes*, for consistency with the nomenclature used in the neuroimaging community.

Spectral decomposition of graph signals

A graph signal defined on the nodes of a graph can be represented as a vector $\mathbf{x} \in \mathbb{R}^N$, where the i-th element, $\mathbf{x}[i]$, is the signal value at the i-th node of the graph. The spectral representation of $\mathbf{x}$, commonly known as the graph Fourier transform (GFT) of $\mathbf{x}$, is given as $\tilde{\mathbf{x}} = \mathbf{U}^\top \mathbf{x}$. The signal can be perfectly recovered through the inverse GFT operation as, $\mathbf{x} = \mathbf{U}\tilde{\mathbf{x}}$, which can be expanded as $\mathbf{x} = \sum_{i=1}^{N} \tilde{\mathbf{x}}[i]\mathbf{u}_i$. As such, any given graph signal can be seen as a linear combination of the orthonormal set of Laplacian eigenvectors. In particular, $|\tilde{\mathbf{x}}[i]|^2$ gives the energy spectral density of the signal associated to the i-th eigenmode. For this work, we define our graph signals in terms of their demeaned and normalized version, obtained as

$$\mathbf{x}' = (\mathbf{x} - \mathbf{u}_1^\top \mathbf{x}\mathbf{u}_1) / ||\mathbf{x} - \mathbf{u}_1^\top \mathbf{x}\mathbf{u}_1||_2, \tag{4.2}$$

which ensures $|\tilde{\mathbf{x}}'[1]|^2 = 0$ and $\sum_{i=1}^{N} |\tilde{\mathbf{x}}'[i]|^2 = 1$.

Spatial saliency of graph Laplacian eigenvectors

In classical Fourier analysis, the complex exponentials entail a specific notion of saliency by their frequency. Similarly, the eigenvalues of a graph Laplacian carry a notion of frequency, which is directly linked to the extent of spatial saliency manifested by their corresponding eigenvectors. To understand this link, a metric known as total variation (TV) [Sandryhaila and Moura, 2014] can be computed for each eigenvector. For a given graph signal $\mathbf{x} \in \mathbb{R}^N$, the TV of $\mathbf{x}$ is defined as $TV(\mathbf{x}) = \mathbf{x}^\top \mathbf{L}\mathbf{x}$, which is a measure that quantifies the extent of variation observed in $\mathbf{x}$ relative to the underlying graph structure. Given that the eigenvectors are orthonormal, *i.e.,* $\mathbf{u}_i^\top \mathbf{u}_i = 1$, and that $\mathbf{L}\mathbf{u}_i = \lambda_i \mathbf{u}_i$, it follows that the TV of Laplacian eigenvectors reduces to

$$TV(\mathbf{u}_i) = \mathbf{u}_i^\top \mathbf{L}\mathbf{u}_i = \lambda_i, \tag{4.3}$$

showing that the variability of each Laplacian eigenvector is reflected by the associated eigenvalue, or in other words, that the Laplacian eigenvectors associated to larger eigenvalues reflect a greater extent of spatial variability.

An alternative way to quantify the degree of spatial variability of graph Laplacian eigenvectors is to compute their number of zero-crossings. In particular, we define a weighted zero-crossing measure as

$$ZC(\mathbf{u}_k) = \frac{1}{2}\sum_{i\neq j} a_{ij} H(-\mathbf{u}_k[i]\mathbf{u}_k[j]), \tag{4.4}$$

where $H(\cdot)$ is the Heaviside step function and a_{ij} is the edge weight that connects voxels v_i and v_j. The higher the zero-crossing metric, the greater the associated variability in the eigenvectors' spatial patterns.

4.2.3 Materials and Methods

Datasets

We constructed and studied brain graphs using MRI data downloaded from the publicly available HCP dataset—the 100 unrelated subjects, WU-Minn Consortium. MRI acquisition protocols of the dataset and preprocessing guidelines for diffusion MRI are extensively described elsewhere [Glasser et al., 2013]. The diffusion and anatomical data were downloaded in a readily preprocessed format. Each subject's resting-state and task fMRI scans were then realigned to their mean images, and were registered and resampled onto the diffusion data through rigid-body registration using SPM[1]. Two signal reconstruction methods were applied to the diffusion data: (1) DTI tensor fitting using FSL[2], and (2) ODF estimations using DSI Studio[3]. These steps were done to extract the associated signal representation in individual voxels across the brain, cf. sections 4.2.3 and 4.2.3.

Brain graph design

For each subject, we define a weighted brain graph characterized by a node set $\mathcal{N} = \{1,\ldots,N\}$ defined based on the set of voxels that fall within the subject's brain mask, covering gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF), representing a 3D mesh arrangement. In particular, each node i is associated to a voxel, denoted v_i, with coordinates (x_i, y_i, z_i). The graph edges are defined based on the adjacency of voxels within the Moore neighborhood cubic lattice of size $3\times3\times3$ and $5\times5\times5$, where the latter size is only used in the ODF-based design. For the $5\times5\times5$ design, voxels in the outer layer that fall in parallel to the voxels within the inner layer were excluded, enabling encoding of connections to a maximum of 98 voxels/directions in the neighborhood of each focal voxel, whereas the $3\times3\times3$ design enables

[1] https://www.fil.ion.ucl.ac.uk/spm/software/spm12/
[2] https://fsl.fmrib.ox.ac.uk
[3] http://dsi-studio.labsolver.org

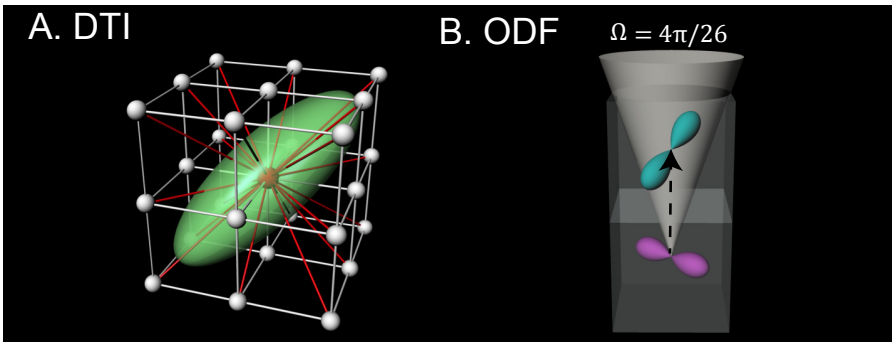

Figure 4.1 – Quantification of diffusion orientation for defining graph edge weights. (A) In the DTI-based graph, the diffusion tensor model associated to each voxel is discretized along the direction of voxels within its $3 \times 3 \times 3$ neighborhood, using a frequency formulation. (B) In the ODF-based design, the diffusion ODF model associated to each voxel is discretized by computing the area of the ODF that falls within a solid angle subtended from the center of the voxel along the direction of voxels within either the $3 \times 3 \times 3$ or $5 \times 5 \times 5$ neighborhood.

encoding connections to a maximum of 26 different directions. As such, the 5-connectivity design trades localization for better angular resolutions. With this definition of edges, each node i entails a neighborhood set of the indices of nodes in $\mathcal{N}$ that are adjacent to it, denoted $\mathcal{N}_i$, where $|\mathcal{N}_i| \leq 26$ and ≤ 98 for the 3- and 5-connectivity designs, respectively.

We define the edge weights based on a measure of *inter-voxel fiber coherence* across all pairs of adjacent voxels. In particular, we propose the use of two different signal representation models extracted from DW-MRI data: (1) diffusion tensors and (2) diffusion ODFs; we denote the resulting edge weighting schemes as the DTI-based and ODF-based methods, respectively. Let $\mathbf{r}_{ij}$ denote the vector pointing from the center of the voxel v_i to the center of the voxel v_j. Let $\{p(i,\mathbf{r}_{ij})\}_{j\in\mathcal{N}_i}$ denote estimates of the extent of diffusion at voxel v_i in directions $\{\mathbf{r}_{ij}\}_{j\in\mathcal{N}_i}$. In the following, we first present a DTI-based and an ODF-based approach to estimate $\{\{p(i,\mathbf{r}_{ij})\}_{j\in\mathcal{N}_i}\}_{i=1}^{N}$, cf. Sections 4.2.3 and 4.2.3, respectively. We then use these estimates to define a measure of *inter-voxel fiber coherence* to define as the entries of the adjacency matrix, a_{ij}, cf. Section 4.2.3.

DTI-based quantification of diffusion orientation. DTI is a model-based method for reconstructing the diffusion signal from DW-MRI. The assumption in DTI is that the diffusion pattern follows the shape of a 3D ellipsoid. The molecular displacement of water at voxel i in the direction $\mathbf{r}_{ij}$ can be approximated by a 3D Gaussian distribution with real symmetric diffusion tensor $\mathbf{T}$ as the covariance matrix:

$$P(\mathbf{r}_{ij}) = \frac{1}{\sqrt{(4\pi)^3|\mathbf{T}|}} \exp(-\frac{1}{4}\mathbf{r}_{ij}^{\top}\mathbf{T}^{-1}\mathbf{r}_{ij}), \tag{4.5}$$

where $|\mathbf{T}|$ is the determinant of the diffusion tensor. The calculation of $p(i,\mathbf{r}_{ij})$ requires a discretization step that guarantees a one-to-one mapping between the (continuous) multivariate Gaussian model and the (discrete) weighting of nodes in the brain graph. For the $3 \times 3 \times 3$

neighborhood encoding scheme, for a given voxel v_i, the set of values $\{p(i,\mathbf{r}_{ij})\}_{j\in\mathcal{N}_i}$ can be arranged into $3\times3\times3$ discrete representation, which mimics the structure of a 3D finite impulse response (FIR) filter. As such, the problem of obtaining $\{p(i,\mathbf{r}_{ij})\}_{j\in\mathcal{N}_i}$ can be alternatively seen as that of obtaining the coefficients of an FIR filter, denoted $\mathbf{h}[\cdot,\cdot,\cdot]$, characterized by it's frequency response $H(\omega)$ as

$$H(\omega)=\sum_{v=-1}^{1}\sum_{l=-1}^{1}\sum_{k=-1}^{1}\mathbf{h}[v,l,k]e^{-j\omega v}e^{-j\omega l}e^{-j\omega k}. \tag{4.6}$$

Given (4.5), we compute its Fourier transform, which is another Gaussian function. The coefficients $\mathbf{h}$ can then be obtained by a one-to-one mapping from the continuous domain to the discrete domain through the matching of corresponding frequencies. In particular, noting that $e^x=\sum_{k=0}^{\infty}x^k/k!$, both representations can be approximated as low order Taylor series, and thus, the filter coefficients are obtained by matching the coefficients of the lowest order polynomial terms. We refer to Appendix 4.2.7 for a proof.

ODF-based quantification of diffusion orientation. Unlike diffusion tensors, ODFs do not follow a specific model and shape. Thus, a one-to-one discretization of a continuous function similar to that presented for the DTI model cannot be applied. Here we build on the construction previously presented by Iturria-Medina [Iturria-Medina et al., 2007]. Within standard spherical coordinates, parametrized by $(r,\theta,\phi,)$, let $O_i(\hat{u})$ denote the ODF associated to voxel v_i with its center of coordinate being the voxel's center, with $\hat{u}(\theta,\phi)=(sin(\theta)cos(\phi),sin(\theta)sin(\phi),cos(\theta))^\top$ denoting the unit direction vector. Given $O_i(\hat{u})$, a measure of the extent of diffusion at voxel v_i along direction $\mathbf{r}_{ij}$ can be obtained as

$$p(i,\mathbf{r}_{ij})=\frac{1}{\Omega_{ij}}\int_{\Omega_{ij}}O_i^n(\hat{u})d\Omega, \tag{4.7}$$

where the exponent $n>0$ is a desired power factor to sharpen the ODF, Ω_{ij} denotes a given solid angle around $\mathbf{r}_{ij}$ and $d\Omega=sin(\theta)d\theta d\phi$ denotes the infinitesimal solid angle element; in particular, a solid angle of $4\pi/26$ and $4\pi/98$ is used for the 3- and 5-connectivity schemes, respectively. As such, $p(i,\mathbf{r}_{ij})$ gives an average measure of the surface area of the ODF within a spherical cap defined by the solid angle. Given a discrete representation of $O_i(\hat{u})$ in form of N_o samples, denoted $\{O_{i,k}\}_{k=1}^{N_o}$, along N_o spherical directions from the center of voxel v_i, (4.7) can be approximated as

$$p(i,\mathbf{r}_{ij})\approx\frac{1}{|\mathscr{D}_{ij}|}\sum_{k\in\mathscr{D}_{ij}}O_{i,k}^n, \tag{4.8}$$

where $\mathscr{D}_{ij}$ denotes a subset of direction indices $\{1,\ldots,N_o\}$ whose associated set of directions fall within Ω_{ij}. The normalization by the cardinality of $\mathscr{D}_{ij}$ is due to the difference in the number of ODF samples that fall within the solid angle subtended along the different neighborhood directions. Figure 4.1 gives a visual aid in quantifying diffusion orientation to define edge weights for the DTI design and the ODF design.

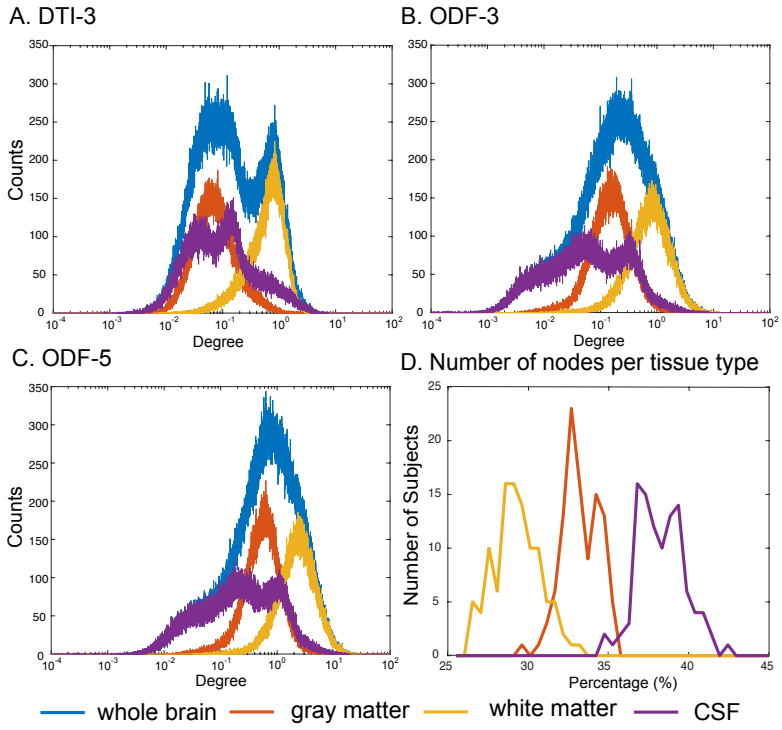

Figure 4.2 – Degree distribution of the (A) DTI-based, (B) 3-connectivity ODF-based, and (C) 5-connectivity ODF-based voxel-wise brain graphs for a representative subject. (D) The distribution of the number of nodes in each tissue type across all 100 subjects.

Brain graph edge weights. The graph edge weights a_{ij} are defined by using the estimates of diffusion orientation at the associated voxels v_i and v_j, i.e., $p(i,\mathbf{r}_{ij})$ and $p(j,\mathbf{r}_{ji})$, as well as the strength of anisotropy at those voxels. In particular, for a given voxel v_i, let $F(v_i)$ and $Q(v_i)$ denote the voxel's fractional anisotropy (FA) and quantitative anisotropy (QA), respectively. FA can be calculated directly from the eigenvalues of the diffusion tensor [Alexander et al., 2007] and QA pertains to the amount of diffusion anisotropy along the fiber orientation as originally defined by Yeh et. al. [Yeh et al., 2010]. Using these measures, we define the graph edge weights a_{ij} as

$$a_{ij} = \overbrace{\frac{P_{\text{mag}}(v_i)P_{\text{mag}}(v_j)}{\alpha^2}}^{\text{magnitude}} \overbrace{\left(\frac{p(i,\mathbf{r}_{ij})}{2\beta_i} + \frac{p(j,\mathbf{r}_{ji})}{2\beta_j}\right)}^{\text{orientation}}, \tag{4.9}$$

where $\alpha = \max_k\{P_{\text{mag}}(v_k)\}$, $\beta_k = 2\max_{l\in\mathcal{N}_k}\{p(k,\mathbf{r}_{kl})\}$ and $P_{\text{mag}}(v_i)$ denotes the magnitude of anisotropy at voxel i, defined as

$$P_{\text{mag}}(v_i) = \begin{cases} F(v_i), & \text{DTI-based design} \\ Q(v_i), & \text{ODF-based design.} \end{cases} \tag{4.10}$$

The connectivity structure of the graph is then characterized in $\mathbf{A}$, such that $a_{ij} > 0$, given as in (4.9), if nodes i and j are connected through an edge, and $a_{ij} = 0$ if otherwise.

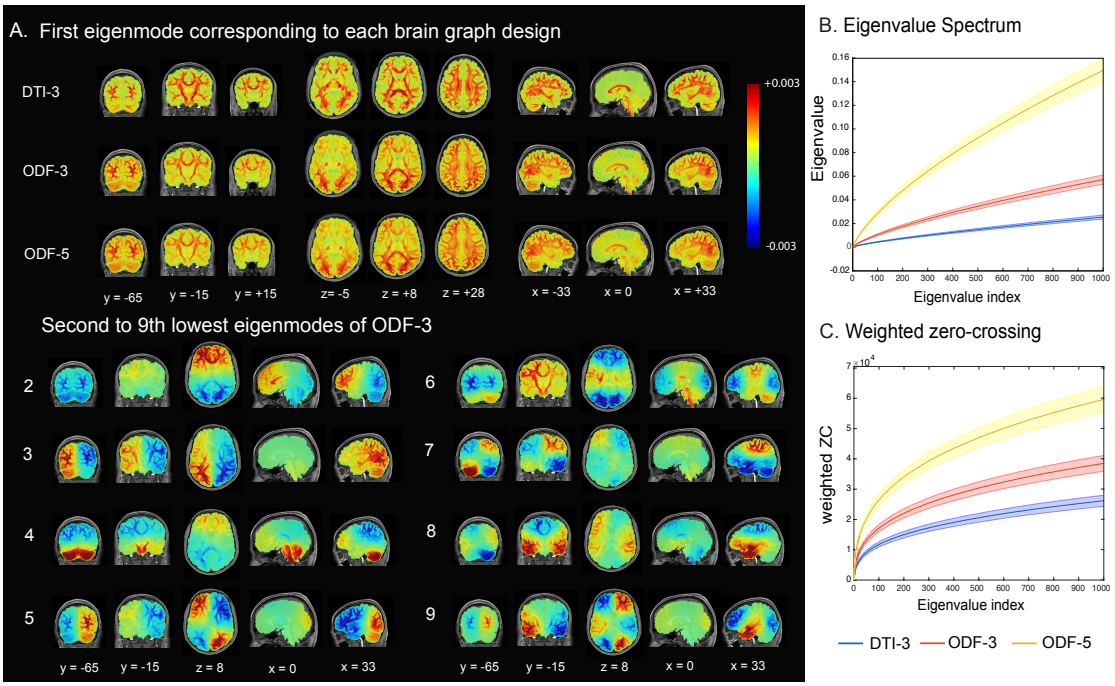

Figure 4.3 – (A) First Laplacian eigenmode of 3-connectivity DTI, 3-connectivity ODF and 5-connectivity ODF brain graphs of a representative subject. Below these are the next lowest frequency eigenmodes corresponding to the 3-connectivity ODF brain graph. (B) Lower-end spectra of DTI-3, ODF-3 and ODF-5 graphs, consisting of their first 1000 eigenvalues, and (C) weighted zero-crossing of their corresponding eigenmodes; solid lines show the mean value and shades show the standard deviation across the 100 subjects.

The *orientation* term in (4.9) gives a measure of diffusion orientation coherence between the two connected voxels, whereas the anisotropies give a *magnitude* measure that is useful in delineating tissues. For example, given two adjacent voxels that exhibit highly coherent diffusion orientations, a large weight is associated to the connection only if the two voxels also exhibit notably large anisotropies. This interplay between the orientation term and magnitude term enables, for example, to prevent associating large weight to an edge between a WM voxel and a CSF voxel. Furthermore, the normalizations incorporated in the definition, i.e., the α and β_k terms, ensure having an unbiased definition of weights relative to the structure of the diffusion tensors/ODFs across the brain, and, mathematically, they impose bounds on the orientation and magnitude terms—both terms bounded to $[0,1]$, which in turn results in a_{ij} also being bounded to $[0,1]$.

Group-level eigenmodes

The voxel-wise nature of the proposed graphs renders their size excessively large, with $7.8 \pm 0.7 \times 10^5$ nodes across the 100 subjects considered. The sheer size of the proposed graphs impedes computing the full eigendecomposition of the graph Laplacian, and, therefore, we

compute and study the first leading 1000 eigenvectors corresponding to the lowest spectral frequencies, which can be efficiently computed using the Krylov–Schur algorithm [Stewart, 2001]. To preserve subtle subject-specific spatial details, we construct all graphs in the native space of each subject's diffusion data. To enable inter-subject comparison of eigenmodes, the DARTEL normalization algorithm [Ashburner, 2007] implemented in SPM12 was used to define a group-level template coordinate space, based on the group's T1-weighted MRI data. This results in a structural T1-weighted template as well as a set of transformation maps per subject. Each subject's eigenmodes are then transformed into the template space using the subject-specific transformations, resulting in inter-subject spatially aligned eigenmodes.

Consistent inter-subject ordering of eigenmodes

Although DARTEL-normalized eigenmodes are in the same template space, the ordering of eigenmodes is not necessarily consistent across subjects, including sign ambiguity and linear combinations between modes with close eigenvalues. To obtain a consistent ordering of the eigenmodes, and enable inter-subject comparison of individual eigenmodes, we used the Procrustes transform [Goodall, 1991], which finds the optimal rotation, translation, and/or reflection between two linear subspaces. We implemented a scheme of Procrustes transformation, where the subspace is defined by the first K DARTEL-normalized eigenmodes of any M subset of subjects. Let $\mathbf{u}_{i,m}$ denote the i-th eigenmode of subject m, and let $X_{K,m} = [\mathbf{u}_{1,m}\, \mathbf{u}_{2,m}\, \cdots\, \mathbf{u}_{K,m}]$.

Algorithm 1: group-level matching of eigenmodes

$X_{avg} \leftarrow X_{K,1}$
for $j = 1$ **to** J **do**
  for $m = 1$ **to** M **do**
    $X_{K,m} \leftarrow$ Reorder columns in $X_{K,m}$ such that the resulting permuted matrix best matches X_{avg}
  end for
  $X_{avg} \leftarrow$ average $\{X_{K,1}, \ldots, X_{K,M}\}$
end for

The *reordering* is done based on Procrustes transformation estimates, and J denotes the optimal number of iterations to ensure that X_{avg} does not remain biased towards its initial value, i.e., $X_{K,1}$. The columns of the resulting $\{X_{K,m}\}_{m=1}^{M}$ are reordered in such a way that they optimally match each other and can thus be compared across the subjects.

Quantification of the extent of inter-subject structural variability as encoded in brain graphs

The precision of the Procrustes transformation can be evaluated by quantifying the cosine similarity between corresponding eigenmodes of different subjects, after DARTEL normalization and Procrustes transformation. For any pair of subjects, a symmetric cosine similarity matrix is obtained, where the deviation of the off-diagonal elements of the matrix from zero

quantifies inter-subject structural variability. If the set of eigenmodes of two subjects has been ideally matched, the cosine similarity matrix should be the identity matrix. Therefore, to quantify the mismatch between a pair of subjects m and n based on their first K eigenmodes, we define a measure of the extent of inter-subject structural variability, termed *Procrustes error*, using cosine similarity between eigenmodes as

$$E_{m,n}(K) = \frac{1}{2}\sqrt{\sum_{i=1}^{K}\sum_{\substack{j=1\\ j\neq i}}^{K}\left(\frac{\mathbf{u}_{i,m}^{\top}\mathbf{u}_{j,n}}{\|\mathbf{u}_{i,m}\|\,\|\mathbf{u}_{j,n}\|}\right)^2}. \tag{4.11}$$

In particular, we used this error term to determine J in Algorithm 1; see Figure B.1 of Appendix Bs.

To validate the performance of the Procrustes transformation, we implemented a bootstrap scheme, which successively applies the transformation on the first K eigenmodes of two randomly chosen subjects; the implementation is summarized in the following algorithm:

Algorithm 2: Procrustes validation
for $K = 100$ **to** 1000 **step** 100 **do**
  for $j = 1$ **to** 1000 **do**
    $\{m, n\} \leftarrow$ randomly select 2 values $\in \{1, \ldots, 100\}$
    $\{X_{K,m}, X_{K,n}\} \leftarrow$ **Algorithm 1** on $\{X_{K,m}, X_{K,n}\}$
    $\epsilon_j \leftarrow E_{m,n}(K)$
  end for
  $\mu_K \leftarrow$ mean of $\{\epsilon_j\}_{j=1,\ldots,1000}$
  $\sigma_K \leftarrow$ standard deviation of $\{\epsilon_j\}_{j=1,\ldots,1000}$
end for

When comparing two brain graph designs, the design that results in a higher μ_K, with reasonably small σ_K, is interpreted as capturing more subject-specific structural features.

4.2.4 Results

We constructed and studied three different brain graphs for each subject: the 3-connectivity DTI-based graph, denoted DTI-3, and the 3- and 5-connectivity ODF-based graphs, denoted ODF-3 and ODF-5, respectively. Figure 4.2 shows the degree distributions of the graphs across the different tissue types. The connectivity strength is highest in WM nodes compared to GM and CSF nodes. The degree distribution of ODF-5 graphs is a shifted version of that of the ODF-3 graphs towards a higher degree, reflecting the larger number of connections possible with the ODF-5 design. In addition, ODF-3 bears larger differences across tissue types, especially compared to DTI-3. Figure 4.2(D) shows the distribution of nodes for the different tissue types across all subjects; the median number of nodes for GM, WM, and CSF were 254 299, 221 964, and 294 028, respectively, with standard deviations 25 322, 28 579, and

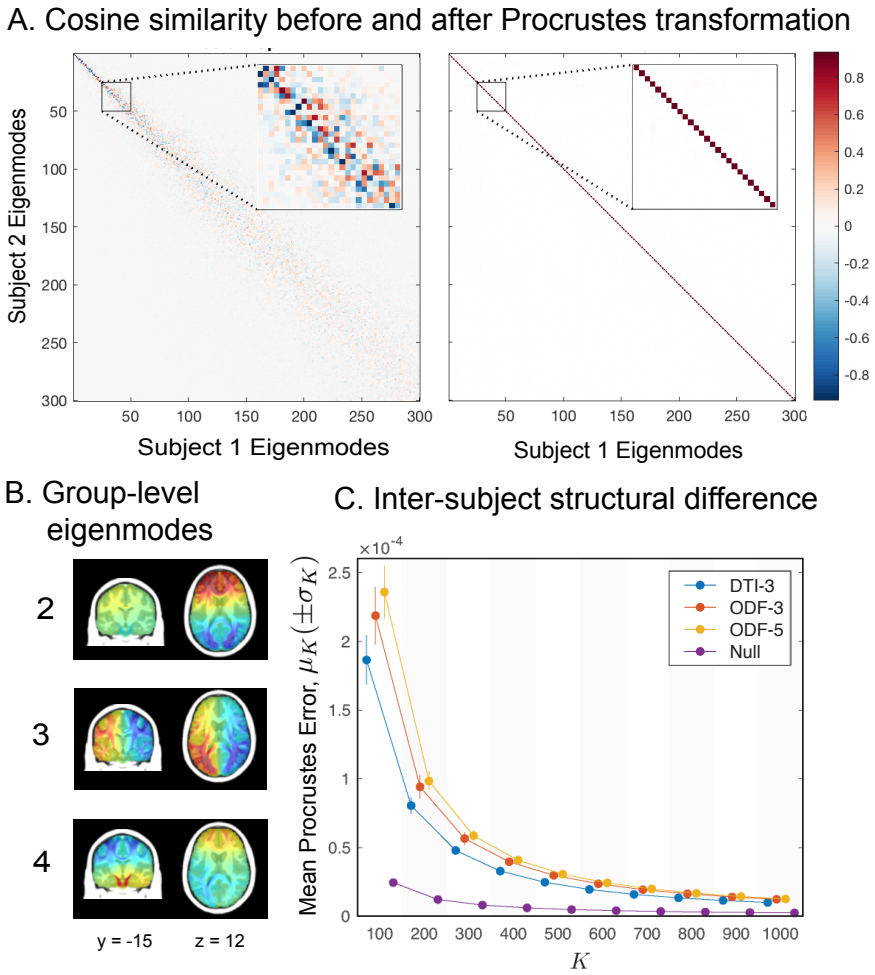

Figure 4.4 – (A) Cosine similarity measures of the first 300 eigenmodes before (left) and after (right) Procrustes transformation of two representative subjects. (B) Group-averaged eigenmodes— 2, 3, and 4 — across all transformed subjects. (C) Quantification of the extent of inter-subject structural variability captured by the eigenmodes, cf. Algorithm 2.

25 742, respectively.

Spectral Properties of the Brain Graphs

Figure 4.3(A) (top) shows the first Laplacian eigenmodes obtained from the three brain graphs of a representative subject. Noting that the first eigenmode of $\mathbf{L}$ is a function of the graph nodal degrees—$\mathbf{u}_1 = \mathbf{D}^{1/2}\mathbf{1}$ where $\mathbf{1}$ denotes the constant function that assumes the value of 1 on each node, the spatial pattern manifested by the first eigenmodes is a corroboration of the results shown in Figure 4.2(A)-(C), demonstrating that the distinction between tissue types naturally arises from the assignment of the connectivity weights in the brain graph, in particular through the assignment of the magnitude term in (4.9). The second and third eigenmodes shown in Figure 4.3(A) (bottom) manifest global morphological organization

of the brain, contrasting the posterior and anterior, and the left and right brain regions, respectively. The next eigenmodes exhibit a greater extent of spatial variability and localized information.

Figure 4.3(B) shows the lower-end graph spectra, manifesting the rate of increase in the first 1000 eigenvalues of the graphs. The ODF-5 graph has relatively larger eigenvalues than the ODF-3 and DTI-3 graphs. Given that the eigenvalues entail a notion of spatial saliency (see Sect. 4.2.2), the increase in local connectivity in the ODF-5 implies that the associated eigenmodes have greater "degrees of freedom", and as such, spatial saliency can become higher. An alternative way to verify the eigenmodes' spatial saliency is to compute their weighted zero-crossing, cf. (4.4). Figure 4.3(C) shows a trend that is consistent with that of the eigenvalues, thereby confirming the general notion that higher indexed eigenmodes encompass a larger extent of spatial saliency.

Procrustes Validation

The cosine similarity matrices of two sets of eigenmodes before and after applying Procrustes transformation are shown in Figure 4.4(A). The diagonal structure of the cosine similarity matrix associated to the transformed eigenmodes shows the effectiveness of the transformation in aligning eigenmodes of the same neuroanatomical spatial nature. The insets show traces of flipped signs and unordered eigenmodes in the original vectors, which were corrected after Procrustes transformation. The precision of the Procrustes analysis, evaluated through $E_{m,n}$ cf. 4.11, improves after several rounds of the transformation. The evolution of the Procrustes error for each successive iterations of Procrustes transformation is displayed in Figure B.1. After the first iteration of Procrustes transformation, we see a big drop in the error (about 60% change), whereas after three rounds (for $J = 3$, 94% drop of error), we obtained the final set of group-averaged eigenmodes, the first few of which are displayed in Figure 4.4(B). Major WM features that are relatively smoother than subject-specific bundles are preserved, whereas subject-specific local structures are lost.

While the Procrustes transformation maximally matches eigenmodes between subjects, it cannot account for inherent inter-subject variability. As such, the off-diagonal elements of the cosine similarity matrix are not zero even after performing several rounds of the transformation. We therefore implemented a Procrustes validation scheme that characterizes inter-subject variability of the Laplacian eigenspace as explained in Sect. 4.2.3. Furthermore, to serve as a null for comparing and validating the brain graphs, we synthesized 100 random orthonormal vectors of the same dimension as each of the subject-specific eigenmodes, applied DARTEL normalization, and then subjected the resulting vectors to Procrustes validation. In Figure 4.4(C), all three brain graphs show a decreasing trend, and they get closer to that of the null upon reaching higher K-values. Despite the differences in the Procrustes errors across the three brain graphs for low values of K, all errors eventually converge to a single point for high K. The knee points of the curves for the 3 designs are at the same $K = 300$. In addition, the ODF-based designs show higher Procrustes errors than the DTI-based design, across K.

Spectral decomposition of fMRI

To provide a proof-of-concept result on the applicability of the proposed voxel-wise brain graphs for the analysis of brain function on top of brain structure, we evaluated the extent to which fMRI data are spatially shaped by the underlying structure of the brain as encoded by the proposed graphs. In particular, we constructed fMRI graph signals through extracting voxels from within the functional volumes associated to graph nodes, from six functional tasks as well a resting-state acquisition, across 100 subjects. We then studied the energy spectral density of the extracted graph signals associated to the first 1000 spectral indices of the ODF-3 graph. To serve as a null, we also evaluated the energy spectral density of synthesized shuffled fMRI signals—obtained through random permutation of voxel indices of each fMRI volume to destroy spatial order, as well as white Gaussian noise signals. Figure 4.5(A) shows the resulting ensemble energy spectral densities. The energy spectral density of the fMRI data is characterized by a power-law behavior. The lowest frequency components of the spectrum capture the majority of the energy content of the fMRI, which is about two orders of magnitude more than what is captured by the 1000th eigenmode. This observation is more apparent by inspecting the cumulative energy spectral densities, see Figure 4.5(B), particularly in comparison to the energy profiles of shuffled fMRI data and Gaussian noise. Functional brain activity is expressed preferentially by lower-frequency components; approximately up to 45% of the energy content is captured by the first 1000 eigenvectors, wherein the contributions from the first eigenmode are equal to zero due to that the data were demeaned, cf. (4.2).[4] Despite the large size of the voxel-wise brain graph, the energy spectral density of fMRI signals shows clear evidence that functional patterns are smooth over the structural connectome, consistent with previous findings in voxel-resolution gray matter graphs [Behjat et al., 2015, Behjat and Larsson, 2020] and region-wise connectomes [Preti and Van De Ville, 2019, Atasoy et al., 2016].

4.2.5 Discussion and Perspectives

The proposed brain graphs encode subject-specific structural connections at the resolution of voxels. The Laplacian eigenmodes of both DTI- and ODF-based designs manifest major WM bundles as well as inter-gyri connections. Low-frequency eigenmodes display geometrical information of the head shape, demonstrating low spatial variability in spatial features. Eigenmodes corresponding to higher eigenvalue indices show a greater extent of spatial variability. These findings were supported by the eigenvalue spectra and the observed weighted zero-crossings of the eigenmodes, confirming the general notion that GFT-based decompositions allow a similar interpretation of variability across regions of the brain.

We evaluated each brain graph design by characterizing the inherent inter-subject variability encoded in the graphs, through the Procrustes transform, which finds the optimal configura-

[4] Approximately 70% of the total signal energies were captured by the first eigenmodes, which were removed in the demeaning step, and as such, the total amount of energy captured by the first 1000 eigenmodes amounts to approximately 85% of total signal energies, i.e., $70 + 0.45 \times 30$.

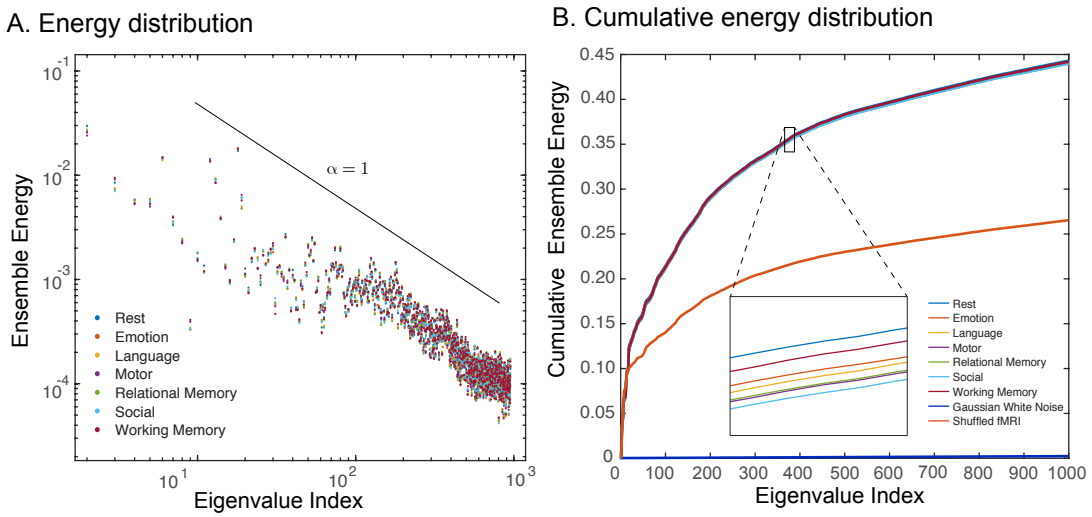

Figure 4.5 – Ensemble energy spectral density of fMRI data with respect to the first 1000 eigenmodes of ODF-3 brain graph: (A) ensemble energies associated to individual eigenmodes, (B) cumulative ensemble energies.

tion that matches single subject eigenmodes to an averaged set. This step, however, cannot account for subject-specific WM tracts encoded in the eigenmodes, as manifested by the cosine similarity analysis of pairs of subjects. The ODF-based designs captured higher Procrustes errors than the DTI-based design. The ODF-5 design slightly outperformed the ODF-3 design, suggesting the benefit of using the larger neighborhood in encoding fiber orientations with better angular resolution.

An important assumption of the Procrustes validation analysis is the presumed consistent precision and accuracy of spatial normalization using DARTEL across subjects and graph designs. The normalization step in itself aims to maximize the degree of spatial agreement across subjects through matching individual subjects' data to the group-level template. Maximizing image registration's precision therefore requires minimizing the amount of residual spatial variability across individuals. All eigenmodes were registered using the same subject-specific deformation, and as such the performance of DARTEL can be assumed consistent across the three designs. The residual spatial variability is then carried on to the group-level matching using Procrustes transform, and are subsequently accounted through the Procrustes errors.

Although the use of voxel-wise brain graphs comes with a high dimensionality problem, we show through the energy spectral density of rest and task fMRI data (Figure 4.5) that human brain activity can be well approximated by a small subset of low-frequency components of the graph Laplacian spectra. This observation is, on the one hand, consistent with the energy profile of fMRI graph signals on voxel-wise gray matter graphs [Behjat et al., 2015, Behjat and Larsson, 2020], as well as region-wise brain graphs [Preti and Van De Ville, 2019], which is reminiscent of a power-law behavior, and on the other hand, can be linked to the decreasing trend observed in the Procrustes validation errors (see Figure 4.4(D)), where increasing K-

values reduce the error close to that of a randomly generated graph. This is interesting in light of the increasing evidence of scale-free behavior in human brain activity [Van De Ville et al., 2010, He et al., 2010, Ciuciu et al., 2012] observed in both temporal and spatial scales [Maxim et al., 2005]. The practical implication of such energy profiles is that the lower-spectral-end energy content of fMRI data on the proposed voxel-wise graphs can provide meaningful information, which significantly reduces the computational complexity associated to diagonalization of the graph Laplacian. Alternatively, to study the energy profile across the spectrum, graph filter design schemes that adapt to the ensemble signal content can be used to efficiently filter the data [Behjat et al., 2016], which can be made computationally efficient through leveraging a polynomial approximation scheme to estimate the ensemble energy content [Behjat and Van De Ville, 2019], thus, obviating the very need to diagonalize the Laplacian matrices.

The ODF-based design is an extension of previous work in [Iturria-Medina et al., 2007], modified in three ways. First, by using the QA measure, we amplify the strength of connection between adjacent nodes that exhibit high anisotropy. Instead of using quantitative anisotropy (QA) derived from the DW-MRI data, the design in [Iturria-Medina et al., 2007] uses probabilistic tissue maps to delineate the magnitude of each node among tissue types; similar to the use of QA, we use FA for the DTI design. This modification in the design has been inspired by diffusion tractography frameworks [Conturo et al., 1999, Staempfli et al., 2006], where FA maps are used as a mask to constrain fiber trajectories. Second, we sharpen the ODFs to boost the fiber coherence between aligned anisotropies, given the limited degree to which diffusion ODFs can differentiate fiber orientations [Jones et al., 2013]. Third, we include the entire set of voxels within the brain mask, rather than just GM and WM. Although including CSF voxels entails a greater computational complexity, the inclusion was due to several reasons. By including all voxels within the brain mask, the delineation of the tissue types is naturally reflected by the diffusion information encoded in the graph edge weights, as verified, particularly, by the spatial patterns manifested by the first eigenmodes, see Figure 4.3(A). Moreover, this modification bypasses the need for tissue segmentation, which is not convenient within clinical studies including patients with structural brain deficits, for which normalization of segmented tissue maps is also challenging [Glass et al., 2008]. Nevertheless, if high-resolution structural data are available, it is possible to exclude the CSF, through pruning edges that connect voxels that lie on opposite sides of the brain pial surface [Behjat and Larsson, 2020] and through morphological operations to remove the ventricles. Furthermore, in terms of the encoding of structural data, within-voxel axonal orientation is currently assumed to be symmetric, thus, a potential improvement of the brain graph design could be an extension to the use of asymmetric fiber orientation estimates [Bastiani et al., 2017].

Voxel-wise brain graphs hold the potential to open new possibilities for the study of the brain. The eigenvalues of voxel-wise brain graphs encoding cortical morphology have been found to provide distinguishable features of gender and hemispheric asymmetry [Maghsadhagh et al., 2019]. Similarly, the extent of variability observed in the lower-end spectra of the proposed graphs across subjects, particularly for the ODF-based design—see Figure 4.3(B), suggests their

potential for deriving markers to distinguish population groups and individuals. The proposed brain graphs may also prove beneficial in deriving novel structural connectivity measures that account for both direct and indirect pathways, as for example proposed by a conductance model in [Frau-Pascual et al., 2019], through using shortest path algorithms from graph theory. Moreover, identifying individuals using patterns of functional brain connectivity [Finn et al., 2015, Amico and Goñi, 2018]—a.k.a. connectome fingerprinting—has been recently generalized to include structural connectomes [Mars et al., 2018]. Information embedded in structural connectomes can be explored not only for individual fingerprinting, but also for studying neuroplasticity in longitudinal data [Yeh et al., 2016].

In relating brain structure to function, in a similar manner that low-frequency eigenmodes show strong association to low-sensory functions in region-wise integration of structure and function, the recruitment of low-frequency components in the voxel-wise Laplacian can also be associated to long-distance interactions in the brain. Functional connectivity has often been associated to Euclidean distance [Achard et al., 2006, Ercsey-Ravasz et al., 2013, Alexander-Bloch et al., 2013], whereas using the eigenmodes, distance can be better interpreted in terms of the spectral representations of functional signals. Functional variations that are captured using low-frequency eigenmodes are smooth with respect to long-distance WM bundles, whereas localized and short-distance functional associations are expected to be dominated by higher frequency components. GSP measures such as the energy spectral density and total variations highlight these differences and can be straightforwardly applied to distinguish spatial variations in fMRI data across different functional tasks [Behjat and Larsson, 2020]. The voxel-to-voxel ODF-based encoding proposed in this work has been exploited for inpainting WM fMRI from cortical fMRI data [Tarun et al., 2020b] through graph signal recovery [Chen et al., 2015a], revealing the mediation of WM structures across temporally-coherent gray matter areas. Furthermore, preliminary results have shown the effectiveness of using the voxel-wise diffusion informed graphs in enhancing fMRI activation mapping in WM [Abramian et al., 2020], in-line with previous findings on voxel-wise gray matter graphs [Behjat et al., 2015]. Finally, a potential application of the proposed voxel-wise graphs is in assessing the extent to which the brain structural-functional relation holds at a spatially finer mesoscale using the concept of graph Slepians [Van De Ville et al., 2017, Bolton et al., 2018, Petrovic et al., 2019], which are linearly recombined sets of eigenmodes that are constrained to be localized in a subset of nodes. This allows to put focus the representation on particular nodes given specific hypotheses about brain function, at a higher level of resolution than conventional region-wise applications.

4.2.6 Conclusion

We proposed two methods for constructing voxel-wise brain graphs from DW-MRI data. Through a Procrustes validation scheme that reflects inter-subject structural differences, it was shown that low-frequency eigenmodes reflect the highest amount of structural information from DW-MRI. This finding was corroborated by the manifested energy spectral density of

functional signals showing the preferential expression of human brain activity onto lower frequency components. Potential applications of the proposed brain graph were discussed, particularly in relation to integrating brain structure and function through applications of GSP to functional brain imaging.

4.2.7 Proof: Cartesian discretization of a Gaussian using its Fourier domain representation

For simplicity of notation, we present the discretization process for a 1D Gaussian, which can be readily extended to higher dimensions. The 1D Gaussian is given as

$$G(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}x^2} \tag{4.12}$$

where σ^2 is a constant which describes the variance of the curve and $\frac{1}{\sqrt{2\pi\sigma^2}}$ is a normalization factor. Its Fourier transform is given by:

$$G(\omega) = e^{-2\pi^2\sigma^2\omega^2}, \tag{4.13}$$

Obtaining a 3-sample discretization of the Gaussian using it's Fourier domain representation is a problem that can be equivalently seen as finding the coefficients of a 3-tap FIR filter, $\mathbf{h}$, matching the Frequency response given in (4.13), with its frequency response $\mathscr{H}(\omega)$ given as

$$H(\omega) = \sum_{n=-1}^{n=1} \mathbf{h}[n] e^{-j\omega n} = \mathbf{h}[-1] e^{j\omega} + \mathbf{h}[0] + \mathbf{h}[1] e^{-j\omega}. \tag{4.14}$$

Noting that $e^x = \sum_{k=0}^{\infty} x^k / k!$, Taylor series expansion of (4.13) and (4.14) up to the second order gives

$$G(\omega) \approx 1 - 2\pi^2\sigma^2\omega^2 + \frac{1}{2}(-2\pi^2\sigma^2\omega^2)^2 \tag{4.15}$$

and

$$\begin{aligned} H(\omega) \approx \mathbf{h}[-1] + \mathbf{h}[0] + \mathbf{h}[1] + (\mathbf{h}[-1] - \mathbf{h}[1])\, j\omega \\ -\frac{1}{2}(\mathbf{h}[-1] + \mathbf{h}[1])\omega^2, \end{aligned} \tag{4.16}$$

respectively. Matching the coefficients of the polynomials in (4.15) and (4.16), a system of linear equations is obtained as

$$\begin{aligned} 1 &= \mathbf{h}[-1] + \mathbf{h}[0] + \mathbf{h}[1] \\ 0 &= \mathbf{h}[-1] - \mathbf{h}[1] \\ 2\pi^2\sigma^2 &= \tfrac{1}{2}(\mathbf{h}[-1] + \mathbf{h}[1]), \end{aligned}$$

where the solution gives the desired FIR impulse response as

$$\mathbf{h} = [2\pi^2\sigma^2, 1 - 4\pi^2\sigma^2, 2\pi^2\sigma^2]. \tag{4.17}$$

In the case of the 3D multivariate Gaussian density, the exponent in (4.12) becomes $-\frac{1}{2}\mathbf{x}^\top\Sigma^{-1}\mathbf{x}$—a quadratic form in the 3D vector variable $\mathbf{x}$ covering $x-$, $y-$, and $z-$directions—and σ^2 is replaced by a 3×3 invertible covariance matrix Σ that is equivalent to the tensor $\mathbf{T}$ in (4.5). The goal is then to find the coefficients of a $3\times 3\times 3$ FIR filter by matching its frequency response $\mathscr{H}(\omega)$, as given in (4.6), to the Fourier transform of the 3D Gaussian, in a similar way as done for the 1D case above.

5 Recovery of whole-brain structure-function networks

5.1 General Motivation

In Chapter 3, we provided a comparison between structural and functional connectivity measures, across different brain states and different tasks, by directly computing the correlation of SC and FC at the region-wise level. Meanwhile, in Chapter 4, we introduced a new technique that extends region-wise characterization of brain structure to voxel-level analysis. We also demonstrated the potential of the proposed brain graph to be used in GSP analysis by decomposing resting-state and task fMRI data in terms of the lowest frequency components of the brain graph's Laplacian eigenspectra. The energy spectral densities gave a direct measure of how much information in fMRI is captured by the anatomical data from DW-MRI. However, the exact manner in which dynamically evolving functional brain networks observed during rest and sleep (see Chapter 3) are supported by the underlying brain anatomy remains elusive. *Which white matter tracks are implicated in the formation of resting-state networks?*

In this chapter, we finally address this question by providing an extended picture of human brain activity, not only within the GM but also in the WM. Leveraging on the voxel-level brain graph introduced in Chapter 4, we interpolate GM signals onto the WM using the structural information encoded in the brain grid to guide the process. This results in a 4D functional volume that includes interpolated WM signals, which can be analyzed using conventional static and dynamic FC approaches reviewed in section 2.2.2. We focus the analysis into unraveling the structure-function co-activating patterns related to the default mode network (DMN), a set of brain regions known to be predominant during rest [Greicius et al., 2003]. The results allow us to observe various combinations of WM tracks that mediate the interactions of these functionally connected brain areas. In particular, we demonstrate how different brain areas and associated WM tracks corresponding to the default mode subsystems are recruited in different phases of working memory and relational task paradigms. At the end of the chapter, we enumerate two main research avenues for future structure-function studies based on the proposed technique.

5.2 Journal Article: Structural mediation of human brain activity revealed by white-matter interpolation of fMRI

A. Tarun, H. Behjat, T.A.W. Bolton, D. Abramian, and D. Van De Ville
Post-print version of the article published in NeuroImage 2020, vol. 213, https://doi.org/10.1016/j.neuroimage.2020.116718

Author Contributions

A. Tarun conceptualized research and performed all stages of the analysis, from methodology, validation, formal analysis, writing of the original draft and editing of the manuscript, **H. Behjat** and **T.A.W. Bolton** edited and reviewed the manuscript, **D. Abramian** contributed in the preliminary version of the voxel-level brain graph used, and **D. Van De Ville** conceptualized research, offered continuous guidance, edited and reviewed the manuscript.

Abstract

Understanding how the anatomy of the human brain constrains and influences the formation of large-scale functional networks remains a fundamental question in neuroscience. Here, given measured brain activity in gray matter, we interpolate these functional signals into the white matter on a structurally-informed high-resolution voxel-level brain grid. The interpolated volumes reflect the underlying anatomical information, revealing white matter structures that mediate the interaction between temporally coherent gray matter regions. Functional connectivity analyses of the interpolated volumes reveal an enriched picture of the default mode network (DMN) and its subcomponents, including the different white matter bundles that are implicated in their formation, thus extending currently known spatial patterns that are limited within the gray matter only. These subcomponents have distinct structure-function patterns, each of which are differentially observed during tasks, demonstrating plausible structural mechanisms for functional switching between task-positive and -negative components. This work opens new avenues for the integration of brain structure and function, and demonstrates the collective mediation of white matter pathways across short and long-distance functional connections.

5.2.1 Introduction

The coordination of distant neuronal populations gives rise to a vast repertoire of functional networks that underpin human brain function. Using functional magnetic resonance imaging (fMRI), temporally coherent activity can be investigated using measures of functional connectivity (FC) [Friston, 1994]. On the other hand, inter-regional communication is mediated by the anatomical scaffold which can be conveniently summarized by structural connectivity (SC) extracted from diffusion-weighted MRI (DW-MRI) [Bullmore and Sporns, 2009]. Over the past decade, a particular focus has been dedicated to unraveling how the human brain maintains its vast repertoire of FC states despite being constrained by a rigid anatomical substrate. To bridge the gap between SC and FC, several methods have been proposed. The simplest ones are the seminal works that directly probed the statistical interdependence (*e.g.,* Pearson correlation) between separately defined SC and FC measures [Andrews-Hanna et al., 2007, Hagmann et al., 2008, Honey et al., 2009, Supekar et al., 2010, Horn et al., 2014]. Limited by the bivariate and summarizing nature of this analysis, the effect is capturing only a general trend of correlation. In order to directly capture white matter structures that connect temporally coherent areas, several studies initially compute FC and specify implicated regions of interests (ROI) as a seed for estimating fiber streamlines [Greicius et al., 2009, van den Heuvel et al., 2009, Chamberland et al., 2015]. Most of these works were specifically directed to the analysis of the default mode network (DMN), a set of brain regions that are known to be more engaged during rest [Greicius et al., 2003]. In contrast to extracting SC from FC priors, a number of studies have attempted to reproduce brain activity from predefined structural connectomes through numerical simulations [Galán, 2008, Honey et al., 2007, Deco et al., 2011, Deco et al., 2012, Deco et al., 2013, Goni et al., 2014].

Studies that extract SC from FC or *vice versa* are mostly hypothesis-driven and entail many explicit assumptions (see [Damoiseaux and Greicius, 2009, Zhu et al., 2014] for early reviews). To understand how distributed patterns of functional activity arise from a fixed underlying structure, a need for research methodologies that are observer-independent and data-driven is of utmost importance. A common approach for data-driven methods is based on blind-decomposition techniques. By combining diffusion anisotropy (*e.g.,* fractional anisotropy, axial and radial diffusivities) and classical FC to build a joint SC-FC measure [Sui et al., 2013, Sui et al., 2015], structural and functional alterations in healthy and clinical populations are extracted using multimodal canonical correlation analysis and joint independent component analysis (ICA) [Amico and Goñi, 2018]. On the same line, using ICA, concurrent white matter (WM) bundles and gray matter (GM) networks from a tractography-based data matrix have been proposed [O'Muircheartaigh and Jbabdi, 2018]. Also using tractography, Calamante and colleagues [Calamante et al., 2017] introduced the concept of track-weighted dynamic FC (TW-dFC), where as a first step, WM streamlines from two structurally connected voxel endpoints are obtained. Then, these streamlines are weighted according to the FC at their endpoints. To render the approach dynamic, FC measures were computed over a sliding-window whose length is arbitrarily chosen. TW-dFC then produces a set of four-dimensional volumes showing the averaged FC between corresponding streamline endpoints onto the WM at the temporal

resolution of the window used. While several studies have made considerable progress in linking structure and function, not only with fMRI but also with other imaging modalities such as electroencephalography (EEG) and magnetoencephalography (MEG) [Deslauriers-Gauthier et al., 2019, Fukushima et al., 2015], none of these methods jointly and simultaneously integrates functional and diffusion data into a *single* integrated framework, which allows the natural emergence of *full-brain* spatial patterns that cover both the WM and the GM.

Lately, to further transcend our current understanding of structure-function relationships, attention has been set on studying SC and FC through the lens of more technical frameworks borrowed from other research fields, such as propagator-based methods [Robinson, 2012, Robinson et al., 2016, Atasoy et al., 2016], and control network theoretical tools [Gu et al., 2015]. Graph signal processing (GSP) for neuroimaging is another emerging field [Huang et al., 2018a], where initially, a graph is defined by identifying regions in GM as nodes, and the strength of their SC through WM as edge weights (*e.g.,* number of streamlines connecting GM regions). Functional data are then defined atop the graph and are interpreted as time-dependent graph signals, on which connectome-informed signal processing operations can be performed. This framework incorporates connectivity through the WM, but only as SC between pairs of GM regions (*i.e.,* WM regions are not explicitly defined as nodes).

In this work, we advance the GSP concept by extending the existing approach from region-wise analysis to a high-resolution (*i.e.,* 850,000 voxels) voxel-level perspective. This translates to modeling the WM explicitly as nodes of the brain graph for which local connectivity is known from DW-MRI. We then define the blood–oxygenation level-dependent (BOLD) time-series from resting-state fMRI as dynamic signals residing on the nodes of the brain graph. By generalizing principles of classical signal processing in regular domains [Buades et al., 2010, Rudin et al., 1992, Chambolle, 2004, Candès et al., 2006] to irregular graphs [Shuman et al., 2013, Chen et al., 2015b], we interpolate functional signals, measured on the GM, into the WM, using a whole-brain voxel-wise connectome to guide the process. This allows relating measures of brain activity on the GM with their mediation through WM; *i.e.,* how particular patterns of brain activity jointly rely on an ensemble of WM pathways. We then apply the WM interpolation on all volumes of resting-state data, thereby generating a set of 4D-volumes on which conventional static and dynamic FC tools can be applied. Using the posterior cingulate cortex (PCC), a well-known subcomponent of the DMN, as a seed region, we retrieved new structure-function networks that extend the well-known GM-limited DMN pattern to include the WM. The functional relevance of these structure-function networks is further validated through investigations of their relation to task fMRI. Our results illustrate the structural mechanism for the dynamic switching between task-positive and -negative subsystems of the DMN in different phases of working memory and relational task paradigms.

5.2.2 Materials and Methods

Data

We used a total of 51 subjects obtained from the publicly available Human Connectome Project dataset (HCP 1200 release), WU-Minn Consortium. Preprocessed diffusion MRI, original functional MRI, and anatomical data are downloaded following a random subject selection scheme while maintaining a uniform demographic distribution (male and female, ages 22-50). MRI acquisition protocols of the HCP and preprocessing guidelines for diffusion MRI are fully described and discussed elsewhere [Glasser et al., 2013]. On the other hand, fMRI data were acquired using Gradient-echo EPI sequence (repetition time [TR] /echo time [TE]/flip angle = 720 ms/ 33.1 ms/ 52 deg). Following the standard preprocessing protocol [Van Dijk et al., 2010], four sessions of resting-state fMRI scans (1200 volumes each session, giving a total of 4800 volumes per subject) were realigned to their mean images. The results were then registered and resampled onto the diffusion data using rigid-body registration (SPM12, https://www.fil.ion.ucl.ac.uk). The registered images were de-trended (*i.e.,* constant, linear) and smoothed using a Gaussian kernel (FWHM = 6mm). The first 10 frames were removed to achieve steady-state magnetization of the fMRI data. Meanwhile, the diffusion data were processed using DSI Studio (http://dsi- studio.labsolver.org) to estimate the orientation density functions (ODFs) associated with individual voxels across the brain. T1 images were downloaded from the HCP in a readily aligned format corresponding to the diffusion subject space. Individual tissue maps (WM, GM, and cerebrospinal fluid (CSF)) were then segmented from the T1 image using SPM12. The GM probability maps outputted from the segmentation algorithm were thresholded at 0.3 and were used to mask GM voxels.

Building the voxel-wise brain grid

Similar to classical brain graphs, we define our grid as $\mathcal{G} := (\mathcal{V}, \mathbf{A})$, where $\mathcal{V} = \{1, 2, 3, ..., N\}$ is the set of N nodes representing the brain voxels (*i.e.,* $N = 700 - 900$ thousand), and $\mathbf{A} \in N \times N$ is an adjacency matrix that encodes the likelihood of water molecules to diffuse from their current position to neighboring voxels. We used an ODF-based weighting-scheme on a three-dimensional 26-neighborhood mesh. ODF is an empirical distribution of water diffusion at different orientations introduced to solve fiber-crossing confounds that are typically observed in model-dependent reconstruction techniques from DW-MRI [Tuch et al., 2003]. For more details regarding the choice of the signal reconstruction technique and the design of the brain graph, we refer to Chapter 4 and [Tarun et al., 2019]. The topology of the brain grid is then validated by defining the graph Laplacian matrix in its symmetric normalized form ($\mathbf{L} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}}$), whose eigendecomposition leads to a complete set of orthonormal eigenvectors that span the graph spectral domain with their corresponding real, non-negative eigenvalues. The top eigenfunctions corresponding to the lowest spatial variation are provided in Figure 4.3.

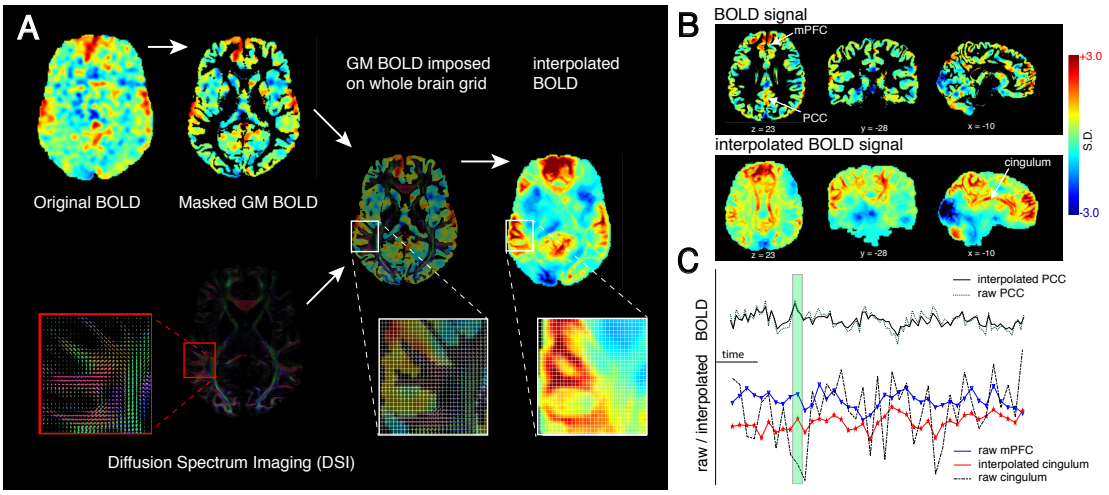

Figure 5.1 – **Workflow of the graph signal recovery framework.** (A) GM BOLD signals are extracted from fMRI volumes, one signal per time instance, through masking the volume by the thresholded probabilistic GM map (threshold=0.3). ODFs associated to all voxels are extracted from the diffusion MRI data. The ODFs are then embedded on a three-dimensional, 26-neighborhood connectivity grid, forming a probabilistic connectome at voxel-level resolution. A signal value can be associated to each voxel on the grid. To initiate signal recovery at each time instance, the associated GM BOLD signal is mapped to the GM voxels within the grid, and the remaining voxels are set to zero. The interpolated BOLD volume is obtained through minimizing a cost function that balances between (1) retaining the GM signal fixed and (2) obtaining a smooth signal over the entire brain grid relative to the underlying structure. (B) Axial, sagittal, and coronal views of a single slice of a BOLD volume compared to its corresponding interpolated BOLD volume. (C) Original and interpolated BOLD signal time-series of PCC, mPFC and cingulum; the selected time frame in (B) corresponds to the highlighted time point.

Graph signal recovery formulation

The overall goal of this work is to be able to interpolate signals from the GM into the WM based on the structural information encoded in the voxel-wise connectome. To do this, we consider a cost function that includes a least-squares data-fitting term equal to the residual sum of squares (RSS), and an L2 regularization term that takes into account smoothness with respect to the brain grid, given by

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{B}\mathbf{x}\|^2 + \lambda \mathbf{x}^T \mathbf{L}\mathbf{x}, \tag{5.1}$$

where $\mathbf{y}$ is a vector of length N containing initial BOLD values within GM nodes and zero elsewhere, $\mathbf{B}$ is an indicator matrix that selects the GM nodes, and λ is the regularization parameter. $\mathbf{L}$ is the symmetric normalized graph Laplacian, which encodes the voxel-level anatomical structure. The cost function balances between (1) minimizing the RSS and retaining the original GM signals, and (2) imposing smoothness with respect to the structure of the graph. The balance between the two is dictated by the choice of regularization parameter λ. The parameter λ was selected by evaluating the cost of the regularization term $\mathbf{x}^T\mathbf{L}\mathbf{x}$ across

a sufficiently large range of λs, specifically, $\lambda = \{1,2,3,...,60\}$. This procedure allowed for a subject-dependent and data-driven choice of the regularization parameter. Figure S3 shows the plot of the evaluation for all 51 subjects considered in this study; the optimal λ were 9(10), 10(39), 11(2), where the values in the parentheses indicate their frequency of occurrence. The solution to Equation 5.1 gives the interpolated volume for one time-point in the fMRI. The interpolation is then done repeatedly at each time instance (*i.e.,* each volume) in the functional data.

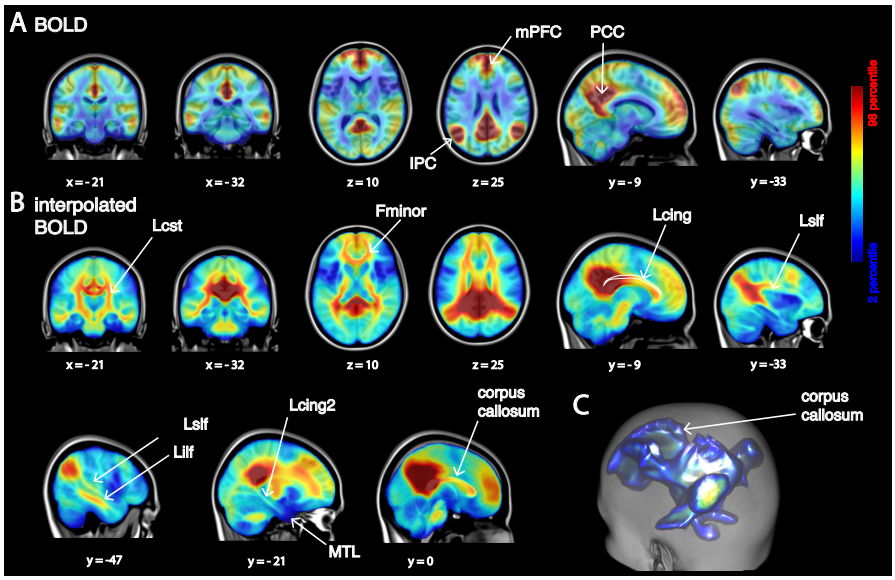

Figure 5.2 – **PCC-seed correlation** for the (A) original BOLD volumes and the (B) interpolated volumes, averaged across all subjects. WM bundles connecting temporally-coherent GM regions are uncovered from interpolated volumes. (C) 3D view of the observed WM structure. PCC - posterior cingulate cortex, mPFC - medial prefrontal cortex, Lcst - left cortico-spinal tract, Fminor - forceps minor, Lcing - cingulum, Lslf - left superior longitudinal fasciculus, Lilf - Left inferior longitudinal fasciculus, Lcing2 - hippocampal cingulum.

Static functional connectivity analysis

The output of the graph signal recovery problem is a set of 4D whole-brain functional volumes at a frame-wise resolution. We used the classical method of static FC analysis (*i.e.,* Pearson correlation) in the interpolated volumes, wherein the averaged BOLD time-course of an ROI is correlated to the activity profiles of all brain voxels to locate temporally coherent areas. We used the PCC, a core region of the DMN, as a seed to observe the joint structural-functional patterns associated with it. This is done to each of the four sessions of resting-state (RS) fMRI sequences per subject, and the corresponding Fisher Z-transformed correlations were averaged across the group.

Temporal decomposition of DMN-related co-activation patterns

Next, we extended the analysis from static to dynamic FC perspective, using the so-called co-activation pattern (CAP) analysis [Liu and Duyn, 2013]. We used the PCC to extract time-varying structural mediation of functional cross-talks implicated with the DMN. To obtain the CAPs, we performed a two-step temporal decomposition: (1) subject-level clustering and (2) group-level matching. For the first part, the top 15% significant frames corresponding to the PCC are extracted in each of the four sessions of resting-state scan from the HCP. Significant volumes that survived the thresholding are concatenated together in each subject. The choice of the threshold is motivated by finding the optimal level of similarity between averaged significant PCC volumes at different thresholds to that of the PCC-seed connectivity outcome (see Figure C.4(A)). These significant volumes are then decomposed into multiple co-activation maps by running the k-means clustering algorithm which groups and averages together those with spatially similar frames. The optimal number of clusters was determined using consensus clustering, a resampling based procedure for optimal class discovery [Monti et al., 2003] (see Figure C.4(B)). We used cosine as a distance metric and performed 10-folds replicates to obtain the final clustering at the individual level. After obtaining the co-activation maps for each subject, we performed the group-level matching by first transforming all individual maps into a common template (MNI), followed by a one-to-one matching and averaging of the correspondingly similar spatial patterns across subjects.

When doing the clustering at the subject-level, the assignment of CAP indices in one subject does not necessarily correspond to the cluster ordering of another subject. Therefore, when doing the group-level averaging, it is necessary to perform a one-to-one CAP matching and index re-labeling. The one-to-one group-level CAP matching is done using the Hungarian algorithm [Kuhn,], which solves for the optimal pairing that has the most similar spatial combinations as defined by their cosine similarity. The first matching is done on the first two subjects, matching spatially similar CAPs and re-labeling their respective frame indices. Then, the average of the first two is computed, producing a temporary group level CAP from which the third subject is paired with. This process continues until we reach the last pair of subjects to match. The whole process is repeated again, but this time, the first subject is realigned with the overall average obtained from the first round of group-level CAP matching. Two rounds of subject-to-subject Hungarian matching were done to remove bias on the effect of the first subject.

5.2.3 Results

Recovery of activity in white matter that is compatible with structural connectivity

Signal recovery is a classical task that relates to image denoising [Buades et al., 2010], signal inpainting [Rudin et al., 1992, Chambolle, 2004], and sensing [Candès et al., 2006]. Here, we embed the three-dimensional voxel-level grid into a graph where the nodes are voxels, and the edges encode the voxel-wise strength of structural connectivity measured by DW-MRI. Each

node has maximally 26 edges and is connected to its closest neighboring voxels. Functional volumes are preprocessed and upsampled to the same resolution as this structural grid. The values from the functional volumes are assigned to each node in the GM, while the remaining nodes, in particular, those in the WM, remain unassigned. The values of the unassigned nodes are then recovered by solving an optimization problem that relies on two assumptions: (i) values in the GM nodes should strongly influence the data fitness (and thus remain reasonably unchanged); (ii) whole-brain signals should maintain smoothness according to the graph structure. The assumption of smoothness with respect to the graph ensures that signals are interpolated according to the brain's structural scaffold. It is important to note, however, that the term *interpolation* does not denote newly generated information, but rather a resulting pattern retrieved by incorporating GM measures with the WM structure.

The overall pipeline is illustrated in Figure 5.1(A). For an illustrative frame of resting-state fMRI, we show in Figure 5.1(B) an example of the original and the interpolated volumes. Visually, we observe high activity in the PCC and medial prefrontal cortex (mPFC) in the original BOLD signal. After performing the WM interpolation, we recovered additional patterns, such as the cingulum bundle (Lcing) that mediate PCC and mPFC. In addition to major WM bundles that are captured, we also find short-range WM connectivity within the cortical foldings of the GM [Koch et al., 2002]. For a quantitative comparison, we used a probabilistic WM atlas provided by Zhang and Arfanakis [Zhang and Arfanakis, 2018] to compute the average signal within the cingulum bundle (R/Lcing), and plotted its overall trend throughout the whole timecourse together with the mean signals in the PCC and mPFC. Figure 5.1(C) shows that the original PCC-averaged signal bears close similarity (cosine = 0.95) to its interpolated counterpart, a direct consequence of the first constraint imposed in the signal recovery formulation (*i.e.*, retaining the signals within the GM nodes fixed). In contrast, we observe comparatively lower similarities in the mean signals of the interpolated BOLD within the R/Lcing compared to the original BOLD (cosine similarity = 0.34). Meanwhile, PCC, mPFC and the recovered R/Lcing signals all show similar trends (all cosine similarity pairs > 0.9), while the original R/Lcing signal displays incoherent activity (cosine similarity with PCC = 0.036). These observations demonstrate that the interpolated functional signals are smooth over the structural brain grid, reflecting in this particular case the existing anatomical pathway (R/Lcing) running from the rostral aspect of the PCC towards the mPFC [Greicius et al., 2009, van den Heuvel et al., 2009].

PCC-seed connectivity map

In order to capture the joint structural-functional connection of the DMN, we perform a PCC seed-connectivity analysis [Biswal et al., 1995] on the interpolated volumes derived from four sessions of resting-state scans in a population of 51 subjects (244,800 volumes in total). The superior axial slice (z=25) in Figure 5.2(A) and (B) shows the PCC and the mPFC, including the bilateral inferior parietal cortex (IPC). In addition to the expected GM FC pattern, Figure 5.2(B) also shows the distinct WM structures that support the structural connections

between these temporally coherent cortical regions. The most prominent connections are the R/Lcing that connect the PCC and the mPFC, the forceps minor (Fminor) that provides intra-connectivity between gyri within the mPFC, and the left and right superior longitudinal fasciculi (R/Lslf) that support the long-range connection between the IPC, and the posterior and frontal regions. Interestingly, we also find the two-sided corticospinal tracts (R/Lcst) that traverse the brainstem all the way to the motor cortex, the genu of the corpus callosum connecting the two hemispheres, as well as the bilateral hippocampal cingulum (R/Lcing2) that exit the caudal aspect of the PCC and continue towards the medial temporal lobe (MTL).

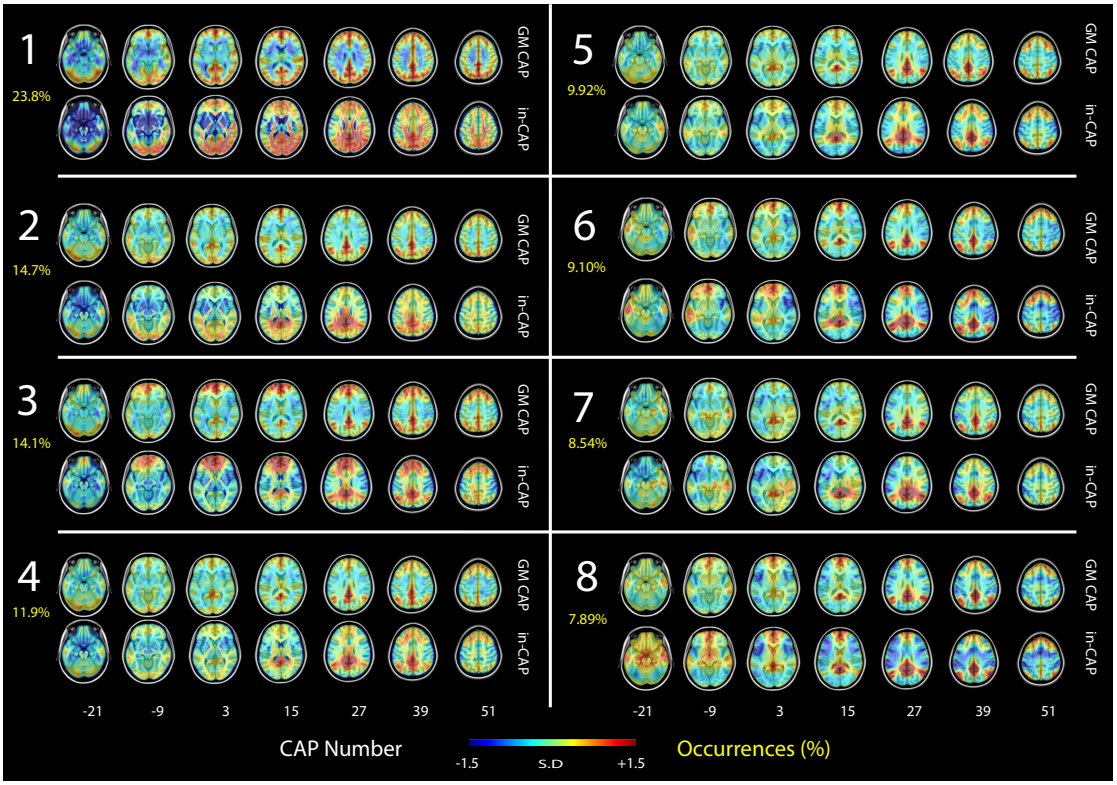

Figure 5.3 – **PCC-seed co-activation patterns (CAPs).** Eight co-activation maps obtained by temporal decomposition of the top 15% significant frames related to the posterior cingulate cortex (PCC), extracted from (1) the original BOLD volumes (GM CAPs) and (2) interpolated volumes (in-CAPs). The in-CAPs are numbered according to their frequency of occurrence, which is computed by counting the number of frame assignments corresponding to the CAP. Visual inspection of the in-CAPs reveals strong GM similarity with the observed GM CAPs. We found eight distinguishable structure-function networks related to the default mode network (DMN), each of which varies in terms of the observed WM structures that conjoin distinct PCC-coherent GM regions.

Structural mediation of dynamic functional connectivity

The axial, sagittal, and coronal slices shown in Figure 5.1(B) come from a volume corresponding to a time point with significant activity in the PCC (green shade, Figure 5.1(C)). A particular

limitation of the PCC-seed connectivity is the inherent static assumption of FC. It has been suggested that the relevant information about RS FC can be condensed into events or short periods of time (*e.g.*, where fMRI timecourses exceed a chosen threshold) [Tagliazucchi et al., 2012a]. In line with this idea, FC obtained from the conventional seed-correlation analysis has been shown to be approximated by averaging all frames for which the seed has high activity [Liu and Duyn, 2013]. This congregation of significantly active frames can then be temporally decomposed into several distinct and meaningful co-activation patterns (CAPs). We applied the CAPs analysis in all interpolated fMRI volumes on each subject. For each session, we determined the significant frames corresponding to the top 15% of the PCC. We then combined all significant volumes on each subject and decomposed them into eight PCC-related CAPs. Then, a group-level matching procedure was performed to obtain the final group-level CAPs.

Figure 5.3 presents the CAPs obtained from the interpolated volumes, termed *interpolated CAPs* (in-CAPs), as well as the CAPs generated from the original BOLD signal, denoted GM CAPs. The resulting spatial patterns reveal a set of WM structures unique to each in-CAP. In-CAP 1 is the most frequently occurring and is characterized by many negative ventral regions, in contrast to its highly positive occipital and frontoparietal regions. In-CAP 2 shows similar characteristics, but with much less pronounced activation in the frontal region, as also seen in its weaker Fminor and R/Lcing. In-CAP 3 exhibits very strong activation in the frontal lobe and negative signals in motor and sensory areas. In-CAPs 4, 5, 6 and 7 all bear high resemblance with the PCC-seed connectivity pattern, but vary in their strength of activity in the occipital and frontoparietal regions; these variations are consequently also reflected in Fminor, Fmajor, R/Lcing, R/Lslf, R/Lilf, as well as the bilateral fronto-occipital fasciculus (R/Lifo). In-CAPs 6 and 7 contain highly negative right and left frontoparietal signals, respectively, demonstrating strong hemispherical asymmetry in their activation patterns. Finally, in-CAP 8 is characterized by positive activation within the ventral brain, specifically in the MTL, but highly negative frontoparietal and somatosensory signals—the implicated WM structures are less pronounced in this in-CAP.

Spatial and temporal properties of in-CAPs

Whereas GM CAPs reveal instantaneous co-activation of multiple brain regions, in-CAPs complete the structure-function picture through the addition of distinct and functionally relevant WM structures that interpose spatially-distant GM co-activations. To assess the correspondence of the in-CAPs with the DMN subsystems, we computed the average signal within major WM tracts using the probabilistic WM atlas [Zhang and Arfanakis, 2018], and compared them to the main components of the DMN within the GM, as classified by Andrews-Hannah and colleagues [Andrews-Hanna et al., 2010]. The probabilistic WM atlas was thresholded to 0.3 to generate the mask (see Figure C.4(D) for details of finding the optimal threshold). The signal averaging within WM and GM atlases are computed in both the classical GM CAPs and the inpainted CAPs to be able to compare the two. Figure 5.4 reveals good correspondence in the

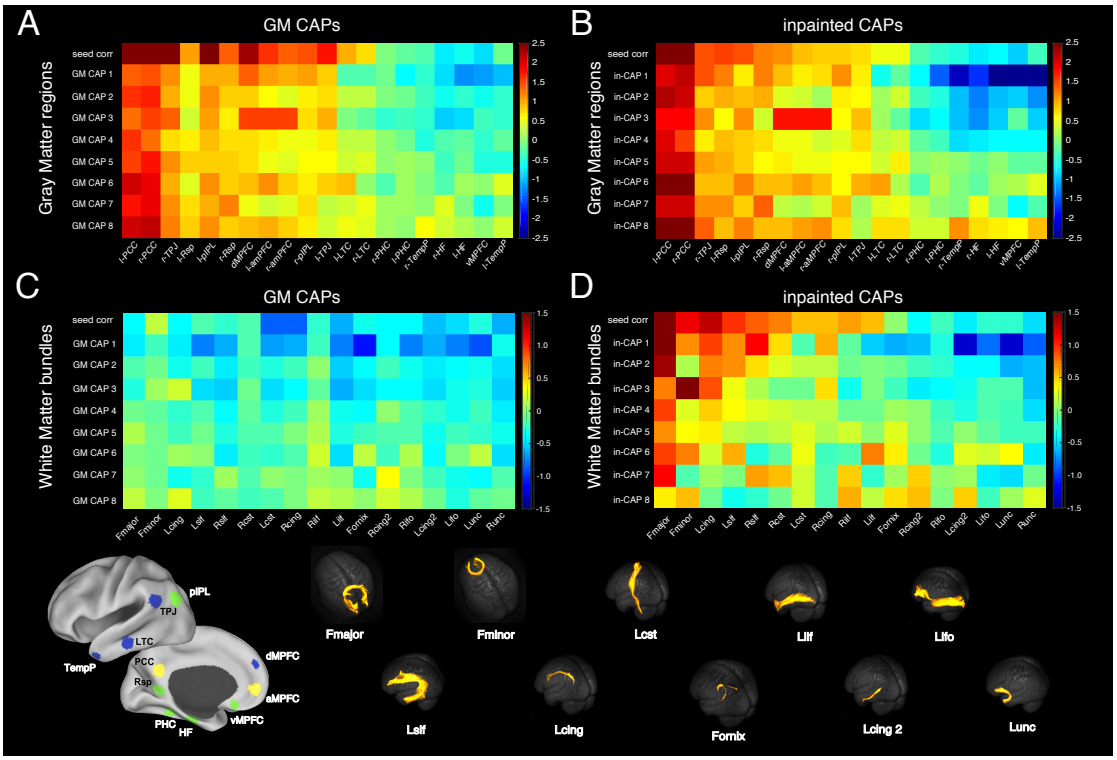

Figure 5.4 – **Spatial characteristics of in-CAPs compared to classical GM CAPs.** The overall spatial patterns of the in-CAPs and GM CAPs are explored through spatial averaging of signals within the (A & B) GM sub-components of the DMN [Andrews-Hanna et al., 2010] as well as the (C & D) 17 major WM bundles. A similar spatial trend is observed in the subcomponents of (B) DMN GM in-CAPs and (D) DMN WM in-CAPs, while there is almost no correspondence between the signal averages within (A) DMN GM regions of GM CAPs and the (D) DMN WM bundles of GM CAPs. The arrangement of WM structures and DMN sub-components based on the polarity of the inpainted CAPs with respect to the baseline also shows a concordant pattern of positive-to-negative transition from dorsal-to-ventral regions. Full labels of GM regions and WM bundles are summarized in Tables. C.1 and C.1 of Appendix C.

spatial patterns of GM regions and WM tracts in the inpainted CAPs, especially in terms of activation polarity (with respect to baseline) from the dorsal (positive-valued) to the ventral (negative-valued) side. This is in stark contrast to the spatial patterns of WM tracts in classical GM CAPs showing no accordance with their corresponding GM averages.

Moreover, we further validated the observed tracts in the in-CAPs and PCC seed connectivity based on how much they follow tractography-based results [Greicius et al., 2009, van den Heuvel et al., 2009, Figley et al., 2015]. Using predefined ROIs associated to the DMN as seeds, tractography is performed to obtain the WM fiber bundles that connect these regions. We computed the signal averages within the different WM tracts obtained from fMRI-based probabilistic tractography corresponding to the DMN ROIs [Figley et al., 2015]. The plot of the signal average is displayed in Figure C.5, side by side with the signal averages computed within

the associated DMN ROIs. In general, the signal averages corresponding to the PCC-seed connectivity display positive values in all DMN-related WM tracts, suggesting that the WM interpolation successfully captured tractography-verified WM bundles. On the other hand, the in-CAPs display varied signal averages (*i.e.*, some tracts are positive, some are negative), thereby reflecting the distinct structural-functional organization of the brain at different time points when the PCC is active. These DMN-related WM tracts are all physically connected, but are differentially wired at different time instances.

Next, we assessed the concurrence of the GM CAPs and the in-CAPs by computing the amount of temporal overlap in each CAP. This is done by counting the number of times that a single frame is assigned to a matching in-CAP and GM CAP pair. Figure 5.5(A) displays generally matching frame assignments between the GM CAPs and the in-CAPs, reflecting a strong temporal correspondence between the two CAP types. This finding demonstrates that the addition of the interpolated signals does not interfere with the segregation during the clustering, thereby showcasing the ability of the proposed method in successfully capturing distinct WM structures connecting temporally varying PCC-related networks without compromising information from the original data.

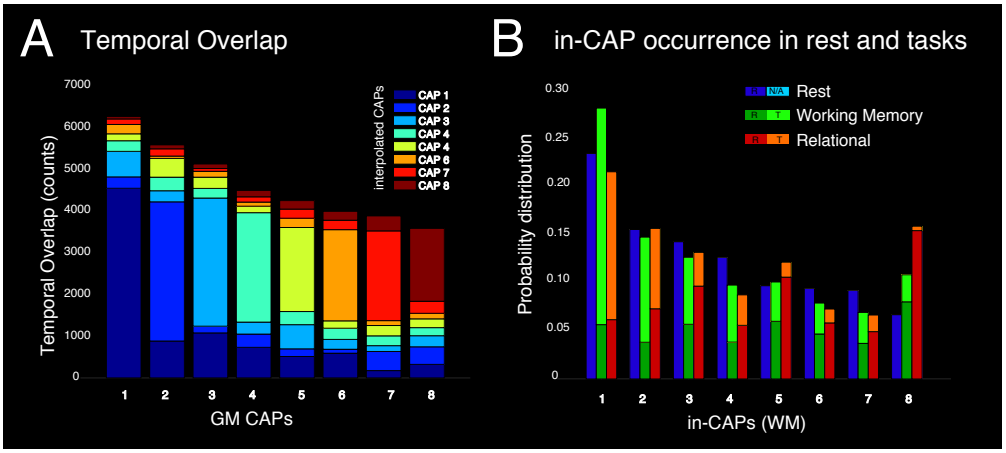

Figure 5.5 – **Temporal characteristics of in-CAPs and their relevance to rest and tasks.** (A) Temporal overlap between the interpolated CAPs and the GM CAPs; disregarding the color-code, the height of each bar represents the number of frames that were identified to belong to either of the eight GM CAPs. The sub-bands within each bar show the distribution of the frames linked to a particular GM CAP when classified based on in-CAPs. (B) The probability of occurrence of in-CAPs in resting-state, working memory, and relational tasks. Each bar in the plot corresponds to the probability of that particular in-CAP to appear in the top 15% significant volumes of rest and task scans. The stack of colors denotes the proportion of frames that occur during blocks of rest (R) when there is no stimulus presented, and during blocks of task (T) in which subjects are expected to perform memory or relational tests. Results were verified using two-factor ANOVA (main effects of CAP type and task, and interaction between both), with all p-values less than 0.0001.

Relevance of WM interpolated patterns during rest and task

Visual inspection of the spatial averages computed within the WM of all in-CAPs reveals specific patterns that are unique for each in-CAP. In-CAP 8 for instance, displays positive mean signal across all DMN sub-components, and in contrast to all other in-CAPs, also shows a positive bilateral uncinate fasciculus (L/Runc). Previous works have found that activations in both the hippocampus and ventral mPFC during retrieval-mediated learning support novel inference [Zeithamova et al., 2012]. We hypothesize that the positive activity of L/Runc in in-CAP 8 reflects its potential role in mediating the coherent activity between the positive temporal poles, hippocampus, and ventral mPFC. To verify this association and to further explore the functional meaning of the obtained in-CAPs, we computed their occurrence probability upon working memory and relational tasks (Figure 5.5(B)). To do so, we interpolated signals into the WM as before, extracted frames with high activation in the PCC, and assigned them to their closest in-CAP. We found that although in-CAP 1 is the most recruited at rest, it mostly occurs during task blocks (as opposed to rest intervals). In contrast, in-CAP 8 occurs the least during resting-state, but it is vividly recruited during the rest epochs of the task paradigms. These results are evaluated using a two-factor ANOVA on the occurrences of the different interpolated CAPs during rest and tasks. Results showed highly significant main effects of CAP type and task, as well as the interaction between both, with all corresponding p-values less than 0.0001.

5.2.4 Discussion

General findings

We modeled the brain grid using local, voxel-to-voxel probabilistic connections based on DW-MRI as a large graph onto which brain activity in GM is interpolated into WM. The interpolated brain volumes revealed WM structures that mediate distributed patterns of activity in GM. We explored the neuroscientific relevance of the interpolated volumes by applying conventional static and dynamic FC tools, namely seed-correlation [Biswal et al., 1995] and CAP analysis [Liu and Duyn, 2013], using the PCC as a seed. The resulting in-CAPs unraveled the complexity of recruited WM patterns underlying typically observed GM co-activations (Figure 5.3). In-CAPs also showed strong neurophysiological relevance, as they revealed the segregation of the DMN into task-positive and task-negative sub-components (Figure 5.5B).

Structure-function relationships of PCC-related networks

Interpolating functional signals measured on GM into the WM has revealed new structural-functional relationships that are akin to the DMN. Figure 5.4(D) shows the signal average within WM bundles observed in our PCC seed-correlation map, consistent with the visually observed tracts in Figure 5.2. Prominent tracts include the cingulum bundle (R/Lcing), the superior longitudinal fasciculus (R/Lslf), and the forceps minor (Fminor), which are all in close

agreement with previous findings of WM structures associated with the DMN [Greicius et al., 2009, van den Heuvel et al., 2009]. In addition, the signal averages within DMN-ROI guided tracts (Figure C.5) further reflect the success of the signal recovery approach in capturing WM structures that mediate temporally coherent DMN regions. Although seed-correlation analysis is a straightforward measure of FC, it is debased by its transitive nature. The connectivity obtained from this static approach, therefore, summarizes all major WM connections linked to the PCC. On the other hand, CAP analysis can account for time-dependent behavior; thus, the in-CAPs reveal different sets of WM bundles that are characteristic of each of the observed GM CAPs at different time-points.

While it has been established that the DMN anatomically consists of the anterior and posterior midline, the bilateral parietal cortex, prefrontal cortex, and MTL [Buckner et al., 2008], these regions have been found to functionally dissociate according to the ongoing internal processes [Andrews-Hanna et al., 2010]. We surmise that this functional disintegration is captured by in-CAPs, given that distinct in-CAPs occur in different phases of working and relational memory tasks (Figure 5.5(B)). The MTL activates when internal decisions involve constructing a mental scene based on memory [Andrews-Hanna et al., 2010]. However, despite being dominantly positive in this region, in-CAP 8 appears to be more active during blocks of rest, and instead, in-CAPs 1 and 2 dominate in periods of task blocks when the stimulus is presented. These findings suggest that in-CAP 8 is a characteristic *task-negative* structure-function network that activates during the maintenance phase (*i.e.*, when subjects are focused on consolidating an observed stimulus). Conversely, in-CAPs 1 and 2 are *task-positive* structure-function networks that are engaged during the encoding and retrieval periods when the external stimulus is presented. Altogether, our results provide new insights into the well-established dichotomy of the human brain [Fox et al., 2005, Greicius et al., 2004], and clearly demonstrate that the DMN and task-positive networks (*e.g.*, attention network, working memory network) are not exclusively expressing opposite activity [Piccoli et al., 2015, Karahanoglu and Van De Ville, 2015].

Arranging the DMN sub-components according to the polarity of their activity in each of the in-CAPs reveals a general trend of positive dorsal areas, such as those observed from PCC to temporal-parietal junction in Figure 5.4(B). On the other hand, ventral components (lateral temporal cortex to temporal pole) are more diverse, showing a subtle divide between more negative (in-CAPs 1-4) versus more positive (in-CAPs 8) ventral default network. The implicated dorsal and ventral WM bundles also show the same trend in Figure 5.4(D). In addition, we observe the sagittal slices (Figure C.1) of in-CAP 8 to show strong activation in the ventral PCC together with the R/Lcing2. This bundle runs from the caudal aspect of the retrosplenial cortex to the ventrally located hippocampal formation (HF). The bilateral HF, PCC, R/Lcing2, and temporal poles also exhibit positive signals and, therefore, more homogeneous averaged activity. Overall, the above observations highlight the distinctive role of dorsal and ventral parts of the PCC in cognitive control [Leech et al., 2011, Karahanoglu and Van De Ville, 2015]. Furthemore, in line with our hypothesis about the role of R/Lunc in memory consolidation, we observe that in-CAP 8 displays positive signal in this particular

tract, reflecting its role in mediating coherent activity between the hippocampus and the prefrontal cortex [Squire et al., 2015]. Fittingly, it has been found that the deterioration of R/Lunc lowers working-relational memory performance [Hanlon et al., 2012].

Methodological perspectives

The overall aim of the introduced signal recovery framework is to extract WM structures that mediate the interaction of temporally coherent cortical regions. Several methods that have attempted to achieve the same goal have been proposed over the past years. The most straightforward approach is fMRI-guided tractography, where cortical regions corresponding to well-known functional networks are used as seeds to estimate streamlines coming from one ROI to another [Figley et al., 2015, Friman et al., 2006]. One obvious limitation of this approach is the need for a two-step procedure, wherein as a first step, ROIs are extracted from functional data, then tractography is launched to recover the streamlines that connect pre-selected regions. In our approach, we go beyond a combined analysis [Zhu et al., 2014] of functional and diffusion MRI, and instead jointly model the two modalities into a single integrated framework that enables the recovery of structure-function patterns which are fully data-driven and observer-independent. Previous works have also introduced a fusion of fMRI and diffusion MRI in a data-driven manner by relying on blind-decomposition techniques (*e.g.,* ICA) and the use of diffusion anisotropy measures to approximate SC [Sui et al., 2013, Sui et al., 2015]. However, measures of diffusion anisotropy are unable to resolve crossing fibers as they are extracted from low-order diffusion tensor models, making their interpretation misleading and prone to erroneous conclusions [Wheeler-Kingshott and Cercignani, 2009]. Alternatively, a more common approach in defining SC is to count the number of fibers obtained using tractography algorithms [O'Muircheartaigh and Jbabdi, 2018]. SC and FC measures can then be combined together into a huge data matrix that is decomposed into independent hybrid structure-function components [Amico and Goñi, 2018]. Nevertheless, ICA's principal assumption of spatial independence leads to components that have low spatial similarity. The high susceptibility of tractography algorithms in generating false positives [Maier-Hein et al., 2017] also warrants caution in interpreting tractography-based results. In contrast, our approach does not rely on tractography, thereby not only bypassing the complicated task of parameter optimization for extracting WM tracts but also not requiring a parcellation scheme to define ROIs. Our framework keeps the analysis close to the DW-MRI data by working directly with local reconstructions of orientation distribution functions, which are consequently encoded, both in direction and magnitude, to build the brain graph. Overall, we believe that our method offers a more elegant solution to recover WM structures that mediate FC regions, and at the same time, reduces the computational load compared to conventional tractography approaches.

Furthermore, it is worth noting that despite the superficial similarity, there are fundamental differences between our approach and that of the TW-dFC [Calamante et al., 2017]. Specifically, both approaches attempt to give a measure of FC into the WM using GM as a constraint.

Although TW-dFC jointly fuses fMRI and diffusion MRI, the approach still relies on two successive uses of classical methods (*i.e.,* tractography and sliding-window FC), while our work introduces an entirely new concept where recovery of mediation by SC is incorporating the FC itself in a consolidated framework; *i.e.,* the complete pattern of brain activity in GM is constraining the recovery of WM patterns. In addition, our method works at the single-frame level and does not require to choose a temporal window to render the approach dynamic. It is therefore noteworthy to reiterate that the frame-wise analysis using the CAPs was made possible by the fact that the mediation of GM activation is done for each frame. As such, we make use of the maximal temporal resolution of the functional data, which is a characteristic strength of the introduced method that is not yet overcome by tractography-based methods for jointly integrating FC and SC [Amico and Goñi, 2018, Calamante et al., 2017, O'Muircheartaigh and Jbabdi, 2018]. This advantage also translates to easier standardization of results across groups, as classical methods of spatial normalization can be applied in the interpolated volumes in the same manner that is done in original fMRI.

New research avenues for structure-function studies

GSP approaches as applied to functional neuroimaging rely on a *graph shift operator* (typically the Laplacian or the adjacency matrix) that encodes the brain's anatomical information. Its eigendecomposition produces an orthogonal set of eigenvectors, termed *eigenmodes*, some of which are reminiscent of well-established functional networks [Atasoy et al., 2016, Abdelnour et al., 2018] and are useful in effectively estimating the strength of inter-hemispheric interactions in the brain [Robinson et al., 2016]. Moreover, efficient anatomically-informed decompositions of fMRI data using a tractography-based structural connectome [Atasoy et al., 2016], as well as topology encoding GM graphs [Behjat et al., 2015] have been proposed. Brain eigenmodes can be viewed as basic building blocks of increasing spatial variation along the structural brain graph, akin to sinusoids in classical Fourier analysis. Here, we extend the traditionally region-level eigenmodes (defined on a limited subset of GM nodes) to a whole-brain (GM and WM), voxel resolution setting (see Figure C.2), thus enabling to reconstruct structure-function networks at an unparalleled spatial resolution. As an increasing number of operations are generalized from classical signal processing to the graph setting [Shuman et al., 2013, Wang et al., 2017, Huang et al., 2018a, Becker et al., 2018], promising avenues arise to explore the many facets of brain structure and function.

We therefore foresee two direct avenues for future research. The first avenue aims on leveraging the proposed interpolated fMRI data. The interpolated volumes entail additional informative voxels that extend beyond the GM, which is in particular interesting in light of the limitations and skepticisms in interpreting WM BOLD data [Logothetis and Wandell, 2004, Gawryluk et al., 2014]. Previous works have demonstrated the possibility to capture functionally relevant information from the WM BOLD [Peer et al., 2017, Ding et al., 2018, Huang et al., 2018b], despite the well-established findings on the differences of hemodynamic responses in GM and WM [Fraser et al., 2012, Li et al., 2019]. At its current form, we interpolate functional signals

into the WM as solely constrained by the functional signals from the GM. Alternatively, one may rather consider combining signal interpolation with signal recovery of weak signals in the WM. In our current application, the goal is to observe WM pathways that support distributed patterns of FC in GM, and not on the functional organization of WM itself. The interpolated volumes can then be readily explored using existing tools for dynamic FC analyses [Preti et al., 2017], such as sliding-window correlation, ICA [Beckmann et al., 2005] and principal component analysis [Leonardi et al., 2013], to probe functional brain dynamics at a much larger scale.

The second foreseen research direction is to exploit the proposed anatomically-informed brain grid to implement novel GSP operations on fMRI data. High-resolution eigenmodes enable a spectral graph-based analysis of task-based and resting-state fMRI data at an unprecedented level of detail, and across the whole brain, in contrast to the conventional region-wise analysis, which is typically also limited to the cortical layer. Anatomically-informed spectral graph decomposition of fMRI data using these high-resolution brain eigenmodes is anticipated to open up new perspectives on the brain's functional organization not only within GM but also WM. As a case in point, FC has often been associated to Euclidean distance [Ercsey-Ravasz et al., 2013, Alexander-Bloch et al., 2013]; thanks to our framework, the notion of distance becomes more meaningfully defined in terms of the spectral representation of functional signals [Medaglia et al., 2018, Preti and Van De Ville, 2019], enabling a deepened understanding of the roles of short- and long-range connectivity in improving the efficiency of inter-areal brain communication [Betzel and Bassett, 2018].

Technical limitations

While it is true that our approach bypasses the use of tractography algorithms in visualizing WM tracts that support coherent GM activations, our framework is still heavily reliant on the estimation of the orientation density functions. Consequently, our approach is highly dependent on the quality of the diffusion data, as in any other case of DW-MRI studies. We also point out that because we interpolate signals into the WM based on the signals from GM, any discrepancy that is associated with the GM BOLD will likely influence the results of the interpolation. Furthermore, we believe that although jointly integrating functional and diffusion MRI offers several benefits, a potential disadvantage can arise when there is a need to disentangle the effect of the two. This is particularly true when the approach is applied to clinical populations that are hindered by complex mixes of structural and functional alterations, such as in Multiple Sclerosis patients [Giorgio et al., 2017]. In such cases, depending on the hypothesis that can be drawn from the initial results of the inpainting, a detailed case to case analysis is advised. For instance, one may perform a null comparison test, such as contrasting the WM interpolation of a healthy control data to a functionally impaired patient, both of which are interpolated over a structurally deficient brain grid.

5.2.5 Conclusion

This study presents a new framework for studying structure-function relationship that has several key advantages. The interpolation of fMRI activity into the white matter enables observing key WM structures that link interacting GM regions at the single-frame resolution. We believe that the introduction of highly-resolved human brain eigenmodes can (i) shift the existing trend of constructing region-wise connectomes to that of voxel-vise connectomes and (ii) expand the use of the GSP repertoire in the context of functional brain imaging. More importantly, we anticipate the proposed joint structure-function characterization to offer unprecedented benefits for the study of clinical populations, particularly those born with structural deficits but preserved functional efficiency, such as in patients with agenesis of the corpus callosum [Tovar-Moll et al., 2014, Owen et al., 2013].

6 Graph signal processing applications on functional brain imaging: a voxel level perspective

6.1 General Motivation

The construction and development of voxel-level structural brain graphs in Chapter 4 was largely motivated by its huge potential to integrate brain structure and function through graph signal processing (GSP) analyses of functional data. Meanwhile, the interpolation of GM signals onto the WM in Chapter 5 enabled the observation of key WM structures that mediate short- and long-distance functional interactions in the brain. In the perspective of brain network topology, the distinction between short- and long-distance, remains loosely defined, and often times do not consider the actual length of white matter structural connectivity.

In this last chapter of the thesis, we develop a novel framework that redefines short- and long-ranged functional connectivity in terms of a quantitative measure that collectively describes the overall structural mediation of temporally coherent GM areas. We propose to give a more meaningful definition for cortical distance using a canonical model of graph diffusion. Using methods borrowed from spectral graph theory, we implement a voxel-wise graph diffusion of GM signals onto the WM based on the structure encoded in the brain graph. Noting that the diffusion of gray matter signals follows the structure encoded in the brain grid, we postulate that different regions should have different optimal diffusion time lengths that is proportional to the overall range of its anatomical connectivity to all other brain areas which display temporally coherent activity. To confirm this hypothesis, we first demonstrate that the graph signal recovery framework (GSR), which we used to interpolate GM signals onto the WM in Chapter 5, gives a stable solution to the graph diffusion. Next, we define different regions of a functional atlas as ROI in a seed-connectivity analysis that is applied to all diffused 4D fMRI volumes of varying diffusion time lengths. In each ROI, we survey the optimal seed-connectivity map associated to different diffusion time lengths that conforms best to the one obtained using the interpolated volumes from the GSR. The optimal diffusion time length of each ROI is then converted to spatial distance in millimeters, finally obtaining the anatomical range of each ROI with respect to temporally coherent brain areas.

6.2 Journal Article: Anatomical range of functional interactions meaningfully differentiates brain regions

A. Tarun and D. Van De Ville

Full Research Article: Preprint to be Submitted.

Author Contributions

A. Tarun conceptualized research and performed all stages of the analysis, from methodology, validation, formal analysis, writing of the original draft and editing of the manuscript and **D. Van De Ville** conceptualized research and offered continuous guidance.

Abstract

Functionally coupled brain regions recruit different combinations of white matter structures to accommodate local and distant inter-areal communications. Here, we introduce a new framework that quantifies the range of structural mediation between temporally coherent gray matter areas. We use graph signal diffusion as a tool to provide a canonical example of the spread of gray matter activity along a structurally-informed high-resolution voxel-level brain grid. The diffused volumes are evaluated using a seed-connectivity analysis which enables the observation of key white matter structures that are associated to each region of a functional brain atlas. We observed that each region displays different optimal diffusion time length that is proportional to its overall structural connectivity with all other brain areas that display temporally coherent activity i.e., regions that are dominated by local, unimodal patterns of functional connectivity require shorter time for their signals to be completely diffused. As such, the optimal diffusion time length provides a proxy of the anatomical range of functional network interactions. Our results reveal that the observed structure-function range meaningfully organizes the cortex according to a behaviorally-relevant macroscale gradient that differentiates regions associated to low-level primary sensorimotor from high-order cognitive functions.

6.2.1 Introduction

The measure of cortical *distance* between communicating brain areas remains loosely defined in neuroimaging. One of the most widely sought questions to answer in this field is the understanding of the architectural principles that support brain function. Connectomics has made it well-known that across scales and species, the brain follows a *small-world* property, similar to various social, economic, and biological processes [Bassett and Bullmore, 2006]. This description simply means that the brain is composed of highly connected short-distance connections, with a few long-distance links that serve as a short-cut for efficient information transfer, therefore reducing the overall network path length [Sporns and Zwi, 2004, Bassett et al., 2010, Bassett et al., 2009]. As such, the brain follows an optimized spatial embedding that balances wiring cost and topological value which paves the emergence of global neuronal coordination [Bassett and Bullmore, 2017, Bullmore and Sporns, 2012].

Euclidean distance between regional nodes is typically interpreted as the currency when talking about the cost of inter-areal communication [Bullmore and Sporns, 2012, Bassett and Bullmore, 2017]. It is the most straightforward measure of cortical functional distance, and it has been shown to delineate regions that are mostly functionally coupled with their local neighborhood to those that show preferential connections to distant areas [Achard et al., 2006,Sepulcre et al., 2010]. This is typically performed by correlating mean brain activity within a region of interests (ROIs) to all other brain areas (a.ka. functional connectivity (FC) [Friston, 2011]), and obtaining the most significant connections following a thresholding step. Then, the Euclidean distance between the centroid coordinates of regions that survived the thresholding are calculated with respect to the ROI. Such analyses demonstrate the tendency of primary sensory and motor areas to display high local connectivity, consistent with the brain's modular organization. Moreover, connection lengths between brain regions have also been successfully used as biomarkers for clinical studies comparing brain network properties of healthy controls to patients suffering from schizophrenia [Alexander-Bloch et al., 2013].

The use of Euclidean distance in the studies above effectively removes the contribution of the brain's anatomical configuration and is therefore reliant only in its geometrical shape. First, cortical foldings are ignored, and second, the physical wirings of the white matter (WM) fiber bundles are not taken into account. To overcome the first constraint, geodesic distance, which measures connection length along the cortical sheet, has been proposed as an alternative to the widely used euclidean distance [Wagstyl et al., 2015]. It has been applied to delineate short- and long-range connectivity in humans [Oligschläger et al., 2017] and macaque [Oligschläger et al., 2019], and has also been demonstrated to show higher accuracy in participant fingerprinting [Venkatesh et al., 2020] compared to conventional FC measures [Finn et al., 2015, Amico and Goñi, 2018].

Due to the direct computation of Euclidean distance, there is currently a very limited number of studies that take into account the contribution of the WM fiber connectivity itself in the currency of physical wiring cost for inter-areal communication. One of the earliest studies

that incorporated structural features within the WM has shown that WM connectivity between brain regions derived using diffusion tensor imaging (DTI) decays as a function of the Euclidean distance [Lewis et al., 2009], in a similar manner that the FC strength computed using functional magnetic resonance imaging (fMRI) also decays with distance [Salvador et al., 2005]. This relationship has been observed to persist even after controlling for the strong association between structural connectivity (SC) and FC [Honey et al., 2009]. Recently, WM fiber lengths were projected on each region of the cortical surface to characterize how structural brain organization relates to functional resting-state networks [Padula et al., 2017, Bajada et al., 2019]. They used the distribution of track lengths to differentiate closely connected local modules to those of functionally more distant brain areas.

In this work, we propose a new definition of functional distance based on a sophisticated integration of both WM anatomical connectivity from diffusion-weighted MRI (DW-MRI) and gray matter (GM) BOLD activity from fMRI. Leveraging on the emerging field of graph signal processing (GSP) applied to functional brain imaging [Huang et al., 2018a] that was recently extended to voxel-level perspective [Tarun et al., 2020b, Tarun et al., 2019], we derive the anatomical range of functional network interactions. In particular, we utilize a graph diffusion model to obtain a canonical measure of the time it takes for a signal to be diffused from one region of the brain to another. Each region in the cortex reveals a distinct optimal diffusion time that is proportional to the length of its physical separation from other brain areas that display temporally coherent activity.

To do this, GM BOLD signals are first interpolated into the WM using a graph signal recovery (GSR) framework introduced in a previous work [Tarun et al., 2020b]. The interpolation is guided by a voxel-level brain graph that encodes the structural information from diffusion-weighted MRI (DW-MRI) [Tarun et al., 2019]. The result is a 4D interpolated fMRI volume, which can be analyzed using classical FC methods (*e.g.*, seed-correlation map) to obtain a set of entire-brain structure-function networks. These spatial maps are taken as the ground truth or the final solution to the graph diffusion which we used to model the anatomical range of functional interactions. Using methods from spectral graph theory, we diffuse signals from the original GM BOLD data onto the WM using different lengths of time. As before, seed-connectivity maps are extracted for each of the 4D diffused frames, one set for each diffusion time that we evaluate. We then assess which among these seed-connectivity maps conform best with the seed-connectivity maps extracted using GSR. These values are converted to an approximate measure of distance in millimeters through empirical investigations of their one-to-one correspondence. We observe that, depending on the seed, different brain areas exhibit different optimal diffusion times. Our results demonstrate, that the anatomical range of functional interactions follows a functional differentiation that divides the cortex into low-level primary sensorimotor areas and more complex and integrative higher-order association and limbic networks.

6.2.2 Methods

Data

We used a total of 50 subjects obtained from the publicly available Human Connectome Project dataset (HCP 1200 release), WU-Minn Consortium. Preprocessed diffusion MRI, original functional MRI, and anatomical data are downloaded following a random subject selection scheme while maintaining a uniform demographic distribution (male and female, ages 22-50). MRI acquisition protocols of the HCP and preprocessing guidelines for diffusion MRI are fully described and discussed elsewhere [Glasser et al., 2013]. No further preprocessing was applied to the diffusion data. On the other hand, original and unprocessed functional MRI data underwent preprocessing following standard protocol [Van Dijk et al., 2010]. We used four sessions of resting-state fMRI scans (1200 volumes each session, giving a total of 4800 volumes per subject). Functional images of each subject were realigned to their mean images. These were then registered and resampled onto the diffusion data using rigid-body registration (SPM12, https://www.fil.ion.ucl.ac.uk). After which, the registered images were detrended and smoothed using a Gaussian kernel (FWHM = 6mm). We discarded the first 10 frames to achieve steady-state magnetization of the fMRI data. Finally, we z-scored the functional data in the temporal domain by demeaning each voxel time-course and dividing it to its standard deviation.

White matter interpolation of fMRI via graph signal recovery

We built upon a recently introduced framework based on GSR for interpolating GM signals into the WM using a voxel-wise connectome to guide the process. The voxel-wise brain graph is defined as $\mathcal{G} := (\mathcal{V}, \mathbf{A})$, where $\mathcal{V} = \{1, 2, 3, ..., N\}$ is the set of N nodes representing the brain voxels (i.e., $N = 700 - 900$ thousand), and $\mathbf{A} \in N \times N$ is an adjacency matrix that encodes the probability of alignment between two neighboring voxels. We then used orientation distribution functions (ODF) as a weighting-scheme for embedding DW-MRI data into a three-dimensional 26-neighborhood mesh. For more details about the brain graph construction, we refer to [Tarun et al., 2019] and Chapter 4 of the thesis. After the brain graph was constructed, we interpolated GM signals into the WM by minimizing a cost function that includes a least-squares data-fitting term equal to the residual sum of squares (RSS), and an L2 regularization term that takes into account smoothness with respect to the brain grid. The graph structure is represented by the graph Laplacian matrix in its symmetric normalized form ($\mathbf{L} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}}$). The interpolated signal is then given by the solution to:

$$\tilde{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \left\| \mathbf{y} - \mathbf{B}\mathbf{x} \right\|^2 + \lambda \mathbf{x}^T \mathbf{L} \mathbf{x}, \tag{6.1}$$

where $\mathbf{y}$ is a vector of length N containing initial BOLD values within GM nodes and zero elsewhere, $\mathbf{B}$ is an indicator matrix that selects the GM nodes, and λ is the regularization parameter. The cost function balances between (1) minimizing the RSS and retaining the

original GM signals, and (2) imposing smoothness with respect to the structure of the graph. The norm in the first term therefore focuses on keeping the GM signals unchanged, while the second term ensures smoothness with respect to the Laplacian which encodes the structure. The balance is dictated by the optimal λ and the interpolation is done repeatedly for each volume in the functional data. For more details about the white matter interpolation. in particular how the regularization parameter λ is chosen, we refer to our paper [Tarun et al., 2020b] or to Figure C.3 in the appendix.

White matter interpolation of fMRI via graph diffusion

Suppose we initialise a source signal at time $t = 0$ in a particular locality defined by the vector $\mathbf{x}(0) = \{x_0[i], i \in \mathcal{V}\}$, *i.e.,* GM BOLD signals defined on GM nodes, then from spectral graph theory [Kondor and Lafferty, 2002], we expect the diffused signal at a later time t to be governed by a first-order linear network heat equation given by:

$$\frac{\partial \mathbf{x}(t)}{\partial t} = -\mathbf{L}\mathbf{x}, \tag{6.2}$$

which is equivalent to the classical and continuous diffusion equation $\frac{\partial x}{\partial t} = \nabla^2 x(t)$ that is applied in many areas of physics in engineering. The discrete counterpart of the second-order derivative in this case is the graph Laplacian $\mathbf{L}$. If we approximate the solution to the diffential equation in Eq.6.2 through numerical methods, a finite-difference approach can be employed, wherein in this case, the Laplacian matrix acts as a difference operator. One can therefore iteratively apply the Laplacian operator to the signal at each timepoint t until we reach a stable solution $\mathbf{x}(t_{final})$. The t_{final} is equivalent to the characteristic time of spread required to reach the steady-state solution. Finding the optimal value of t_{final} will be the primary goal of this work, but will be implemented using a more straightforward and computationally easier method discussed below.

From the perspective of matrix algebra, the approach discussed above can be implemented faster by introducing discrete diffusion operators and exploiting spectral kernels [Kondor and Lafferty, 2002, Shuman et al., 2013]. The solution to Eq. 6.2 is an exponential given by:

$$\mathbf{x} = \exp(-\tau\mathbf{L})\mathbf{x}_0, \tag{6.3}$$

where $\mathbf{x}$ represents the amount of heat at each vertex of the graph after time τ. Exponentiating the Laplacian matrix with time τ, we have

$$\exp(-\tau\mathbf{L}) = \mathbb{1} - \tau\mathbf{L} + \frac{\tau^2}{2!}\mathbf{L}^2 - \frac{\tau^3}{3!}\mathbf{L}^3 + ... \tag{6.4}$$

Exploiting the orthogonality of $\mathbf{L}$ and diagonalizing, we obtain $\mathbf{L} = \mathbf{V}\Lambda\mathbf{V}^\top$, where $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, ..., \mathbf{v}_N]$ and Λ is a diagonal matrix whose diagonal elements are the corresponding eigenvalues, denoted $\{\lambda_i := \Lambda_{k,k}\}_{k=1}^N$. Substituting this to Eq. 6.4, Eq. 6.3 can then be expressed in terms of

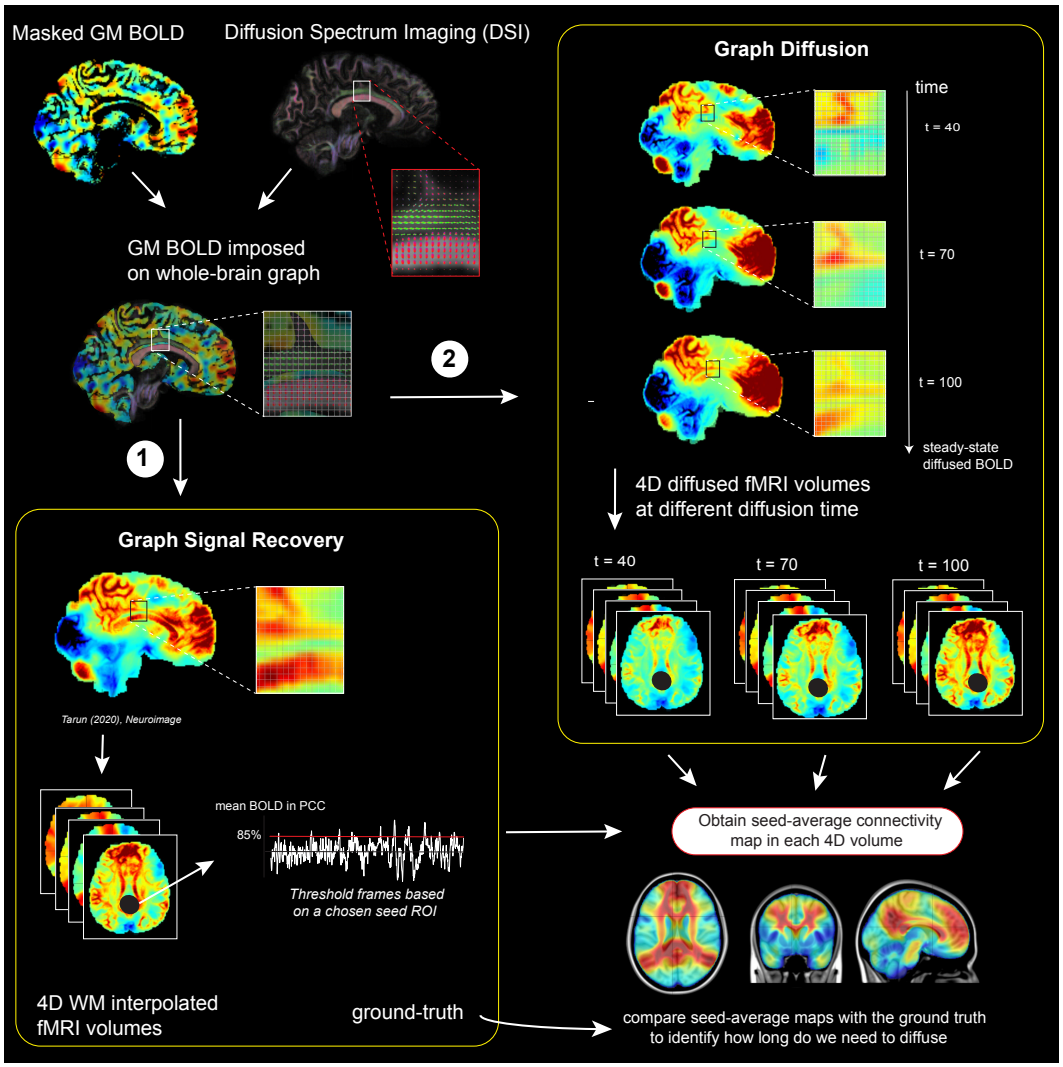

Figure 6.1 – **Overview of the pipeline.** The original GM BOLD from one fMRI volume is first masked using a probabilistic GM mask from the segmented T1 image. Meanwhile, the ODFs are extracted from DW-MRI and local voxel-to-voxel structural connectivity is encoded in a whole-brain graph. GM BOLD signals are imposed on the brain graph, *i.e.,* each GM voxel signal is defined on GM node, and zero elsewhere. Then, two succeeding methods of WM signal interpolation are implemented, (1) one based on graph signal recovery (GSR) and (2) one based on a graph diffusion model. The GSR is implemented by minimizing a cost function that balances between keeping the GM signal fixed and obtaining a whole-brain signal that is smooth with respect to the underlying structure. For graph diffusion model, GM BOLD signals are diffused at different amounts of time. Both methods provide different 4D WM interpolated fMRI volumes. We evaluated different time of spread for each region in the brain, and use the GSR approach as ground truth to identify which among the evaluated diffusion time-scales gives a seed-average connectivity map that conforms best with the seed-average connectivity derived from GSR interpolation. The seed-average map is obtained by averaging frames whose chosen ROI has a mean signal above 85 percentile. PCC - posterior cingulate cortex

matrix exponential:

$$\mathbf{x} = \mathbf{V}\exp(-\tau\Lambda)\mathbf{V}^{\top}\mathbf{x}_0. \tag{6.5}$$

The exponential term is called the diffusion kernel and acts as a spatial blurring on the initial heat source $\mathbf{x}_0$. Intuitively, different values of τ give different descriptions of the flow of heat over the graph, recalling that the rate of flow is dependent on the weights of the edges encoded in $\mathbf{L}$. When τ is small, Eq.6.4 is governed primarily by the first terms and as such, the solution depends on the local connectivity structure or topology of the graph. Meanwhile, for large values of τ, Eq.6.4 is dominated by higher order terms, and the heat kernel becomes dependent on the values of vertices that are farther away, thus, giving a solution that follows the global structure of the graph.

The solution to the graph diffusion equation discussed above is equivalent to the action of a low-pass filter into the initial signal source $\mathbf{x}_0$, where the diffusion kernel, which decays exponentially with the edge weights, is the filter. In this work, we define each GM BOLD signal as the source signal and interpolate them into the WM following the graph diffusion model discussed above. It is also important to establish that the graph diffusion model is only taken as a tool to model the anatomical range of FC, and is used as a canonical example of the spread of GM signals along WM fiber bundles. It is not meant to model the inter-areal interactions at the neuronal level.

General Workflow

Figure 6.1 gives a visual summary of the working pipeline used in this study. First, original GM BOLD data extracted from a 4D set of fMRI volume is masked, one frame per time-instance, using the probabilistic GM mask from the segmented T1 image. Meanwhile, the voxel-level brain graph is constructed by encoding local, voxel-to-voxel SC from DW-MRI in each edge of the brain graph. Then, each signal in the masked GM BOLD is defined on the GM nodes of the graph. The rest of the nodes are set to zero. Then, two succeeding WM interpolation methods are performed: (i) GSR and (ii) graph diffusion. For method (i), the solution is straightforward, which is given by the minimization of the cost function in Eq. 6.1. The final output is a series of 4D fMRI volumes which contains, in addition to GM BOLD signals, the inpainted WM signals. Meanwhile, method (ii) requires the need to identify the proper value of the diffusion time τ in order to obtain the final solution. We evaluated a range of $\tau-$values within [0,10,20,...,200]. The output is also a series of 4D fMRI volumes, one for each of the $\tau-$value assessed. To identify which among the diffusion times gives the steady-state solution to the graph diffusion, we assessed the similarity of their resulting seed-connectivity map to the seed-connectivity generated using the WM interpolated fMRI volumes using GSR.

We applied a seed-connectivity analysis to the 4D WM interpolated and diffused volumes. We used regions of the 17-network Yeo functional atlas [Thomas Yeo et al., 2011], which has a total of 97 regions. Regions that belong to the same network and consists of less than 50 voxels at 1mm atlas resolution are merged. We identified the top 15% significant frames corresponding

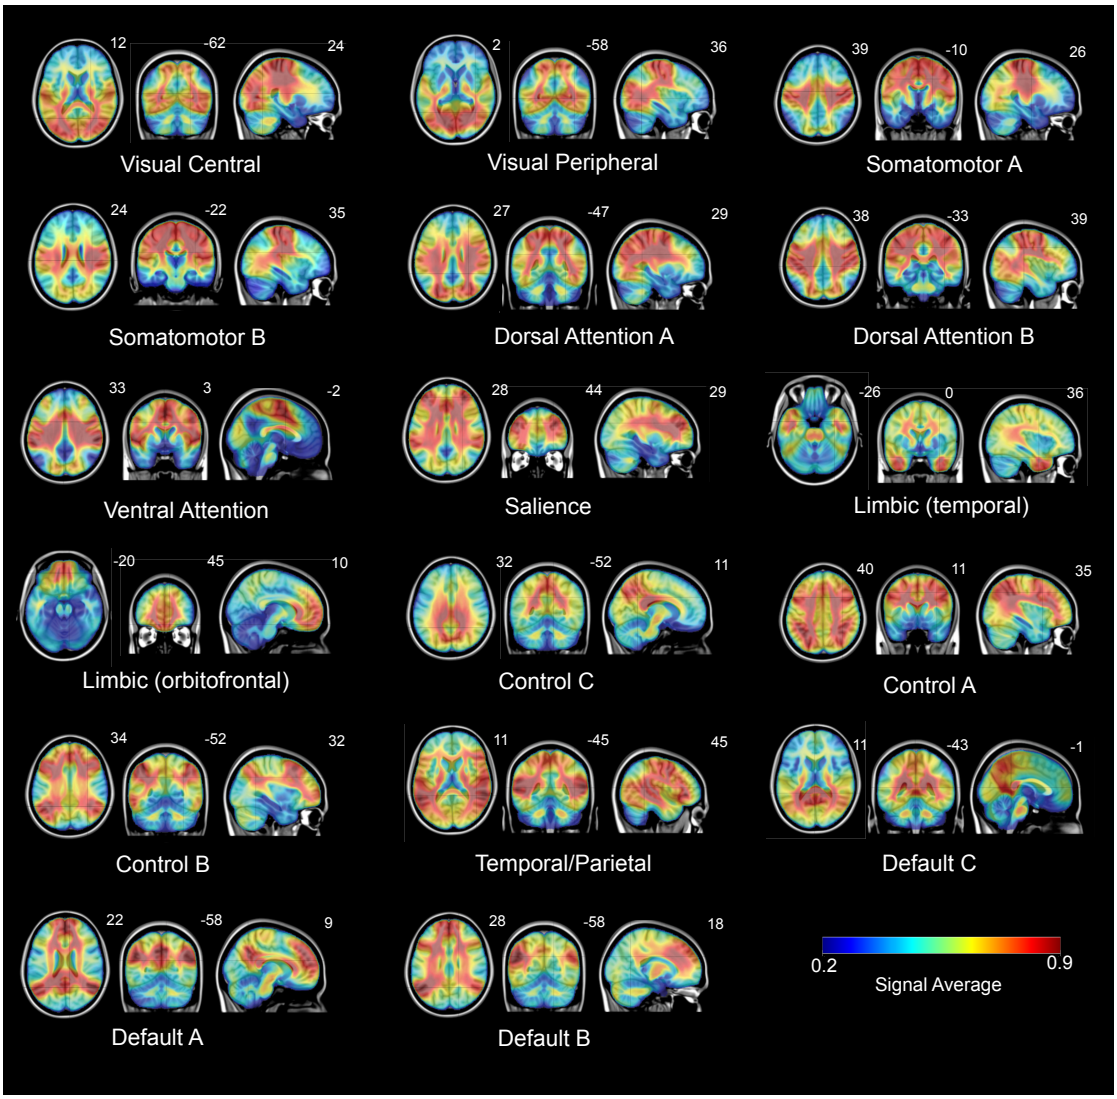

Figure 6.2 – **Whole brain structure-function networks.** 17 functional networks extracted by averaging the top 15% significant frames corresponding to each ROI of the Yeo functional network atlas.

to the mean BOLD signal in each ROI. Only the frames that survived this thresholding are concatenated and are averaged across subjects forming a seed-averaged connectivity map for each ROI. This is done for all 97 ROIs and each of the 4D interpolated/diffused volumes generated: (i) 97 sets for the GSR [97 × 1], one for each ROI, and (ii) 2037 sets for the graph diffusion [97 × 21], one for each ROI and each of the evaluated $\tau-$value.

The pipeline is mostly done in the diffusion subject space, except at the last part of the workflow in which the seed-connectivity maps of each subject are normalized to the Montreal Neurological Institute (MNI) template in order to perform a group-level averaging of the connectivity maps.

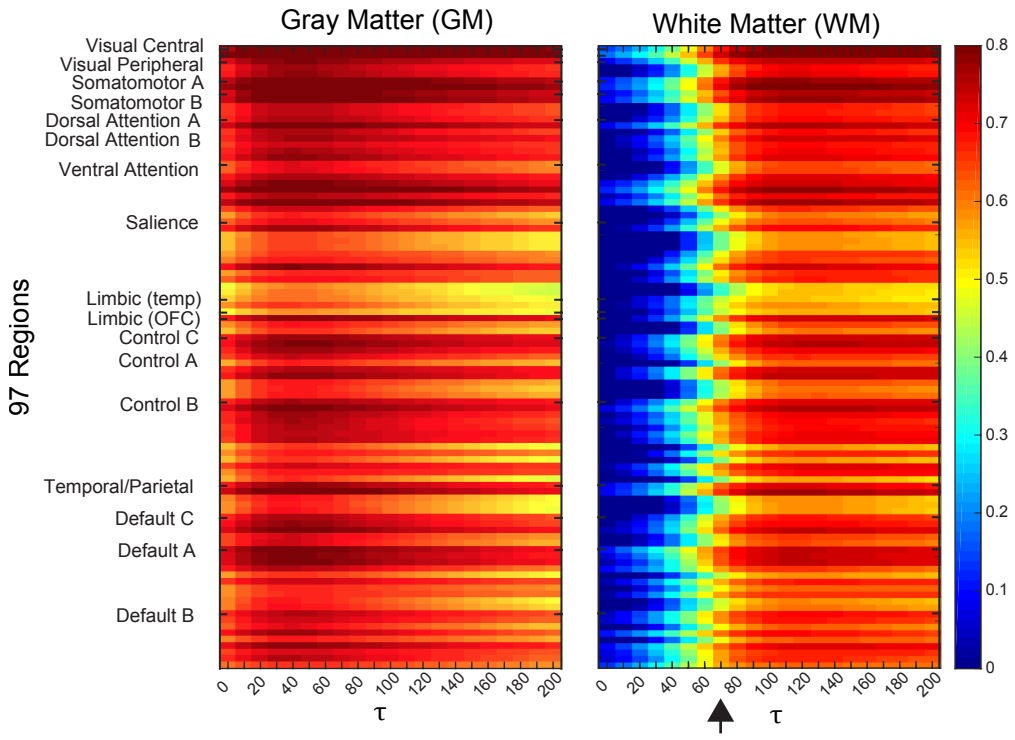

Figure 6.3 – **Cosine similarity of seed-averaged diffused volumes** using different diffusion time τ with respect to the seed-averaged maps obtained using GSR. Similarity matrices are derived separately for the GM voxels and the WM voxels. The black arrow indicates the starting point when a phase-transition can be observed, *i.e.,* structures start to form in the WM, finally showing similarities to seed-averaged maps extracted using GSR.

6.2.3 Results

Graph signal recovery gives the optimal solution to graph diffusion

GSR and graph diffusion can be interpreted as two different approaches of white matter signal inpainting. The GSR minimizes a cost function that is dependent on two constraints – retaining the GM BOLD signal fixed and smoothness with the structure encoded in the Laplacian. This optimization relies on the regularization parameter that balances these constraints. Meanwhile, graph diffusion relies on a parameter that dictates the amount of time required for the signals to be diffused. Figure 6.1 gives a visual aid of the overall pipeline, and also demonstrates for an illustrative frame the resulting interpolation using GSR and the evolution of the graph diffusion at different times. Both the GSR and the graph diffusion rely on the structural brain grid which encodes the structure from DW-MRI. As such it can be observed in the insets of Figure 6.1 that the solution to the GSR follows a similar pattern with the brain grid. The same is observed in the insets of the sagittal slices representing diffused GM signals at different times., *i.e.*, a signal is more likely to diffused in regions where there is a strong structural link. In particular, for these two approaches, we find that a specific WM

structure connects the posterior region of the brain and the frontal region. This structure corresponds to the cingulum bundle, which anatomically connects the posterior cingulate cortex (PCC) and the medial prefrontal cortex [Greicius et al., 2009, Tarun et al., 2020b].

Structural mediation of functionally coupled regions revealed using seed-connectivity map

Seed-correlation analysis is a classical method for extracting functional connectivity networks in the brain, wherein the mean BOLD signal within a specified ROI is correlated to all brain voxels in the GM. This analysis results to a spatial map that displays which regions in the brain are temporally coherent with the chosen seed. It has been shown through a point-process analysis [Tagliazucchi et al., 2012a] approach that only time-points that show significant activation in the chosen ROI contribute to the observed correlation, and as such, one can also take the average of the significantly active frames [Liu and Duyn, 2013] and be able to re-create the seed-correlation map.

As a proof of concept, we first derived the seed-connectivity map for all 17 networks of the Yeo functional atlas, taking the union of all frames that survived the threshold (top 15%) in each of the seed belonging to each network. Figure 6.2 displays all 17 networks corresponding to the solution of the GSR. In addition to the well-known patterns that reveal functionally coupled regions during rest, as displayed in Figure 2.2 in Chapter 2, we can also observe the different WM tracts that are implicated in the formation of these networks.

Finding the optimal diffusion time

Following the observed correspondence between the results of GSR and the graph diffusion, we implemented the voxel-wise graph diffusion in all fMRI volumes from all subjects using different τ values. The goal is to find the optimal value of the diffusion time which gives a solution that conforms best with the solution of the GSR formulation. A higher value of τ corresponds to a longer amount of time that the GM signals are diffused, which can also be interpreted as signals being given more time to reach farther distances. In contrast, a lower value of τ results to less dispersed signals, and thus reaching only shorter distances. Using all the 4D sets of interpolated volumes corresponding to different $\tau-$values, we obtained different sets of seed-connectivity maps. We expect that for low τ, signals are not readily diffused and only reaches the local neighborhood. If we, therefore, assess the similarity of the resulting seed-connectivity maps, within the GM and within the WM, with respect to the ground truth (*i.e.,* the ones derived using GSR), GM signals should show high similarity in low-values of τ. In contrast, there is an optimal value of τ for which the structural mediation of GM interactions starts to emerge within the WM.

Figure 6.3 displays separately, the similarity matrices within the GM and the WM, between the seed-connectivity maps obtained using the diffused signals and the GSR. The results verify the hypotheses laid out above. In particular, a phase transition behavior can be observed in

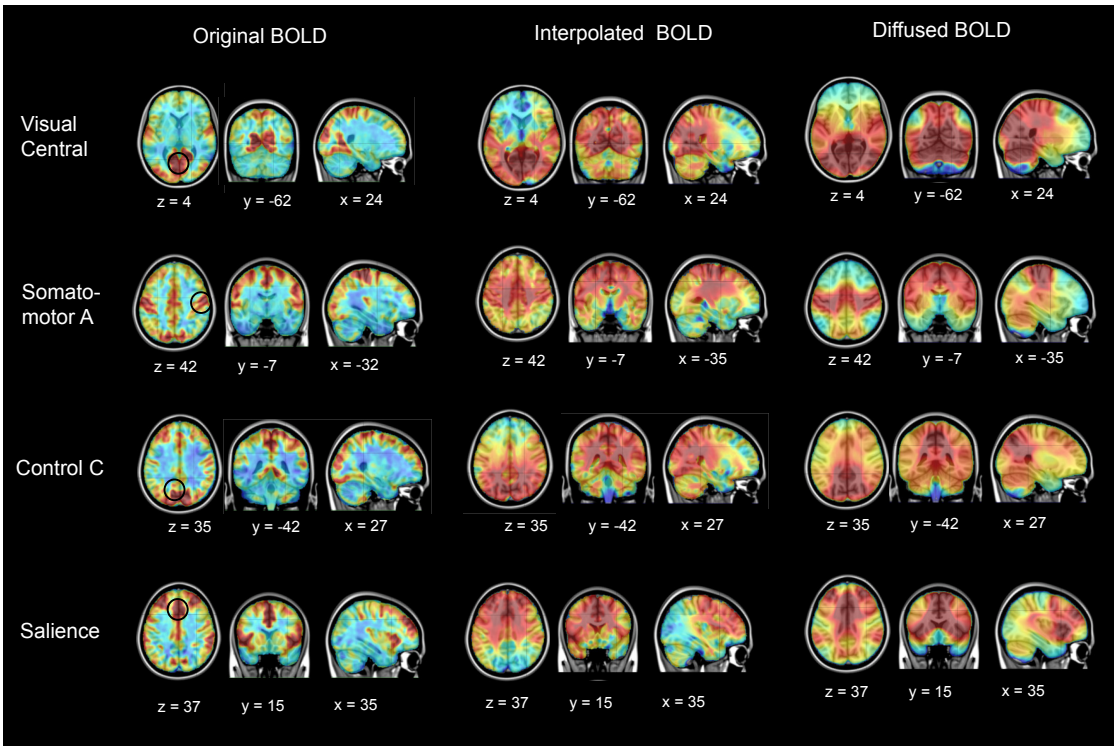

Figure 6.4 – **Seed-average connectivity maps for representative ROIs.** Comparison of the extracted seed-connectivity map in different ROIs using (1) original BOLD data, (2) interpolated BOLD using GSR and (3) diffused BOLD at the optimal diffusion time (Visual Central – 49 ; Somatomotor A – 58 ; Control C – 67; Salience – 74). Seed regions are indicated with a black circle. The associated networks where the ROI belongs to are also displayed in the label on the left.

$\tau-$values between 40 to 90. This regime is indicated by the black arrow in the figure. We then define the optimal diffusion time for each region as the point where correlation starts to be greater than 0.5.

Seed-connectivity map of the original, interpolated, and diffused BOLD data

Now that we have obtained the optimal diffusion time lengths for each ROI, we can now compute the seed-connectivity of the diffused BOLD data and compare it to the seed-connectivity obtained using GSR. Figure 6.4 displays the resulting seed-connectivity maps for representative ROIs extracted using the (1) original BOLD data, (2) interpolated BOLD data using GSR, and (3) diffused BOLD data of different optimal diffusion time lengths. The associated networks where the chosen ROI belongs to are also indicated. The resulting seed-connectivity map of the original BOLD resembles well-known associated networks from the Yeo functional atlas (see also Figure 2.2 in Chapter 2 for comparison). Meanwhile, visual inspection of the connectivity maps obtained using interpolated BOLD and using diffused BOLDs also reveals a high resemblance between the two. It can also be observed that the connectivity maps of

the interpolated BOLD generally display stronger signals in the WM and the GM. In contrast, the GM in the connectivity map of the diffused BOLD is more smooth. This outcome is a consequence of the implementation of the diffusion kernel which acts as a spatial blurring or a low-pass filter. High-frequency components are attenuated, and as such, local details in the GM are also reduced.

Mapping of spectral diffusion rate to vertex diffusion time and its equivalence to spatial distance

We observed the optimal diffusion time lengths to be different for each region. This measure can be interpreted as reflecting the disparate but indirect measure of the anatomical range between different ROIs. The next goal is to obtain the equivalent measure of the diffusion time in terms of physical distance in millimeters.

As discussed in the methods, graph diffusion can be implemented in two ways: (i) vertex domain diffusion and (i) spectral-domain diffusion. The former utilizes the Laplacian matrix as the difference operator in the implementation of the finite difference method when approximating the numerical solution of Eq. 6.2. Meanwhile, the latter approach is equivalent to the low-pass filtering of the signals in the spectral domain. The implementation of approach (i) involves an iterative action of the Laplacian matrix to the signal until the steady-state solution is achieved. Since the voxel-level brain grid uses a $3 \times 3 \times 3$ neighborhood scheme, the difference operator only affects its nearest neighbor at each iteration, and as such, it follows that for every time-point (or iteration), a source signal from one voxel diffuses only to its neighboring voxel. As such, 1 iteration in vertex domain diffusion is equivalent to the voxel resolution of the DW-MRI data – in this case, 1 time-point = 1.25 mm.

The next challenge is to determine the mapping between the vertex domain diffusion iteration and the spectral diffusion time associated with approaches (i) and (ii), respectively. To do this, we performed the signal diffusion taking each of the 17 networks ROI of the Yeo functional atlas [Thomas Yeo et al., 2011] as source signals and implemented the two approaches of graph diffusion. We compared the resulting interpolated maps using cosine similarity as a metric. Shown in Figure D.1 are the similarity matrices for example networks averaged across all subjects. We observe that there is a one-to-one correspondence between t and τ, starting from specific points, for instance, indicated by the white arrow in the similarity matrix corresponding to the Salience network.

Anatomical range of functional network interactions organizes the cortex according to behaviorally-relevant organization

Using the mapping discussed above, we converted the optimal diffusion time lengths observed in each region (Figure 6.3) to their equivalent spatial distance in millimeters. These values are projected to each ROI of a cortical surface which is displayed in Figure 6.5. We observe

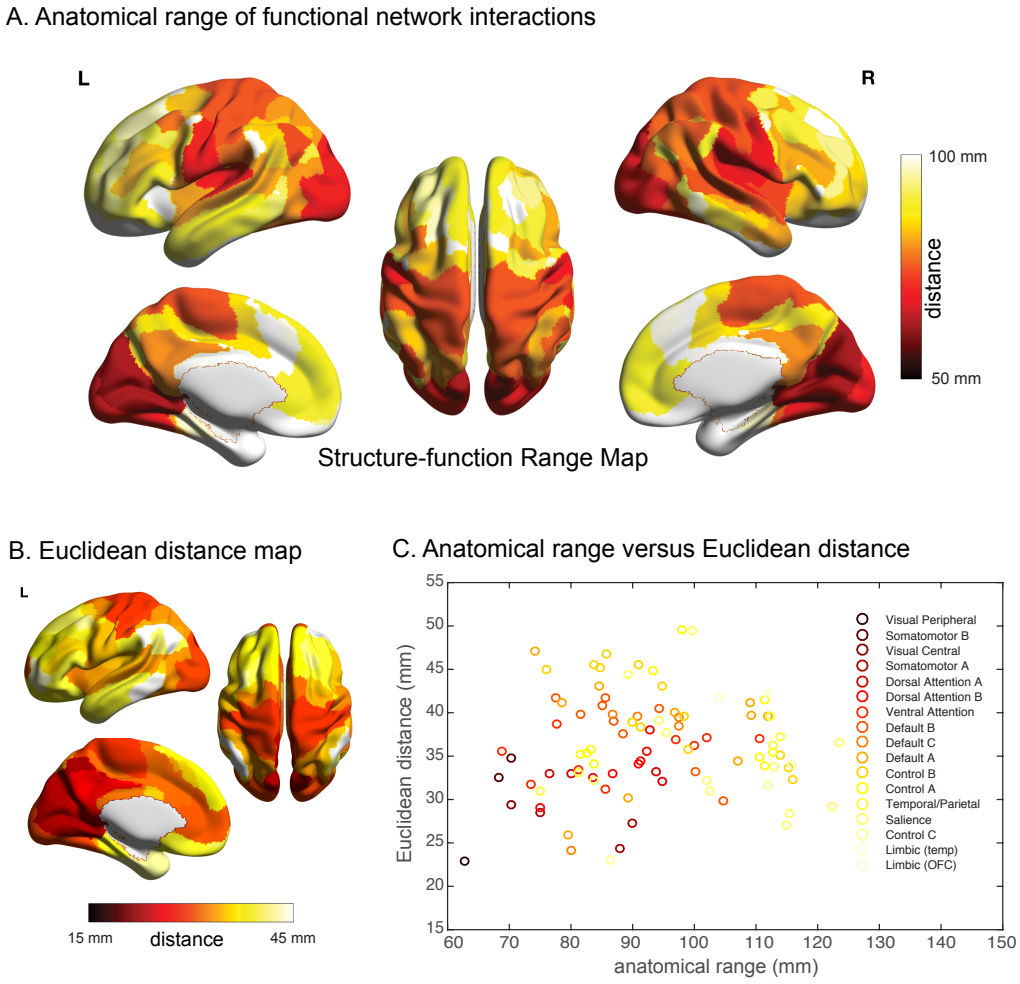

Figure 6.5 – **Anatomical range of functional interactions and its similarity to the Euclidean distance map.** Cortical surface projections of the (A) anatomical range derived using graph diffusion and using (B) Euclidean distance following the general approach proposed by [Oligschläger et al., 2017]. (C) Mean Euclidean distance measures within the 97 regions of Yeo functional network atlas plotted with respect to anatomical range. Colors are ordered according to increasing mean anatomical range per network.

a macroscale organization that divides the cortex into short- and long-distance functional connectivity patterns. Regions belonging to the primary visual and sensorimotor areas consistently display the shortest distances whereas heteromodal areas responsible for high-order cognitive functions such as the control networks, temporal/parietal, and the limbic regions have a more distant anatomical range of connectivity.

To compare with other metrics derived using physical distance measures, we also derived the cortical functional maps for each subject following the approach used by Oligschlager and colleagues [Oligschläger et al., 2017] using Euclidean distance as a metric. Figure 6.5(B) displays projected distances averaged in each region of the Yeo functional atlas. At first glance,

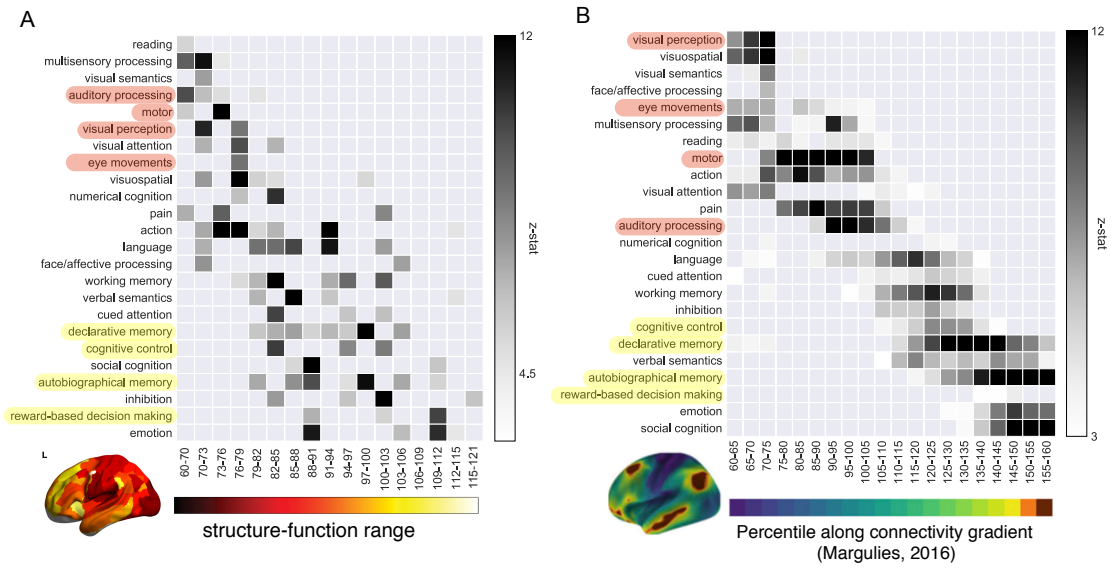

Figure 6.6 – **Anatomical range of functional network interactions reveals behaviorally meaningful organization.** (A) Neurosynth meta-analysis of the ROIs arranged in increasing order according to their anatomical range. (B) Reproduced figure from the analysis conducted by [Margulies et al., 2016].

the Euclidean distance map shows consistent organization with the structure-function range map derived from characteristic diffusion time lengths, especially the low values in primary and sensorimotor regions. A key distinction, however, is the overall pattern around the midline. It can be observed that regions belonging to the default networks do not show substantial differences from primary and sensorimotor areas. Moreover, the overall distance measures are also smaller (*i.e.*, 15 to 45mm) than the values we observed in the structure-function range map (*i.e.*, 60 to 120mm). This range observed using Euclidean distance is also smaller than the range of values observed using geodesic distance as a measure [Oligschläger et al., 2017].

Figure 6.5(C) shows the scatter plot of the Euclidean distance measures in each Yeo functional ROI with respect to the anatomical range derived using graph diffusion. A gradient can be hinted revealing the tendency of regions belonging to the visual, somatomotor, and attention networks to be at the far left of the plot (red); meanwhile, regions from control networks, temporal/parietal, salience, and the limbic networks can be observed in the far right (yellow circles). Midrange values comprise of regions from the default mode networks.

Finally, to verify the link to behavior, we performed a Neurosynth meta-analysis of the 24 topical terms used by [Margulies et al., 2016] and [Preti and Van De Ville, 2019]. Figure 6.6 shows the outcome side-by-side with the result reproduced from [Margulies et al., 2016]. Using very different inputs on the meta-analyses – one based on principal gradients derived from functional data alone [Margulies et al., 2016], one based on structural-function decoupling index [Preti and Van De Ville, 2019], and in our case, one based on the anatomical range

of functionally coordinated brain areas – we observe very consistent patterns of functional differentiation. In particular, regions related to auditory processing, motor control, visual perception, and eye movements are associated with low values of the anatomical range. In contrast, at the higher end of the anatomical range spectrum are areas related to the more complex cognitive manifestations, such as declarative memory, cognitive control, autobiographical memory, and reward-based decision making.

6.2.4 Discussion

General Findings

Here, we utilized a canonical model of signal diffusion borrowed from spectral graph theory to derive the range of anatomical connectivity between functionally coupled brain regions. We first demonstrated that the previously proposed WM interpolation of GM BOLD activity based on GSR [Tarun et al., 2020b] gives a stable solution to the graph diffusion model. We applied a seed-connectivity analysis to the resulting interpolated volumes and extracted 17 whole-brain structure-function networks that reveal the structural mediation of well-known RSNs. Different WM tracts were observed to be recruited in the formation of these networks. Then, we diffused GM signals into the WM at different lengths of diffusion times. We extracted their corresponding seed-connectivity maps and compared them to the seed-connectivity maps obtained using the GSR approach. In doing so, we found region-specific optimal values of diffusion time for when the structural mediation of GM signals begins to appear. These values are then converted to their equivalent spatial distance and are projected to a cortical surface. The overall cortical pattern shows a macroscale organization that is consistent with previously observed cortical connectivity gradient [Sepulcre et al., 2010, Oligschläger et al., 2017, Margulies et al., 2016, Preti and Van De Ville, 2019], which differentiates brain regions according to the complexity of their functional roles – from low-level primary sensorimotor networks to higher-order association networks.

Characteristic time of spread divides the cortex into short- and long-range connectivity

The anatomical range was shortest in primary sensorimotor regions, while it is much longer in higher-order association areas. Notably, regions of the default networks were observed to implicate medium-length anatomical connectivity. These findings are well in-line with the observed cortical pattern obtained using geodesic distance [Oligschläger et al., 2017], the spatial embedding of functional connectivity [Margulies et al., 2016], and the degree of structural and functional coupling [Preti and Van De Ville, 2019]. Unimodal primary and sensorimotor cortices are known to be the first layer in information processing, as they are mainly responsible for the initial synthesis of sensory input. As it gets further away from the primary cortex, associative functions become more general and the associated networks increase in functional complexity [Oligschläger et al., 2017]. Situated in the middle are regions comprising the DMN, and in particular, the posterior cingulate, which reveals a mid-range

value of optimal diffusion time. In has been found that the DMN subregions show a hybrid tendency to prefer both local and distant connectivity [Sepulcre et al., 2010]. In this work, we have specifically observed that the posterior region of the DMN has a preferentially lower anatomical range compared to the medial prefrontal cortex. This aligns well with previous findings concerning the posterior cingulate acting as a connectional hub [Buckner et al., 2008, Fransson and Marrelec, 2008], but consists of both short- and long-ranged connections. Interestingly, regions of the DMN are also observed to be maximally equidistant along the cortical surface from unimodal regions [Margulies et al., 2016].

More than in higher-order association areas, the structure-function range peaked in the limbic networks (95 mm), which consists of the orbitofrontal cortex and the temporal pole. Consistently, these regions have been observed to be the least connected nodes in previous studies using graph theory [Achard et al., 2006], and as such, are considered to be topologically more peripheral. Moreover, these regions have also been found to exhibit the greatest levels of diversity when investigating their connection length distributions [Betzel and Bassett, 2018]. These findings, in particular, the observed behavioral hierarchy associated to the anatomical range of FC, are very consistent with Mesulam's well-regarded theory of functional differentiation, which orders different brain areas according to the complexity of their functional roles, from the most highly differentiated primary sensorimotor areas to the least differentiated limbic structures [Mesulam, 1998].

Furthermore, the cortical brain organization derived from the anatomical range of FC is also observed to be consistent with the measured peak of fiber densities [Bajada et al., 2019], for example, fiber lengths in the motor cortex and in Broca's area are about 50 mm and 80 mm, respectively. Meanwhile, the primary visual cortex has been shown to recruit only a narrow band of short-range fibers, from up to 40 mm. Meanwhile, the range of distance observed using the euclidean distance is much lower than the values we observed.

Lastly, our results are also consistent with the observed patterns of variability across individuals in FC [Mueller et al., 2013] and SC [Chamberland et al., 2017]. When we computed the standard deviation of characteristic time of spread across all subjects and projected them in a cortical surface (see Figure D.2). We also found that the variability of the anatomical range of FC across individuals also scales proportionally with the patter of the structure-function range map.

Methodological Perspective

The overall aim of the introduced graph diffusion framework is to extract the optimal characteristic time of spread for each region of the brain and use it as a proxy to delineate areas that are dominated by local connectivity from those governed by longer and more distant inter-areal connections. Several methods that attempted to establish the anatomical range of different brain regions have already been proposed before. Early studies used Euclidean distance as a proxy to define the metabolic cost of maintaining short- and long-ranged inter-areal connectivity. This is interpreted as one of the graph measures in studies of network theory

that are applied to either functional or anatomical brain connectivity [Bassett and Bullmore, 2017, Bullmore and Sporns, 2012, Achard et al., 2006]. Euclidean distance and geodesic distance are also used outside the context of network theory by relating regions of well-known RSNs and directly observing their overall spatial organization projected in the cortex [Sepulcre et al., 2010, Oligschläger et al., 2017, Oligschläger et al., 2019]. Moreover, tract-tracing studies have also been performed in macaque monkeys to quantify cortico-cortical connections [Markov et al., 2011]. Another family of method that is used in profiling cortical organization is surface-based projections of WM fiber connectivity [Huang et al., 2011, Tozer et al., 2012, Padula et al., 2017], which are used to assess changes in brain development [Ouyang et al., 2016].

A similar theme in these studies is the use of either only functional data from fMRI or anatomical data from DW-MRI to describe the spatial organization of the cortex, and very rarely on integrating the two modalities. Recently, more and more studies have emphasized the strong link between brain structure and function, suggesting that jointly incorporating these two measures together can provide richer and more accurate information about the brain's topological organization. Region-wise GSP approaches have also started to elucidate the importance of structure-function link in supporting coordinated cognition and behavior [Preti and Van De Ville, 2019, Medaglia et al., 2018].

In this work, we proposed a novel framework that takes into account, not only the functional information, but also the contribution of SC in determining the anatomical range of functionally coupled brain areas. We incorporate all measures of GM activity at the voxel level and used region-wise parcellations to identify regions we want to investigate. A clear advantage of the proposed technique is the direct integration of GM BOLD activity and its diffusion onto the WM which is constrained by the SC that is explicitly encoded in the brain grid. Moreover, the use of diffusion kernels gives a two-fold interpretation of cortical organization, one based on the canonical idea of signal propagation, and one based on graph filtering. The former gives a proxy for distance through measures of characteristic length of time that are required for signals to traverse WM fiber bundles from one region of the cortex to another. Meanwhile, interpreting the diffusion kernel as a graph filter translates to directly depicting physical distances that arise from the physical topology of the voxel-level brain grid. It has been shown that the Laplacian eigenvectors or *eigenmodes* provide a basis or building blocks for representing cortical acitvity [Atasoy et al., 2016]. Recent findings also demonstrate that high-frequency eigenmodes in region-wise graphs capture higher-level cognitive processes, while primary sensory areas are expressed using mostly low-frequency components [Medaglia et al., 2018]. Similarly, low-values of τ allows the use of higher-frequency components which outlines local- and short-distance WM connectivity while high values of τ express only low-frequency components that are made up of the more global, long-distance tracts.

The use of graph diffusion models in neuroimaging have mostly been applied to clinical studies. The first study utilized a random-walk graph diffusion model to demonstrate the use of graph Laplacian eigenmodes in the context of activity propagation [Raj et al., 2012]. Their model was able to predict distinct and persistent modes that recapitulate known patterns

of dementia. This was later refined to obtain FC estimates based on SC [Abdelnour et al., 2014], and to investigate the spread of hyperactivity in epilepsy [Abdelnour et al., 2015]. In 2017, the same group extended its application to quantifying the impact of various localized white matter pathologies on cognition and behavior [Wang et al., 2017]. They observed that the eigenmodes obtained from the anatomical connections of healthy subjects are strongly conserved in contrast to the significant deviations observed for improperly developed brains. Recently, the same model was leveraged to analytically demonstrate how functional brain activity can be derived from SC Laplacian eigenspectra [Abdelnour et al., 2018]. All of these works were implemented in region-wise graphs and were motivated by different research questions. Nevertheless, these studies provide a strong foundation for incorporating graph spectral approaches not only in bridging FC and SC, but also in its potential to understand the progression of neurological and cognitive disorders. It is therefore a very promising avenue for a more global approach in probing brain dynamics across different species and modalities.

Technical limitations and future perspectives

An obvious limitation of the study is the reliance on region-wise parcellations to investigate the characteristic lengths of diffusion time in different areas of the brain. The use of parcellations does not fully take advantage of the highly-resolved nature of the proposed voxel-level interpolated volumes. For future studies, it will be useful to evaluate more refined parcellations consisting of up to 1000 regions, or an even better approach would be to extend the analysis to voxel-level resolution. Lastly, an important aspect of this work that was not discussed further is the structural mediation in different RSNs. Case by case investigation of the different WM tracts that are implicated in their formation could be of interest. This could be especially useful in identifying potential neuroplastic responses occurring in populations with brain malformations [Owen et al., 2013].

6.2.5 Conclusion

We proposed a new framework that characterizes the anatomical range of functional network interactions. Our findings demonstrate that the spatial ordering of large-scale functional brain networks are not random, and follows an overarching organization that is constrained by the anatomical substrate. This work is one of the first studies that leveraged recently introduced voxel-wise connectomes, which expand the use of the GSP repertoire in the context of functional brain imaging.

7 Summary & Outlook

In this PhD work, we have introduced three methodological developments that are useful for combining diffusion and functional MRI data into a single integrated framework to analyze the interplay between brain structure and function. Our main contributions are as follows:

1. We developed a novel design to construct voxel-wise brain graphs that encode local structural connectivity from diffusion-weighted MRI. Using this approach, we provided a direct and quantitative measure of how much information in functional MRI can be captured by the global brain structure obtained using diffusion MRI.

2. We proposed a new framework that enables observing key white matter structures that link temporally coherent gray matter regions. This allowed us to extract structure-function networks that reveal the collective mediation of white matter pathways across short and long-distance functional connections.

3. We advanced the state-of-the-art graph signal processing (GSP) methods for functional brain imaging by extending region-wise analyses to voxel-level perspective, thereby offering new research avenues for the joint analysis of brain structure and function. As a first major application, we demonstrated how the anatomical range of functional network interactions follow a behaviorally relevant macroscale gradient that differentiates brain regions according to the complexity of their functional roles.

In addition to the above's methodological contributions, we also characterized brain dynamics through network-level representations of brain function during rest and across non-rapid eye movement (NREM) sleep.

In what follows, we summarize our main findings and discuss potential future research directions.

7.1 Summary of findings

Network-level representations of brain function during rest and sleep. We applied an advanced method that relies on transient brain activity (or moments of activity changes in the BOLD activity) to uncover spatially and temporally overlapping large-scale functional brain networks that occur during wakefulness and across NREM sleep. Our network-level representations of brain function revealed new spatial patterns covering regions that support the physiological organization of sleep and arousal. In particular, we observed visual-sensory areas together with negative activations in the basal forebrain to predominate during wakefulness and NREM stage 1. We also found evidence for changes in neocortical-cerebellar interactions as a function of sleep depth. More generally, the persistence in time of individual brain maps across different sleep stages was significantly altered. Finally, and most importantly, contrary to previously observed strong functional dissociations that accompany the deepening sleep, we instead observed a surge of network activity and cross-network interactions during NREM stage 2, followed by an abrupt decrease of global network communication in NREM stage 3. Thus, by providing a spatiotemporal yet parsimonious description of brain organization, we were able to capture new features of information integration of consciousness during sleep, and provide concrete evidence for the presence of unstable yet distributed global synchronization in NREM stage 2. These results are in-line with previous literature on brain function in different brain states, but also significantly advance current understanding of brain dynamics across NREM sleep. We also explored how changes in brain dynamics can be reflected in terms of their similarity to structural connectivity (SC) from DW-MRI. We observed that, as measures of FC break down in NREM stage 3, its similarity with SC increases. This finding reproduces previous work of [Tagliazucchi et al., 2016], demonstrating the persistence of primary and sensorimotor regions to dominate NREM stage 3.

Dynamics of functional network organization revealed using methods from graph learning. Graphs are versatile representations of data. Typically, relationships between brain regions are described with a graph structure that is predefined to be either structural or functional connectivity. In this section, we applied, for the first time, a recently introduced Graph Laplacian Mixture Model [Petric Maretic and Frossard, 2020] to extract meaningful repeating network patterns in functional brain data. We observed that the probability of each network to occur is consistent with the timings of the task paradigm, and the spatial patterns associated to each epoch are compatible with previously known implicated brain areas obtained using classical regression analyses. We also observed that the extracted graph matrices show high similarity with structure and that this FC-SC link gives a hint of behaviorally relevant macroscale gradient that is consistent with previously observed cortical organization.

Voxel-wise structural brain graphs: design, construction, and application to fMRI. We have extended the spatial resolution of our current analysis into a voxel-level perspective by introducing a new method for jointly integrating brain structure and function. We proposed two designs for constructing high-resolution voxel-level brain graphs based on diffusion tensor imaging and diffusion orientation distribution functions. The graph edge weights are

designed to encode the interplay between the orientation of diffusivity at each voxel and its corresponding magnitude. Spectral analysis of the brain graphs' Laplacian eigenspectra reveals two important findings. First, we observed that despite the huge dimensionality of the proposed brain graphs (*e.g.*, 750,000 voxels), the majority of the structural information is already captured by the Laplacian's lowest frequency eigenmodes. Second, by computing the energy spectral density of functional MRI during rest and task, we observed that about 85% of the information content can also be well approximated by a small subset of low-frequency components of the graph Laplacian spectra. This latter finding gives, for the first time, a direct quantitative measure of the extent to which brain function is captured by the underlying structure. Finally, we discussed the huge potential of the proposed technique not only in characterizing brain structure at a high spatial resolution but more importantly in integrating brain structure and function.

Structural mediation of gray matter activity. We bridged the gap between brain structure and function by proposing a more integrated approach for observing the structural patterns that are implicated in the formation of large-scale functional brain networks. Leveraging on the proposed voxel-wise brain graph in the previous Chapter to guide the process, we interpolated functional signals from the gray matter onto the white matter to obtain the structural links that connect temporally coherent cortical areas. To do this, we utilized a graph signal recovery (GSR) framework that minimizes a cost function which balances the solution between two constraints: (1) retaining the gray matter signals fixed (*i.e.*, ideally unchanged by the process) and (2) making sure that the resulting signals are smoothed over the graph structure. We found that the proposed framework successfully enabled the observation of key white matter structures that link co-activating gray matter areas at a single-frame resolution. The result of the interpolation is therefore a set of 4D fMRI volumes that also include interpolated white matter activity. We performed a co-activation pattern (CAPs) analysis on the interpolated data using the posterior cingulate cortex (PCC) as a seed, a key region of the default mode network (DMN), to obtain time-varying functional connectivity maps. These patterns were enriched through our approach by refining its field of view to include the WM structural mediation of GM networks. The observed CAPs have distinct structure-function patterns, each of which are differentially observed during tasks, demonstrating plausible structural mechanisms that govern the functional switching of the brain between task-positive and -negative DMN subsystems.

The anatomical range underneath cortical functional organization. We proposed a novel framework that utilized a canonical model of graph signal diffusion to define the anatomical range of functional network interactions. We investigated different diffusion time lengths necessary for signals from the gray matter to be diffused entirely onto the white matter. This diffusion time is taken to be proportional to the range of collective structural mediation of temporally coherent gray matter areas. We observed that each region of the brain is characterized by different values of the anatomical range that is consistent with previously observed white matter fiber lengths. Moreover, the meta-analysis conducted on the anatomical range of functional network interactions reveals a behaviorally-relevant macroscale organization in

the cortex that meaningfully differentiates brain regions associated to low-level and high-level cognitive functions.

7.2 Outlook

7.2.1 Further applications to neuroscience

On the persistence of primary and sensorimotor networks to occur during NREM stage 3 and their relation to anatomy. It has previously been observed that the similarity between FC and SC in primary sensory cortices increases during deep sleep, more than in wakefulness [Tagliazucchi et al., 2016]. Meanwhile, higher-order associative cortices display the opposite trend. These findings are well-in line with what observed. In particular, we have shown in Section 3.2 through the normalized relative cumulated durations (RCD) of iCAPs, that the primary and sensorimotor networks dominate activity during NREM stage 3, while higher-order executive control and limbic networks displayed the least RCD. This order is observed to completely opposite of the trend observed during wakefulness. Moreover, we have also observed in Section 3.2.5 that FC is generally decreased while its correlation to SC is increased in NREM stage 3. Taken together, these findings suggest that the brain retains low-level information processing during the deep sleep and that functional connectivity in regions essential for conscious processing breaks down.

An interesting next step would be to investigate whether we can observe these alterations in SC-FC link by studying the structural mediation and the anatomical range of functional interactions across the different NREM sleep stages. We already know from previous literature [Boly et al., 2012] and from our own findings that long-ranged connectivity generally breaks down with increasing sleep depth, *e.g.*, decoupling of the DMN into posterior and anterior regions. It would be interesting to confirm, following these prior findings, whether the overall anatomical range of FC in the cortex should also decrease.

Dynamic FC analyses on interpolated volumes to reveal the structural mediation of well-known resting-state networks. The interpolated white matter volumes entail additional informative voxels that extend beyond the gray matter. It is therefore particularly interesting to apply classical dynamic FC analyses [Preti et al., 2017] on these volumes, such as sliding-window correlation, and independent component analysis [Beckmann et al., 2005], to be able to probe functional brain dynamics at a much larger scale. For example, we presented in Chapter 6 the 17 structure-function networks corresponding to the Yeo functional atlas, each of which reveals key white matter structures that mediate the interactions of the regions comprising these networks. Although it was presented, the characterization of different white matter structures that are implicated in the formation of each network was not discussed in detail. Just as how we have explored the mean signals of the interpolated co-activation patterns (in-CAPs) within major white matter fiber bundles in Section 5.2.3, it would also be interesting to identify which white matter tracks are largely recruited whenever a specific

resting-state network is activated.

Applications of the developed methods to clinical populations. The proposed methodologies in this thesis altogether provide new research avenues for a deeper understanding of brain structure and function, extending state-of-the-art graph harmonic approaches from region-wise analysis [Atasoy et al., 2016] to a finer-grained voxel-level perspective. These methods particularly open a promising perspective for investigations in patients with structural brain deficits, as the need for cortical parcellation and tissue segmentation are eliminated. This is also in light of the recent evidence that region-wise brain eigenmodes provide a compact and robust representation of structure in health and disease [Wang et al., 2017]. A first step towards this analysis is to characterize their brain graph's Laplacian eigenspectra and use scalar GSP measures (*e.g.* total variation) to delineate structural differences between the two groups. More interestingly, applying the proposed white matter interpolation to subjects born with brain malformations (*e.g.,* callosal agenesis) but nevertheless display intact brain function could be very beneficial in identifying neuroplastic responses occurring in the clinical population.

Use of voxel-level brain graphs for individual fingerprinting and observing neuroplasticity in longitudinal data. The amount of structural variability observed in the lower-end spectra of voxel-wise connectomes, as demonstrated in Chapter 4, suggests their potential in identifying groups and individuals. Eigenvalues of voxel-wise graphs that encode cortical morphology have been shown to provide distinguishable features that delineate gender and hemispherical asymmetry [Maghsadhagh et al., 2019]. Identifying individuals through functional connectomes [Finn et al., 2015, Amico and Goñi, 2018] has also recently been introduced to structural connectomes [Mars et al., 2018]. Moreover, as voxel-wise brain graphs can be used to characterize local voxel-to-voxel brain structure, it can also be used as a more targeted approach to study neuroplasticity in longitudinal data.

Relationship of SC-FC link to behavior We have presented a new set of methodologies that could be useful in integrating brain structure and function, and in our latest work on GSP application discussed in Chapter 6, we have demonstrated that the anatomical range of FC follows a behaviorally-relevant macroscale gradient. An important advance would be to investigate whether various voxel-level GSP measures, such as the total variation and energy content extracted from task and rest-state data, can be used as bio-markers for identifying not only the individual task performance, but also their general behavioral profiles.

7.2.2 Methodological Extensions

Use of population-averaged group-level eigenmodes to define the graph Fourier basis for future GSP analyses. The biggest limitation in our work was the lack of diffusion data that could complement our analyses of brain dynamics during sleep. It would be interesting to explore the structural mediation of human brain activity also in different brain states. Without diffusion data, however, this can not be done. The next advance would be to determine

whether high-resolution population-averaged eigenmodes extracted in Chapter 4 from the HCP data retain enough local structural information to be able to perform GSP operations on available fMRI data that lack complementary diffusion data.

Improving white matter BOLD activity via graph signal recovery. Recently, more and more studies have started to demonstrate the possibility of capturing functionally relevant information from the white matter BOLD [Peer et al., 2017, Ding et al., 2018]. This is particularly interesting in light of the controversy surrounding white matter BOLD activity [Logothetis and Wandell, 2004, Gawryluk et al., 2014], which is in part, due to the well-established differences in hemodynamic responses of gray matter and white matter [Fraser et al., 2012]. An interesting extension of our proposed white matter interpolation of fMRI would be to recover weak signals in the white matter. At its current form, the proposed graph signal recovery framework removes all white matter BOLD contributions and retains only the gray matter signals fixed. We can therefore redefine the cost function to include also the original white matter signals.

On the interpretation of graph diffusion as graph filter. The interpretation of the use of graph diffusion to delineate short- and long-ranged functional connections is two-fold. First, as a canonical model of signal diffusion, distance can be interpreted as proportional to the time it takes for implicated white matter structures to link temporally coherent gray matter areas. Second, in the context of GSP, the graph diffusion essentially acts as a low-pass filter, where the filter is equivalent to a decaying exponential parameterized by the diffusion time. Recall that in Chapter 4, we demonstrated that the low-frequency eigenmodes of the Laplacian matrix capture global structural connectivity, whereas high-frequency components capture local, and finer-grained connectivity. Therefore, high-values of diffusion time translates to attenuating high-frequency components and consequently capturing only long-distance white matter structures that link more distant functional connections. Therefore, a potential extension of the work would be to interpret the diffusion kernel and identify the optimal diffusion parameter (equivalent to the average of the ones we observed in Chapter 6) as a low-pass filter. Accordingly, we can then also obtain its complementary high-pass filter to be able to assess not only the spatial patterns that are composed of low-frequency components, but also those that are mostly retaining local and more fine-grained information.

A Supplementary Material for Chapter 3

A.1 Supplementary Materials for Section 3.1

A.1.1 Materials and Methods

Dataset number 1: Sustained sleep experiment

The experimental protocol involved three main phases that were recorded on different days. The first phase was the recording of the structural volumes. The second phase was done in a different session (about 5-7 days later), where subjects were asked to play a face recognition and maze recognition games. After the recording, subjects were taken outside the scanner, were allowed to rest, walk, and eat. At 10 p.m, the subjects were asked to go back to the scanner and were instructed to fall asleep. All participants except for two people were able to sleep. Thirteen participants were able to reach deep sleep (stage N3). The third phase is the memory test phase, wherein participants were asked to perform two days later a memory test assessing their memory for elements of the face and maze games. For the purpose of this work, we limit our study to the sleep session, taking only the recordings when subjects were asked to sleep inside the scanner. The maze and face recordings, as well as the memory tests results will not be included in the analysis.

EEG was recorded in the scanner room simultaneously with fMRI acquisition. The EEG setup included a 64 channels MRI-compatible EEG cap (BrainCapMR ; EASYCAP GmbH, Ammersee, Germany, with 64 ring-type electrodes), two pairs of ECG, horizontal and vertical EOG, and 1 pair of chin EMG (BrainAmp MR plus; Brain Products GmbH, Gilching, Germany). The EEG signal was referenced online to FCz. Electrode-skin impedance was kept under 10 kohms, in addition to the 5-kohm built-in electrode resistance. The EEG signal was digitized at a 5000 Hz sampling rate with 500 nV resolution. Data were analog-filtered by a band-limiting low-pass filter at 250 Hz (30 dB per octave) and a high-pass filter with a 10 s time constant corresponding to a high-pass frequency of 0.0159 Hz.

Dataset number 2: Sleep onset experiment

Ten healthy subjects (4 female) aged 18-30 underwent simultaneous EEG-fMRI recording. All subjects have no history of neurological or psychiatric disease, have declared no previous or current use of psychoactive drugs, and are non-smokers. All participants were allowed to sleep up to 20 minutes inside the MRI scanner. Once Stage 1 sleep was established, participants were awoken and were asked to report what they could remember prior to awakening. Verbal responses were digitally recorded for later transcription.

The preprocessing of the fMRI volumes follow the same standard pre-processing protocol used in the sustained sleep dataset. Functional volumes underwent realignment and co-registration, followed by the regression of 11 parameters (white matter, cerebrospinal fluid, 6 motion parameters, constant, linear, and quadratic drifts). The first 10 volumes were also discarded to achieve steady-state magnetization of the fMRI data. After which, the co-registered T1 images were segmented to obtain probabilistic gray matter mask, thresholded at 0.3.

Deconvolution of fMRI signal via Total Activation

The total activation (TA) framework is the process of deconvolving fMRI time series to detect time points where there are significant changes in brain activity. Fig. SI 1 illustrates the TA framework. At each voxel i, fMRI signals $y(i,t)$ are modelled as a block-like activation activity-inducing signal $u(i,t)$ convolved with the hemodynamic response function $h(t)$, and corrupted by a Gaussian noise (i,t):

$$y(i,t) = h(t) * u(i,t) + \epsilon(i,t) \tag{A.1}$$

The denoised fMRI signal, also called herewith as activity-related signal $x(i,t)$ (or in matrix form X), is obtained by running a combined spatial and temporal regularization that aims to minimize the following cost function:

$$\tilde{\mathbf{X}} = \arg\min \frac{1}{2}||\mathbf{Y}-\mathbf{X}||_F^2 + R_T(\mathbf{X}) + R_S(\mathbf{X}). \tag{A.2}$$

The temporal and spatial regularization are given by the following terms:

$$R_T(\mathbf{X}) = \sum_{(i=1)}^{(N_i)} \lambda_T(i) \sum_{(t=1)}^{(N_t)} |\Delta \mathbf{X}(i,\cdot)|$$
$$R_S(\mathbf{X}) = \sum_{(t=1)}^{(N_t)} \sum_{(i=1)}^{(N_i)} \sqrt{\sum_{i\in S} \Delta_{Lap}\{\mathbf{X}(i,t)^2\}}$$

where $\Delta = \Delta_D \Delta_{L_h}$ combines the deconvolution and the differentiation operation, and Δ_{Lap}

defines the 3D second-order difference operator. The spatial and temporal regularization parameters are given by the λ_S and λ_T, respectively. N_i and N_t denote the total number of voxels and timepoints, while $S(i)$ denote all voxels in the neighborhood of voxel i. For more details about the implementation of the TA paradigm, we refer to [Karahanoğlu et al., 2013] and [Farouj et al., 2017], respectively.

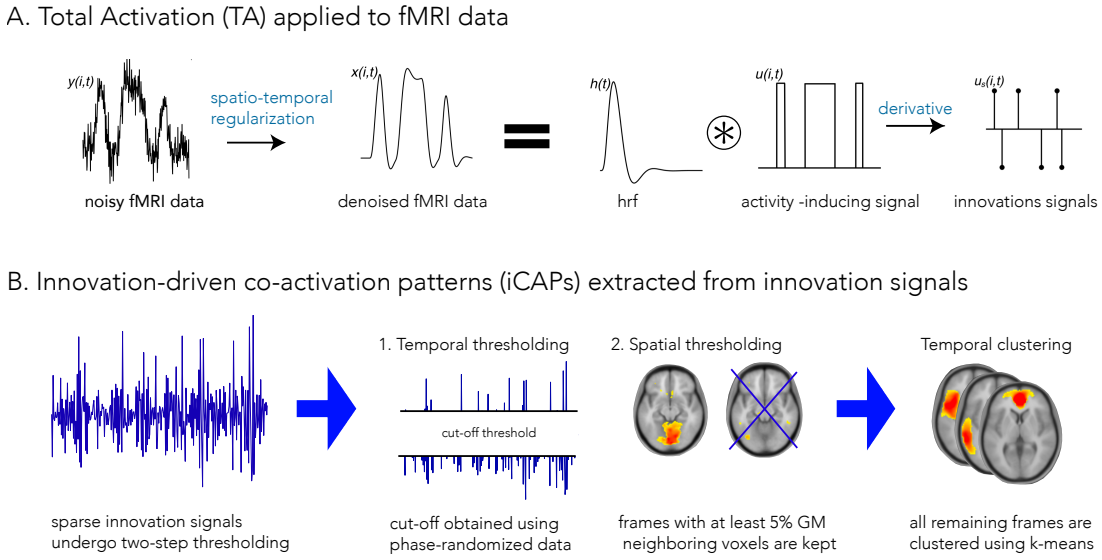

Figure A.1 – **Methodological pipeline of TA-iCAPs.**(A) Noisy fMRI time-courses are denoised using a combined spatio-temporal regularization, followed by a deconvolution from the HRF to obtain block-type activity inducing signals, which is then differentiated to get sparse innovation signals. (B) The resulting innovation signals undergo a two-step thresholding (spatial and temporal), and the remaining frames undergo temporal clustering to extract the iCAPs.

Extraction of innovation-driven co-activation patterns (iCAPs)

After running TA, we extract the sparse innovation signals, $u_S(i,t)$ by computing the temporal derivative of the activity-inducing signals u(i,t), taking only a fraction across brain volumes following a two-step thresholding procedure described in Fig. A.1. The same dataset undergoes phase-randomization, from which we extract the cut-off threshold using the lowest 5th and highest 95th percentile of its corresponding innovation signal. Following the temporal thresholding, we take frames with at least 5% of the total number of gray matter voxels to undergo temporal clustering using k-means. The resulting maps, termed innovation-driven co-activation patterns (iCAPs) are shown in Fig. 2. For the details of the iCAPs extraction, we refer to [Karahanoglu and Van De Ville, 2015].

Temporal Characteristics of iCAPs

Finally, iCAP time-courses for each subject are computed using transient-informed spatio-temporal back-projection of the iCAP maps onto the activity-inducing signals [Zoller et al., 2019]. The clustering of the innovation signals or the transient points for which we observe changes in neural activity allows for extraction of spatially and temporally overlapping spatial maps. This characteristic gives rise to dynamically rich repertoire of functional states, each of which can be explored to observe iCAP-specific temporal characteristics.

The relative total duration pertains to the number of active timepoints in the iCAPs time-courses divided by the length of time that the subject lingered in a particular sleep stage. This measure pertains to how long an iCAP appears in each sleep stage.

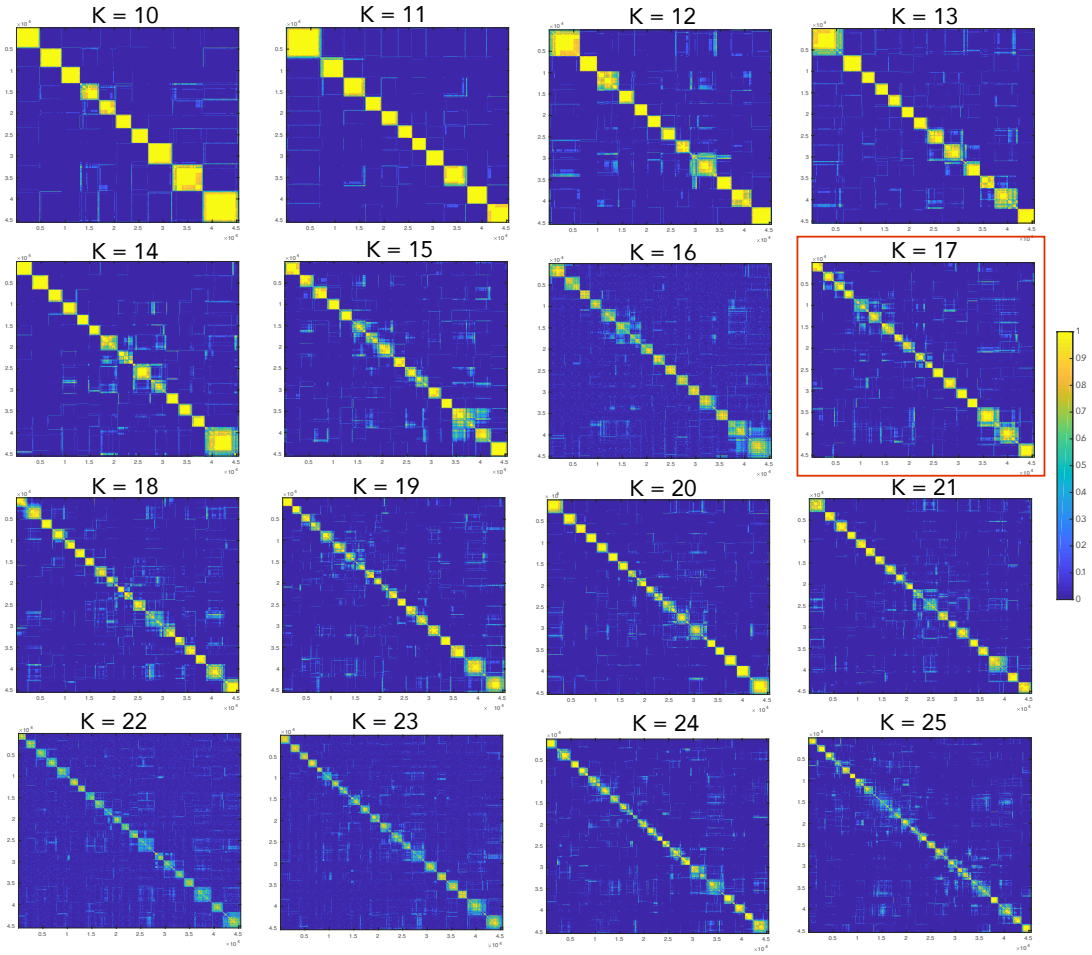

Figure A.2 – **Consensus clustering matrices** for different cluster values $K = \{10, 11, 12, ..., 25\}$. Values in the matrix range from 0 to 1, which indicate the reproducibility of the sampling across multiple runs, with 1 being perfectly re-sampled at all times. We chose K = 17 based on visual inspection and the consensus quality measures displayed in Fig. A.3.

Finding the optimal number of clusters

We used consensus clustering [Monti et al., 2003] to obtain the optimal number of clusters in the concatenated significant innovation frames. The method involved subsampling of the data and multiple runs of the clustering algorithm. The consistency of each frame to be grouped in a similar cluster is monitored through a consensus metric. Fig. SI 2 shows the consensus clustering matrices for all K values evaluated. The cumulative distribution of this metric is displayed in Fig. A.3. We chose K = 17 to be the optimal number of clusters by visual inspection of the clustering matrices, in addition to the shape of the cumulative distributions, and the trend of their corresponding area under the curve (AUC).

Table A.1 – iCAPs functional networks corresponding to the Greicius atlas [Shirer et al., 2012] and regions in the automated anatomical labeling atlas [Tzourio-Mazoyer et al., 2002]. Percentiles indicate the fraction of voxels belonging to a network or region that has a z-score > 1.5. Only networks and regions that have at least 20% surviving percentile are included in the list. iCAP 6 is not included due to the excluded cerebellum in Greicius atlas. Specific regions that are not part of the cerebellum but are co-activated with iCAP 10 (anterior cerebellum) are included

iCAP	Greicius network (%)	AAL Lobe	AAL Region (%)	z-score	voxels
1	Primary Visual (93.12 %)	Occipital	Cuneus L (83.29 %)	2.35	289
	Higher Visual (41.53 %)	Occipital	Calcarine L (72.69 %)	2.41	386
	Ventral DMN (41.06 %)	Occipital	Cuneus R (70.96 %)	2.28	259
	Precuneus (38.27 %)	Occipital	Calcarine R (70.33 %)	2.48	294
		Occipital	Occipital Mid R (53.62 %)	1.99	252
		Occipital	Lingual R (50.94 %)	2.17	271
		Occipital	Occipital Inf L (49.77 %)	2.07	110
		Occipital	Occipital Sup L (48.02 %)	2.03	121
		Occipital	Occipital Mid L (47.02 %)	2.1	363
		Occipital	Lingual L (46.35 %)	2.15	254
		Occipital	Occipital Inf R (44.66 %)	2.03	113
		Parietal	Precuneus R (42.88 %)	2.04	304
		Parietal	Precuneus L (39.09 %)	2.02	310
		Occipital	Occipital Sup R (38.25 %)	1.92	109
		Parietal	Paracentral Lobule R (25.95 %)	1.8	34
		Temporal	Temporal Sup L (21.65 %)	1.78	118
		Occipital	Fusiform R (20.11 %)	2	142
2	Auditory (77.31 %)	Temporal	Heschl R (98.41 %)	2.15	62
	Basal Ganglia (52.46 %)	Limbic	Insula R (96.03 %)	3.12	435
	Anterior Salience (41.40 %)	Limbic	Insula L (93.74 %)	2.89	464
	Posterior Salience (23.94 %)	Temporal	Heschl L (89.86 %)	1.98	62
		Central	Rolandic Oper L (86.75 %)	2.23	216
		Central	Rolandic Oper R (86.52 %)	2.23	276
		Subcortical	Putamen R (86.42 %)	2.29	210
		Frontal	Frontal Inf Tri R (81.98 %)	2.17	314
		Subcortical	Putamen L (77.69 %)	2.18	188
		Frontal	Frontal Inf Oper R (77.27 %)	2.51	238
		Frontal	Frontal Inf Oper L (75.69 %)	2.19	165
		Subcortical	Pallidum R (75.00 %)	2.01	9
		Frontal	Frontal Inf Tri L (49.48 %)	2.04	239

		Temporal	Temporal Sup L (47.71 %)	1.93	260
		Limbic	Cingulum Ant R (44.21 %)	1.74	149
		Subcortical	Pallidum L (43.48 %)	1.96	10
		Frontal	Frontal Inf Orb R (40.23 %)	2.09	138
		Temporal	Temporal Pole Sup R (38.46 %)	2.1	95
		Limbic	Cingulum Ant L (37.53 %)	1.68	137
		Temporal	Temporal Sup R (32.11 %)	1.94	228
		Subcortical	Thalamus R (27.50 %)	1.63	44
		Temporal	Temporal Pole Sup L (27.27 %)	2.03	63
		Frontal	Frontal Inf Orb L (23.97 %)	2.09	93
		Parietal	SupraMarginal R (23.13 %)	1.69	102
		Limbic	Cingulum Mid R (22.46 %)	1.82	117
		Limbic	Amygdala R (21.88 %)	1.69	14
3	Dorsal DMN (76.38 %)	Frontal	Frontal Sup Orb Medial L (96.39 %)	3.51	160
	Anterior Salience (35.23 %)	Frontal	Frontal Sup Medial L (91.98 %)	3.27	436
		Frontal	Frontal Sup Orb Medial R (91.20 %)	3.46	197
		Frontal	Frontal Sup Medial R (91.08 %)	3.14	398
		Limbic	Cingulum Ant R (90.21 %)	3.57	304
		Limbic	Cingulum Ant L (89.86 %)	3.61	328
		Limbic	Cingulum Post L (59.30 %)	1.9	51
		Frontal	Frontal Sup L (57.89 %)	2.3	356
		Frontal	Frontal Sup R (49.35 %)	2.25	344
		Limbic	Cingulum Post R (39.39 %)	1.94	13
		Frontal	Rectus L (35.62 %)	2.37	78
		Frontal	Frontal Mid L (29.92 %)	1.92	289
		Frontal	Rectus R (28.80 %)	2.41	55
		Frontal	Frontal Sup Orb R (24.48 %)	1.85	59
		Frontal	Frontal Sup Orb L (24.44 %)	1.9	55
		Parietal	Angular L (23.59 %)	1.7	67
		Limbic	Cingulum Mid R (20.73 %)	2.05	108
4	Left ECN (85.35 %)	Frontal	Frontal Inf Tri L (94.62 %)	3.5	457
	Language (51.24 %)	Parietal	Angular L (93.31 %)	2.79	265
	Visuospatial (31.20 %)	Frontal	Frontal Inf Oper L (89.91 %)	3.32	196
	Anterior Salience (24.77 %)	Frontal	Frontal Mid L (81.37 %)	2.6	786
		Parietal	Parietal Inf L (66.50 %)	2.69	397
		Frontal	Frontal Sup L (59.84 %)	2.12	368
		Frontal	Frontal Inf Orb L (47.16 %)	2.61	183
		Frontal	Precentral L (43.59 %)	3.2	245
		Frontal	Frontal Sup Medial L (41.98 %)	2.25	199
		Temporal	Temporal Mid L (39.00 %)	2.05	454
		Frontal	Frontal Mid Orb L (33.85 %)	2.67	66
		Frontal	Supp Motor Area L (33.03 %)	2.4	143
		Parietal	Parietal Sup L (20.24 %)	2.29	86
5	Higher Visual (100.00 %)	Occipital	Occipital Inf R (100.00 %)	3.56	253
		Occipital	Occipital Inf L (93.21 %)	3.19	206
		Occipital	Occipital Sup R (84.56 %)	2.83	241
		Occipital	Occipital Mid R (81.49 %)	3.63	383
		Occipital	Occipital Mid L (77.46 %)	3.49	598
		Occipital	Occipital Sup L (76.19 %)	2.69	192
		Occipital	Lingual R (46.62 %)	2.48	248
		Occipital	Calcarine L (46.14 %)	2.5	245
		Occipital	Cuneus R (44.93 %)	2.19	164

		Occipital	Lingual L (39.60 %)	2.58	217
		Occipital	Cuneus L (38.90 %)	2.12	135
		Occipital	Calcarine R (34.21 %)	2.69	143
		Occipital	Fusiform R (30.31 %)	2.4	214
		Occipital	Fusiform L (22.10 %)	2.48	143
7	Right ECN (67.49 %)	Frontal	Frontal Mid R (89.34 %)	3.05	964
	Precuneus (25.98 %)	Parietal	Angular R (81.94 %)	2.96	295
	Posterior Salience (23.67 %)	Parietal	Parietal Inf R (79.42 %)	3.7	274
	Left ECN (23.12 %)	Frontal	Frontal Mid Orb R (71.09 %)	2.66	150
	Anterior Salience (21.15 %)	Frontal	Frontal Inf Tri R (69.19 %)	2.37	265
		Frontal	Frontal Sup R (64.56 %)	2.41	450
		Frontal	Frontal Inf Oper R (62.01 %)	2.28	191
		Frontal	Frontal Sup Medial R (43.02 %)	2.19	188
		Parietal	SupraMarginal R (36.05 %)	2.76	159
		Frontal	Frontal Inf Orb R (34.69 %)	2.38	119
		Frontal	Frontal Sup Orb R (29.05 %)	2.55	70
		Parietal	Parietal Inf L (28.48 %)	1.99	170
		Limbic	Cingulum Ant R (27.30 %)	1.91	92
		Limbic	Cingulum Mid R (25.53 %)	2	133
		Frontal	Frontal Mid Orb L (23.59 %)	1.95	46
		Parietal	Angular L (23.24 %)	1.84	66
8	Auditory (94.91 %)	Temporal	Heschl L (100.00 %)	3.13	69
	Posterior Salience (45.88 %)	Temporal	Heschl R (100.00 %)	3.48	63
	Sensorimotor (28.50 %)	Central	Rolandic Oper R (96.87 %)	3.35	309
		Central	Rolandic Oper L (94.78 %)	2.8	236
		Parietal	SupraMarginal L (74.56 %)	2.6	214
		Temporal	Temporal Sup L (72.48 %)	2.69	395
		Parietal	Postcentral L (59.46 %)	3.38	396
		Temporal	Temporal Sup R (57.61 %)	2.74	409
		Parietal	SupraMarginal R (57.60 %)	2.97	254
		Parietal	Postcentral R (51.89 %)	3.38	343
		Frontal	Precentral R (47.22 %)	2.58	246
		Limbic	Insula R (46.36 %)	2.39	210
		Limbic	Insula L (33.33 %)	2.29	165
		Frontal	Precentral L (22.60 %)	2.44	127
9		Frontal	Rectus L (100.00 %)	3.13	219
		Frontal	Rectus R (100.00 %)	3.05	191
		Frontal	Olfactory R (97.30 %)	2.7	72
		Frontal	Olfactory L (96.00 %)	2.81	72
		Temporal	Temporal Pole Mid L (94.84 %)	2.33	147
		Frontal	Frontal Sup Orb L (94.22 %)	2.47	212
		Frontal	Frontal Sup Orb R (94.19 %)	2.5	227
		Temporal	Temporal Pole Mid R (90.87 %)	2.04	189
		Limbic	Amygdala L (89.66 %)	2.07	52
		Frontal	Frontal Sup Orb Medial L (87.35 %)	2.12	145
		Frontal	Frontal Mid Orb R (86.73 %)	2.54	183
		Frontal	Frontal Inf Orb L (86.34 %)	2.44	335
		Limbic	Amygdala R (85.94 %)	1.94	55
		Frontal	Frontal Mid Orb L (84.62 %)	2.49	165
		Temporal	Temporal Pole Sup L (82.68 %)	2.17	191
		Frontal	Frontal Sup Orb Medial R (81.48 %)	2.13	176
		Frontal	Frontal Inf Orb R (76.68 %)	2.3	263

		Temporal	Temporal Pole Sup R (75.71 %)	2.03	187
		Limbic	ParaHippocampal R (69.61 %)	1.98	197
		Limbic	ParaHippocampal L (69.39 %)	2.03	170
		Temporal	Temporal Inf L (55.67 %)	1.86	447
		Limbic	Hippocampus L (53.60 %)	2.03	119
		Limbic	Hippocampus R (48.23 %)	1.88	109
		Temporal	Temporal Inf R (46.24 %)	1.81	406
		Subcortical	Caudate L (37.43 %)	2.17	70
		Subcortical	Caudate R (36.04 %)	2.13	71
		Occipital	Fusiform L (31.22 %)	1.84	202
		Occipital	Fusiform R (29.18 %)	1.75	206
		Temporal	Temporal Mid L (25.26 %)	1.85	294
10	Sensorimotor (50.65 %)	Occipital	Fusiform R (59.07 %)	2.55	417
		Occipital	Fusiform L (54.71 %)	2.34	354
		Occipital	Lingual R (41.92 %)	2.22	223
		Occipital	Lingual L (39.42 %)	2.19	216
		Limbic	ParaHippocampal R (27.56 %)	2.13	78
		Limbic	ParaHippocampal L (24.90 %)	1.91	61
11	Precuneus (84.88 %)	Parietal	Angular L (97.89 %)	2.86	278
	Ventral DMN (50.52 %)	Limbic	Cingulum Post L (97.67 %)	4.71	84
	Left ECN (30.98 %)	Limbic	Cingulum Post R (93.94 %)	4.28	31
	Dorsal DMN (26.26 %)	Parietal	Angular R (88.61 %)	2.82	319
	Language (23.04 %)	Parietal	Precuneus R (72.78 %)	3.84	516
		Parietal	Precuneus L (64.82 %)	4.13	514
		Parietal	Parietal Inf R (33.62 %)	2.32	116
		Occipital	Cuneus L (32.85 %)	2.83	114
		Limbic	Cingulum Mid L (30.71 %)	3.1	156
		Limbic	Cingulum Mid R (24.95 %)	3.51	130
		Occipital	Cuneus R (24.38 %)	2.29	89
		Parietal	Parietal Inf L (21.44 %)	2.18	128
12	Language (74.17 %)	Parietal	Angular L (86.97 %)	2.68	247
	Auditory (51.16 %)	Temporal	Temporal Sup R (82.68 %)	3.14	587
	Posterior Salience (24.87 %)	Temporal	Temporal Mid R (82.56 %)	3.11	876
		Temporal	Temporal Mid L (77.23 %)	3.05	899
		Temporal	Temporal Sup L (66.06 %)	2.35	360
		Parietal	Angular R (62.22 %)	2.8	224
		Parietal	SupraMarginal L (37.98 %)	2.44	109
		Parietal	SupraMarginal R (31.29 %)	2.51	138
		Parietal	Parietal Inf R (21.45 %)	2.11	74
13	Primary Visual (100.00 %)	Occipital	Cuneus L (97.69 %)	3.82	339
	Precuneus (60.79 %)	Occipital	Cuneus R (89.32 %)	3.63	326
	Ventral DMN (24.59 %)	Occipital	Calcarine R (81.10 %)	4.06	339
		Occipital	Calcarine L (68.55 %)	4.05	364
		Occipital	Lingual R (56.95 %)	3.06	303
		Occipital	Lingual L (49.82 %)	2.88	273
		Parietal	Precuneus R (48.94 %)	2.88	347
		Parietal	Precuneus L (36.19 %)	2.68	287
		Occipital	Occipital Sup L (34.13 %)	2.31	86
		Limbic	Cingulum Post L (31.40 %)	1.74	27
		Occipital	Occipital Sup R (22.46 %)	1.93	64
		Limbic	Cingulum Post R (21.21 %)	1.64	7
14	Primary Visual (100.00 %)	Occipital	Calcarine L (90.96 %)	3.83	483

	Higher Visual (48.91 %)	Occipital	Calcarine R (88.52 %)	3.98	370
		Occipital	Cuneus R (85.75 %)	3.28	313
		Occipital	Cuneus L (83.57 %)	3.58	290
		Occipital	Lingual R (81.39 %)	2.96	433
		Occipital	Occipital Sup L (71.83 %)	2.73	181
		Occipital	Lingual L (67.15 %)	3.04	368
		Occipital	Occipital Sup R (61.75 %)	2.88	176
		Occipital	Occipital Mid R (37.66 %)	1.91	177
		Occipital	Occipital Mid L (29.27 %)	1.95	226
15	Precuneus (59.84 %)	Parietal	Parietal Sup L (87.76 %)	3.88	373
	Visuospatial (47.29 %)	Parietel	Parietal Sup R (86.87 %)	3.58	311
	Ventral DMN (43.44 %)	Parietal	Parietal Inf R (82.61 %)	2.85	285
	Left ECN (27.89 %)	Parietal	Parietal Inf L (71.86 %)	3	429
		Parietal	Angular R (53.61 %)	3.21	193
		Occipital	Occipital Sup R (48.07 %)	3.67	137
		Occipital	Occipital Sup L (47.22 %)	2.8	119
		Parietal	Precuneus L (46.53 %)	2.75	369
		Occipital	Occipital Mid R (44.68 %)	2.6	210
		Parietal	Precuneus R (41.75 %)	2.61	296
		Occipital	Occipital Mid L (32.38 %)	2.56	250
		Parietal	Angular L (30.99 %)	2.1	88
16	Sensorimotor (34.46 %)	Parietal	Paracentral Lobule L (100.00 %)	4.42	166
	Anterior Salience (32.92 %)	Parietal	Paracentral Lobule R (100.00 %)	4.37	131
		Frontal	Supp Motor Area R (86.04 %)	3.59	413
		Limbic	Cingulum Mid L (78.74 %)	2.87	400
		Frontal	Supp Motor Area L (75.98 %)	3.25	329
		Limbic	Cingulum Mid R (67.37 %)	2.99	351
		Frontal	Precentral R (36.08 %)	2.37	188
		Parietal	Postcentral R (33.13 %)	2.73	219
		Parietal	Postcentral L (29.58 %)	2.32	197
		Parietal	Precuneus L (27.99 %)	3.02	222
		Frontal	Precentral L (25.98 %)	2.39	146
		Frontal	Frontal Sup R (25.68 %)	2.33	179
		Parietal	Precuneus R (23.55 %)	2.61	167
		Parietel	Parietal Sup R (20.11 %)	2.21	72
17	Visuospatial (58.76 %)	Parietel	Parietal Sup R (76.26 %)	3.01	273
	Posterior Salience (51.33 %)	Parietal	Postcentral R (68.68 %)	3.83	454
	Sensorimotor (34.46 %)	Parietal	SupraMarginal R (64.85 %)	2.79	286
		Parietal	Parietal Sup L (64.00 %)	2.96	272
		Parietal	Parietal Inf L (60.80 %)	3.39	363
		Parietal	Postcentral L (59.46 %)	3.28	396
		Parietal	Parietal Inf R (55.36 %)	3.43	191
		Parietal	SupraMarginal L (47.39 %)	2.75	136
		Frontal	Precentral R (45.87 %)	2.26	239
		Frontal	Precentral L (40.21 %)	2.23	226
		Frontal	Frontal Sup R (20.66 %)	2.28	144

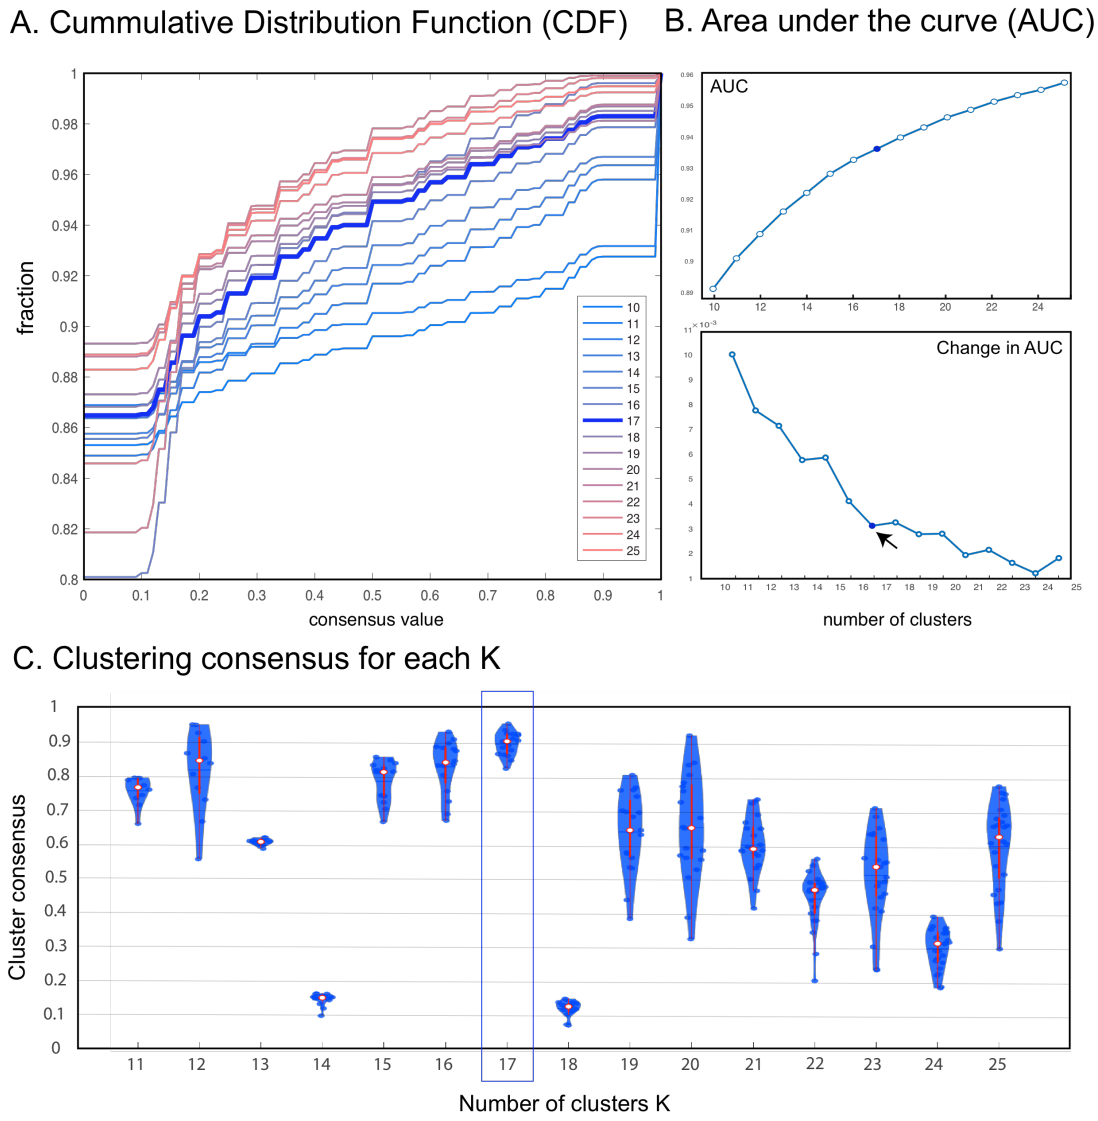

Figure A.3 – **Consensus quality metrics.** The (A) cumulative distribution function (CDF) indicates the extent to which the consensus matrix distribution is skewed toward 0 and 1, with a flat line being the ideal shape (i.e., 0 means two frames are never clustered together while 1 means frames are always clustered together). The (B) area under the curve (AUC) of the CDF and the change in AUC displays the optimal number of cluster K to which there is minimal increase in the AUC. Finally, the (C) clustering consensus gives a measure of the stability of the observed iCAPs with respect to different K over multiple runs of the clustering; $K = 17$ shows the highest cluster consensus. Using all these three consensus quality metrics, we chose $K = 17$ as the optimal number of clusters.

Table A.2 – Test statistics corresponding to Figure 3 (A) and (C) of the main manuscript. Results from two-sample paired t-tests and 1000 rounds of permutation testing. We compared iCAPs RCD and average durations across different sleep stages. Effect size corresponds to Cohen's d statistic.

iCAP		Wake-N1	Wake-N2	Wake-N3	N1-N2	N1-N3	N2-N3
1. visual-sensory							
RCD	t-statistic	-3.92	-6.22	1.69	-2.3	5.61	7.91
	p-value	0.357	0.154	0.715	0.511	0.122	0.025
	Effect-size	-0.34	-0.57	0.15	-0.25	0.6	0.88
average durations	t-statistic	-3.92	-6.22	1.69	-2.3	5.61	7.91
	p-value	0.357	0.154	0.715	0.511	0.122	0.025
	Effect-size	-0.34	-0.57	0.15	-0.25	0.6	0.88
2. insula/thalamus							
RCD	t-statistic	-0.01	-7.8	7.99	-7.79	8	15.79
	p-value	0.999	0.011	0.031	0.031	0.042	0.002
	Effect-size	0	-1.01	0.92	-0.87	0.8	1.71
average durations	t-statistic	-0.01	-7.8	7.99	-7.79	8	15.79
	p-value	0.999	0.011	0.031	0.031	0.042	0.002
	Effect-size	0	-1.01	0.92	-0.87	0.8	1.71
3. anterior DMN							
RCD	t-statistic	6.97	1.79	15.63	-5.17	8.66	13.83
	p-value	0.07	0.581	0.001	0.063	0.009	0.001
	Effect-size	0.74	0.23	1.81	-0.73	1.1	2.21
average durations	t-statistic	6.97	1.79	15.63	-5.17	8.66	13.83
	p-value	0.07	0.581	0.001	0.063	0.009	0.001
	Effect-size	0.74	0.23	1.81	-0.73	1.1	2.21
4. left ECN							
RCD	t-statistic	4.4	3.68	15.34	-0.73	10.94	11.66
	p-value	0.16	0.119	0.001	0.761	0.002	0.001
	Effect-size	0.51	0.66	1.91	-0.12	1.31	2.25
average durations	t-statistic	4.4	3.68	15.34	-0.73	10.94	11.66
	p-value	0.16	0.119	0.001	0.761	0.002	0.001
	Effect-size	0.51	0.66	1.91	-0.12	1.31	2.25
5. secondary visual							
RCD	t-statistic	-5.79	-6.72	-2.34	-0.93	3.44	4.38
	p-value	0.103	0.005	0.364	0.797	0.354	0.139
	Effect-size	-0.74	-1.23	-0.36	-0.11	0.35	0.59
average durations	t-statistic	-5.79	-6.72	-2.34	-0.93	3.44	4.38
	p-value	0.103	0.005	0.364	0.797	0.354	0.139
	Effect-size	-0.74	-1.23	-0.36	-0.11	0.35	0.59
6. posterior CEB							
RCD	t-statistic	-0.21	2.76	7.88	2.97	8.09	5.12
	p-value	0.951	0.356	0.002	0.344	0.008	0.016
	Effect-size	-0.02	0.4	1.15	0.41	1.12	1.02
average durations	t-statistic	-0.21	2.76	7.88	2.97	8.09	5.12
	p-value	0.951	0.356	0.002	0.344	0.008	0.016
	Effect-size	-0.02	0.4	1.15	0.41	1.12	1.02
7. right ECN							
RCD	t-statistic	5.45	3.64	18.2	-1.81	12.75	14.56
	p-value	0.082	0.22	0.001	0.316	0.001	0.001
	Effect-size	0.71	0.55	2.03	-0.39	1.84	2.47
average durations	t-statistic	5.45	3.64	18.2	-1.81	12.75	14.56
	p-value	0.082	0.22	0.001	0.316	0.001	0.001
	Effect-size	0.71	0.55	2.03	-0.39	1.84	2.47
8. auditory/motor							

RCD	t-statistic	-2.07	-7.61	-2.68	-5.54	-0.61	4.93
	p-value	0.308	0.003	0.22	0.034	0.804	0.069
	Effect-size	-0.39	-1.31	-0.48	-0.86	-0.1	0.73
average durations	t-statistic	-2.07	-7.61	-2.68	-5.54	-0.61	4.93
	p-value	0.308	0.003	0.22	0.034	0.804	0.069
	Effect-size	-0.39	-1.31	-0.48	-0.86	-0.1	0.73
9. amygdala/OFC							
RCD	t-statistic	0.48	-2.06	7.45	-2.54	6.97	9.51
	p-value	0.885	0.552	0.022	0.481	0.03	0.001
	Effect-size	0.05	-0.23	0.89	-0.29	0.86	1.25
average durations	t-statistic	0.48	-2.06	7.45	-2.54	6.97	9.51
	p-value	0.885	0.552	0.022	0.481	0.03	0.001
	Effect-size	0.05	-0.23	0.89	-0.29	0.86	1.25
10. anterior CEB							
RCD	t-statistic	1.09	-2.28	7.46	-3.38	6.36	9.74
	p-value	0.784	0.501	0.024	0.38	0.089	0.004
	Effect-size	0.11	-0.26	0.93	-0.34	0.7	1.18
average durations	t-statistic	1.09	-2.28	7.46	-3.38	6.36	9.74
	p-value	0.784	0.501	0.024	0.38	0.089	0.004
	Effect-size	0.11	-0.26	0.93	-0.34	0.7	1.18
11. posterior DMN							
RCD	t-statistic	0.62	-5.07	-2.82	-5.69	-3.44	2.25
	p-value	0.769	0.012	0.151	0.006	0.108	0.284
	Effect-size	0.12	-1.01	-0.56	-1.06	-0.63	0.43
average durations	t-statistic	0.62	-5.07	-2.82	-5.69	-3.44	2.25
	p-value	0.769	0.012	0.151	0.006	0.108	0.284
	Effect-size	0.12	-1.01	-0.56	-1.06	-0.63	0.43
12. language							
RCD	t-statistic	-1.89	-11.53	-4.09	-9.65	-2.2	7.44
	p-value	0.532	0.002	0.217	0.009	0.581	0.05
	Effect-size	-0.25	-1.49	-0.49	-1.09	-0.24	0.79
average durations	t-statistic	-1.89	-11.53	-4.09	-9.65	-2.2	7.44
	p-value	0.532	0.002	0.217	0.009	0.581	0.05
	Effect-size	-0.25	-1.49	-0.49	-1.09	-0.24	0.79
13. precuneus							
RCD	t-statistic	-4.84	-7.15	-4.1	-2.31	0.74	3.05
	p-value	0.13	0.036	0.14	0.516	0.79	0.356
	Effect-size	-0.58	-0.83	-0.59	-0.25	0.1	0.38
average durations	t-statistic	-4.84	-7.15	-4.1	-2.31	0.74	3.05
	p-value	0.13	0.036	0.14	0.516	0.79	0.356
	Effect-size	-0.58	-0.83	-0.59	-0.25	0.1	0.38
14. primary visual							
RCD	t-statistic	-5.75	-8.57	-6.23	-2.82	-0.48	2.34
	p-value	0.025	0.002	0.008	0.313	0.839	0.343
	Effect-size	-0.89	-1.69	-1.11	-0.41	-0.06	0.39
average durations	t-statistic	-5.75	-8.57	-6.23	-2.82	-0.48	2.34
	p-value	0.025	0.002	0.008	0.313	0.839	0.343
	Effect-size	-0.89	-1.69	-1.11	-0.41	-0.06	0.39
15. visuospatial							
RCD	t-statistic	0.75	-4.59	-0.04	-5.34	-0.79	4.55
	p-value	0.834	0.059	0.986	0.049	0.833	0.032
	Effect-size	0.09	-0.79	-0.01	-0.8	-0.11	0.87

average durations	t-statistic	0.75	-4.59	-0.04	-5.34	-0.79	4.55
	p-value	0.834	0.059	0.986	0.049	0.833	0.032
	Effect-size	0.09	-0.79	-0.01	-0.8	-0.11	0.87
16. sensorimotor							
RCD	t-statistic	-0.74	-14.78	-9.18	-14.04	-8.44	5.6
	p-value	0.743	0.001	0.001	0.002	0.002	0.163
	Effect-size	-0.14	-1.71	-1.43	-1.62	-1.32	0.57
average durations	t-statistic	-0.74	-14.78	-9.18	-14.04	-8.44	5.6
	p-value	0.743	0.001	0.001	0.002	0.002	0.163
	Effect-size	-0.14	-1.71	-1.43	-1.62	-1.32	0.57
17. somatosensory							
RCD	t-statistic	-0.37	-7.53	-3.03	-7.15	-2.66	4.49
	p-value	0.854	0.004	0.228	0.014	0.355	0.107
	Effect-size	-0.06	-1.24	-0.47	-1.03	-0.36	0.62
average durations	t-statistic	-0.37	-7.53	-3.03	-7.15	-2.66	4.49
	p-value	0.854	0.004	0.228	0.014	0.355	0.107
	Effect-size	-0.06	-1.24	-0.47	-1.03	-0.36	0.62

Table A.3 – Test statistics corresponding to Figure 4(A) of the main manuscript in Chapter 3.2. Results from ANOVA and the corrected p-values after multiple comparison test. Asterisks mark p-values less than 0.05.

N	F-statistic	Prob >F	Wake - N1	Wake - N2	WAKE-N3	N1 - N2	N1 - N3	N2 - N3
0	2.36	0.08	0.18	0.97	0.79	0.09	0.81	0.58
1	3.24	0.03	0.69	0.02*	0.98	0.29	0.93	0.14
2	4.65	0.01	0.98	0.01*	0.99	0.05*	0.93	0.02*
3	4.74	0.01	1	0.04*	0.81	0.03*	0.91	0.01*
4	4.33	0.01	0.59	0.23	0.49	0.02*	0.99	0.02*
5	6.17	0	0.4	0.18	0.19	0.01*	0.93	0*
6	1.08	0.37	1	0.81	0.72	0.75	0.81	0.3
7	1.26	0.3	0.83	1	0.34	0.83	0.8	0.34
8	0.9	0.45	0.63	1	0.61	0.71	1	0.68
9	1.07	0.37	0.62	1	0.51	0.68	0.99	0.56
10	1.53	0.217	0.55	1	0.32	0.58	0.98	0.35

Table A.4 – Test statistics corresponding to Figure 4 (B) and (C) of the main manuscript. Results from ANOVA and the corrected p-values after multiple comparison test.

	F-statistic	Prob >F	Wake - N1	Wake - N2	WAKE-N3	N1 - N2	N1 - N3	N2 - N3
same-signed	22.71	3.90E-10	4.50E-07	0.0091	3.00E-04	1.30E-09	0.009	4.60E-07
opposite-signed	12.92	1.00E-06	2.30E-06	0.0051	0.0012	7.10E-08	8.80E-04	1.20E-04
Pearson correlation	33.13	4.70E-13	0.2571	1.80E-07	3.30E-09	1.46-7	6.30E-09	0.115

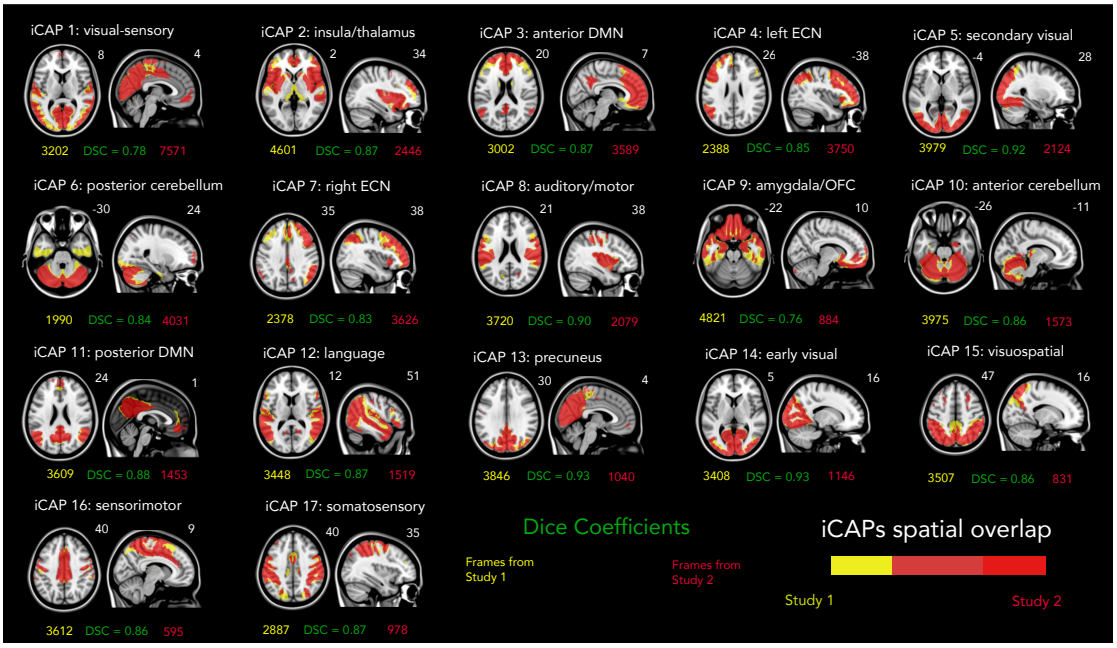

Figure A.4 – **Spatial similarity between iCAPs** obtained by averaging the frame indices coming from the sustained sleep data (yellow) versus iCAPs obtained by averaging the frame indices coming from the sleep onset data (red). The percent number of frames that contributed to the recovery of each iCAP coming from each dataset is written in red and yellow fonts. The iCAP maps are spatially z-scored and thresholded at $|z| > 1.5$. The Dice-Sorrensen Coefficient (DSC) is computed as a quantitative measure for the similarity of the two sets of iCAPs and is written in green font.

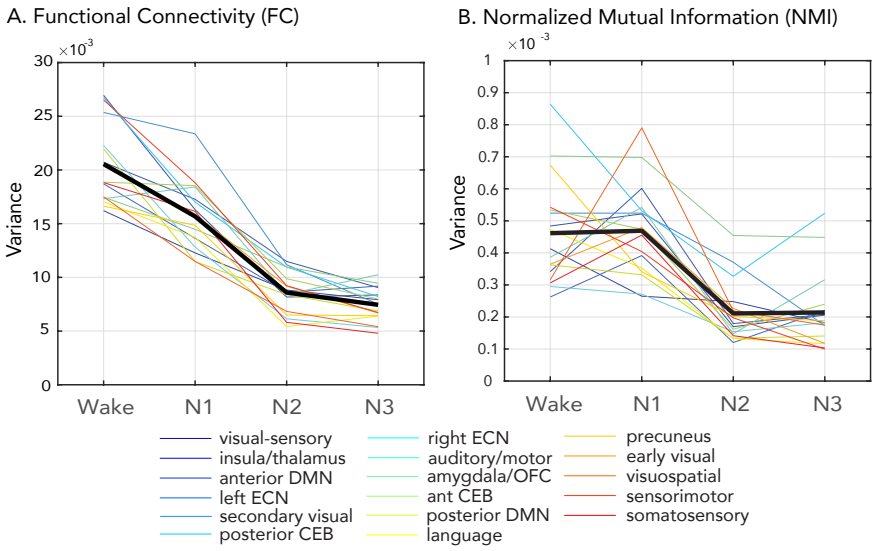

Figure A.5 – **Variance of Network Interactions Measures** (A) functional connectivity (FC) and (B) normalized mutual information (NMI) between iCAP time-courses across the different NREM sleep stages.

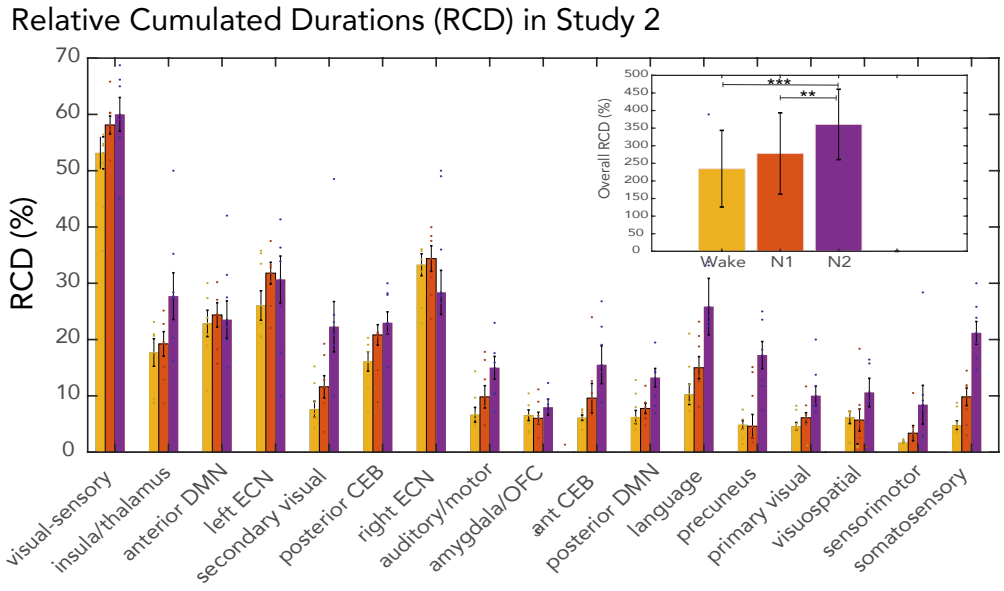

Figure A.6 – **Relative Cumulated Durations (RCD) obtained using the data from Study 2.** Overall, we generally observe a peak in RCD value during NREM stage 2, consistent with the RCD obtained using the data from Study 1. Data are represented as mean +/- SEM. Horizontal lines with 3 stars and 2 stars in the insets represent significant differences with p-values less than 0.001 and 0.01, respectively (corrected using Tukey's range test). Data in the insets are represented as mean +/- SD.

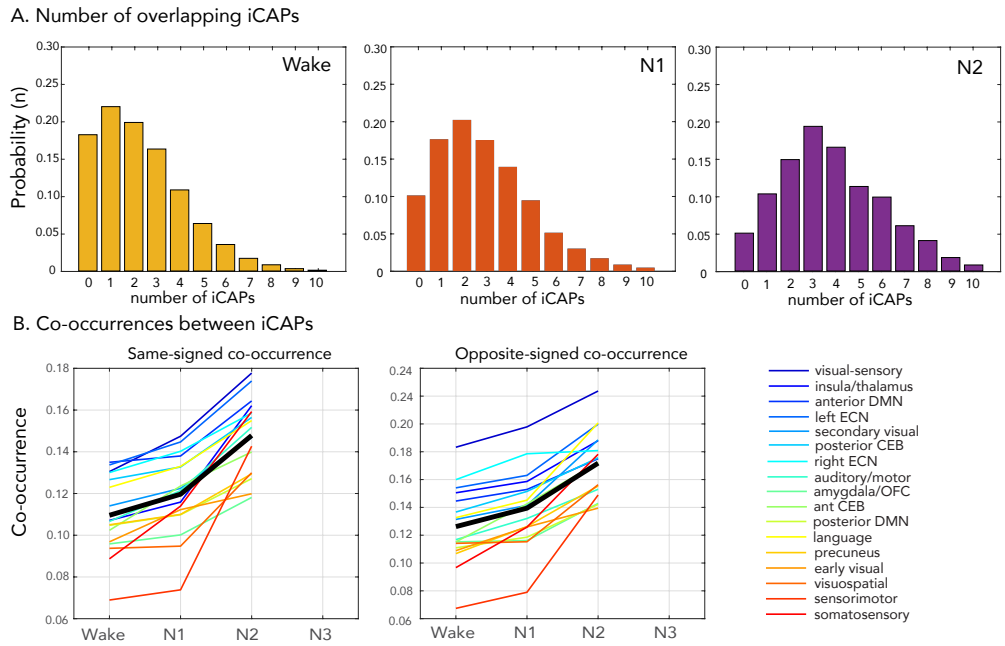

Figure A.7 – **Interactions between iCAPs across NREM sleep in Study 2.** (A) Number of overlapping iCAPs and (B) co-occurrences between iCAPs. Moreover, iCAPs are more likely to overlap and co-occur during NREM stage 2 compared to wakefulness.

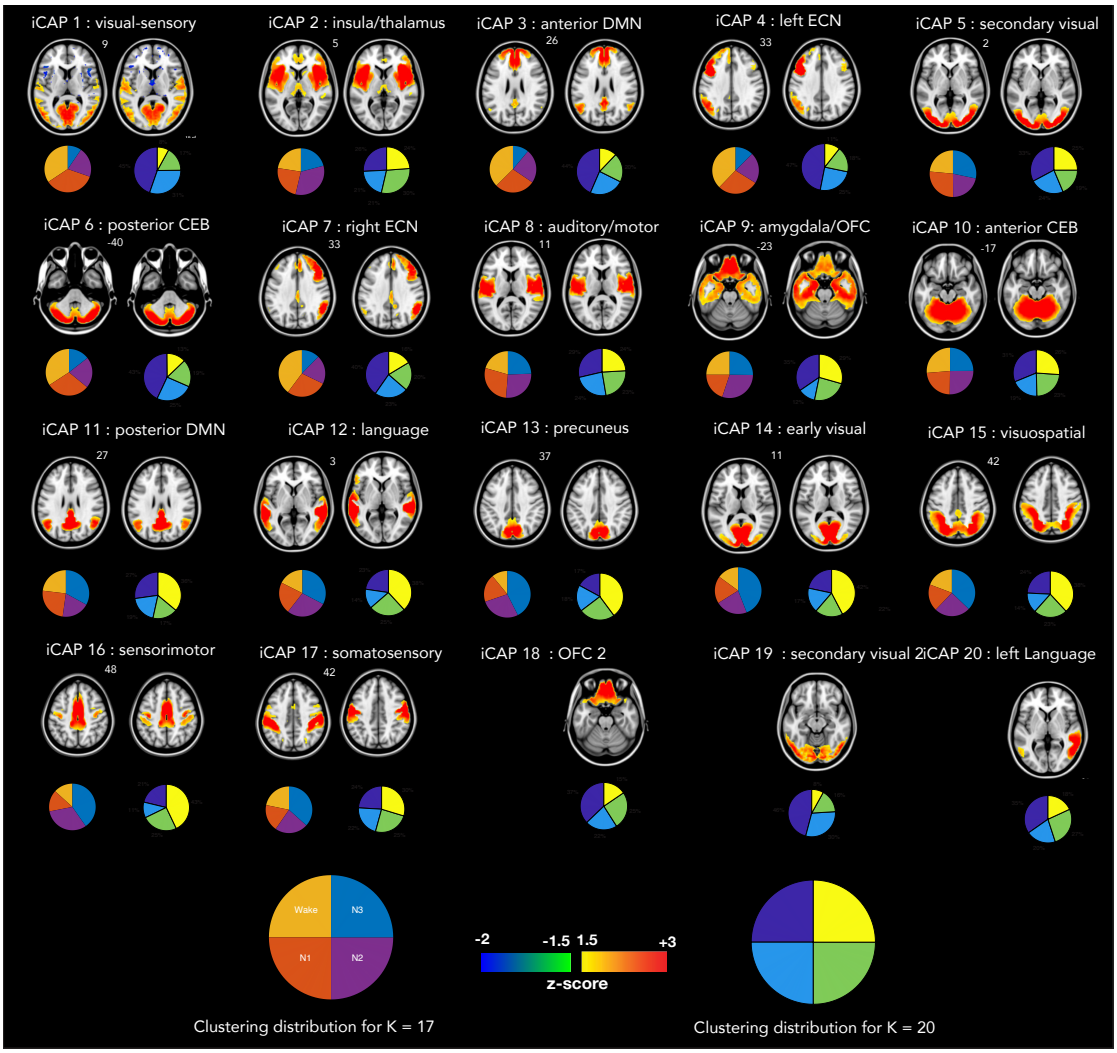

Figure A.8 – **Spatial and temporal characteristics of iCAPs**. Side by side comparison between the results of K = 17 (left) and K = 20 (right). Spatial characteristics of iCAPs generally show similar patterns and their percent distributions (pie chart) across NREM sleep retain similar proportions for K= 17 and K = 20. The 18th, 19th, and the 20th iCAPs for K = 20 repeat patterns that is already present for K = 17, in particular, the OFC2, secondary visual 2, and the left language.

B Supplementary Material for Chapter 4

The extraction of group-level eigenmodes was done through a series of Procrustes transformation, cf. Algorithm 1 in Chapter 4. The transformation, however, is done multiple times to remove bias towards the initial reference value. The optimal number of transformation was based on the off-diagonal errors (cf, Eq. 4.11) observed after several iterations of Procrustes transformation.

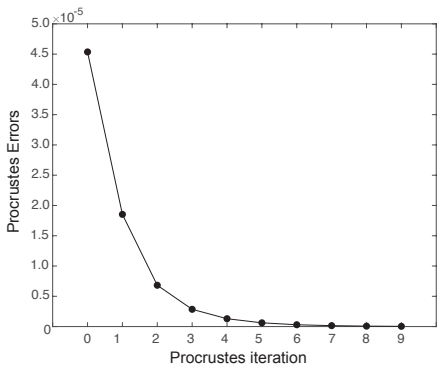

Figure B.1 – Off-diagonal errors of the cosine similarity matrices after several iterations of Procrustes transformation. Three iterations of the Procrustes transform significantly reduce the similarity matrices of inter-subject eigenspace.

C Supplementary Material for Chapter 5

All data used in the study were downloaded from the Human Connectome Project 1200 release, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657).

Motivation of different parameter choices for the extraction and analysis of in- CAPs.

The 15% threshold applied to the PCC was primarily motivated by the choice of the original proponents of co-activation pattern (CAP) analysis [Liu and Duyn, 2013]. To verify this chosen parameter, we scanned cut-off thresholds from 0 to 100% and evaluated the spatial similarity between the averaged frames corresponding to each of the threshold and the seed-correlation map (PCC-seed connectivity) for all subjects. In Fig. C.4(A), the similarity of the averaged frames with the seed correlation map reaches a plateau from 15% onwards and decreases again from 85% onwards. This implies that the overall mean of the 15% fMRI time frames having the highest values at the PCC seed region generates a DMN pattern very similar to that from the PCC-seed correlation map. Next, the optimal number of clusters (K = 8) was also based on the original CAPs analysis corresponding to the DMN [Liu and Duyn, 2013]. We verified the choice of optimal K using a consensus clustering procedure [Monti et al., 2003], which is a resampling-based procedure for optimal class discovery. We ran 100 folds of k-means clustering on randomly resampled data within each subject. The stability of the cluster assignments across several K-values ({2,3,4,...,12}) was evaluated by computing a "consensus" value derived for each pair of data points (i.e.,how often they were clustered together). Fig. C.4(B1) shows the cumulative distribution (CDF) of the consensus matrices across different K-values for an example subject (in case of perfect consensus, consensus matrix contains 0s and 1s only, and the CDF is flat). We identify the optimal number of clusters (K=8) based on the maximally observed change in area under the curve (AUC) of the averaged CDF across subjects (Fig. C.4(B3)). We also determined the stability of the in-CAPs across multiple sessions of fMRI. We divided the four sessions into two batches (sessions 1-2 and sessions 3-4 separately). Then, the temporal decomposition was done separately for the two batches in all subjects. Fig. C.4(C) displays the cosine similarity measure of the 8 in-CAPs from the two batches.

The results show a strong diagonal matrix reflecting the stability of in-CAPs across multiple sessions.

Finally, the probabilistic WM atlas [Zhang and Arfanakis, 2018] that was used in computing the signal average of the in-CAPs was thresholded to 0.3 to indicate which voxels are considered part of a WM tract. The optimal threshold was chosen based on the robustness of the results when compared to other threshold values. We therefore performed the signal averaging of the in-CAPs using multiple thresholds (0,0.1,0.2,...,1) for the WM atlas. Fig. C.4(D) reveals that when taking the cosine similarity between the vectorized signal-averaged values (Fig. 4(D) in manuscript) using the different thresholds, there is a particular threshold (thresh = 0.3) for which the outcome does not change with increasing threshold values. Higher thresholds greater than 0.9 were found to be too extreme.

CAPs back-projection to task data.

The signal recovery is implemented for all functional volumes of the working memory and relational task data. Subject-level CAPs are back-projected to the top 15% significant frames corresponding to the PCC. The frames that survived the thresholding are then labeled based on their extent of spatial similarity to the derived in-CAPs, using cosine similarity as a metric.

C.1 Supplementary Figures

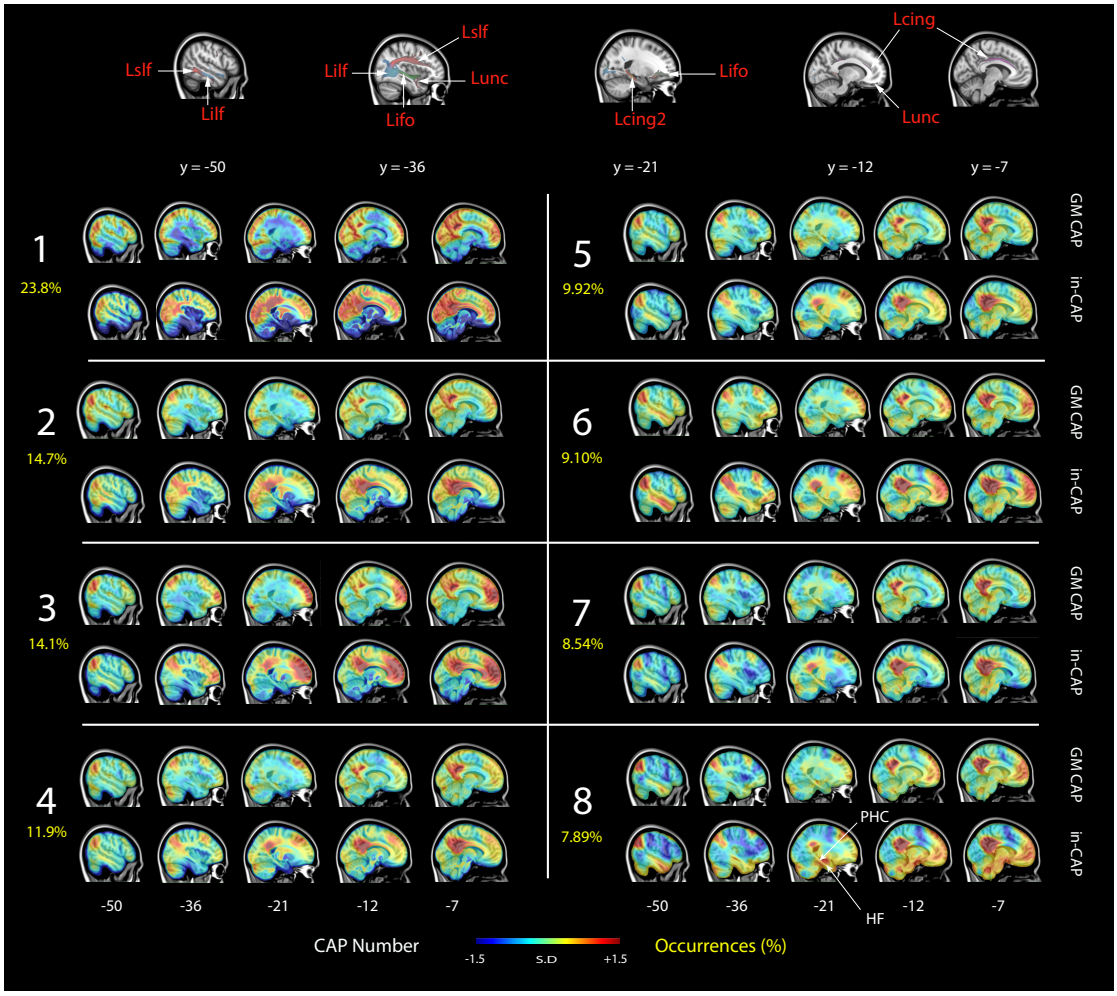

Figure C.1 – **Sagittal view of in-CAP and GM CAP pairs.** The CAPs are arranged according to their frequency of occurrence. To ease visual inspection of the CAPs, overlays of major white matter bundles are illustrated for each of the five sagittal slices. Abbreviations of WM tracts are presented in Table C.2.

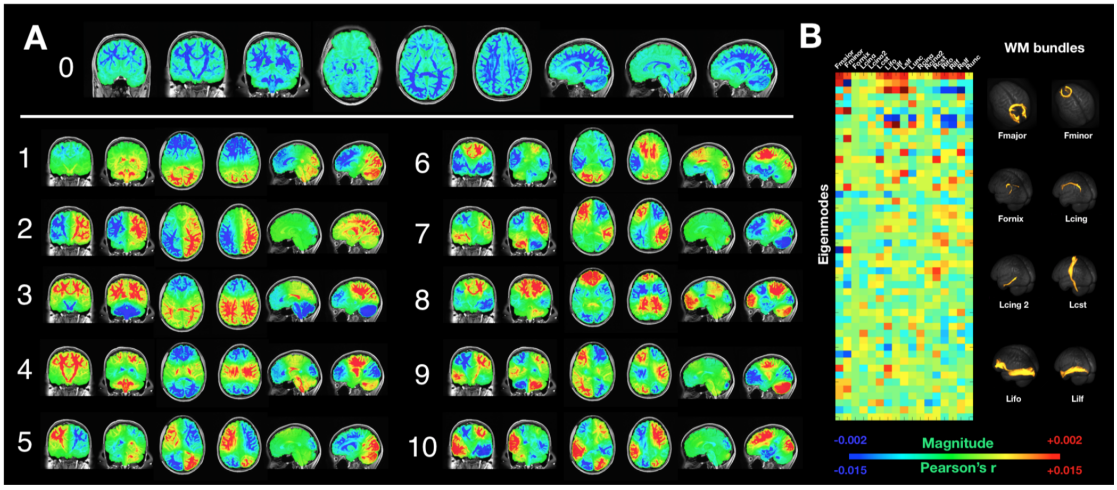

Figure C.2 – **Eigenfunctions of the Laplace operator corresponding to the voxel-level brain graph.** (A) We present examples of the first 11 highly resolved human brain eigenmodes corresponding to the 11 lowest eigenvalues. The top panel is the first eigenmode (corresponding to $\lambda = 0$), showing distinct major white matter structures. The numbers on the left show increasing eigenvalues, pertaining to spatial frequencies, thereby increasing spatial variations across the whole brain. In particular, the second and third lowest spatially varying eigenmodes describe geometrical information of the head shape, dividing the brain into posterior and anterior, and left and right, respectively. (B) Spatial correlation (corrected, p<0.005) of the eigenmodes evaluated with respect to known fiber bundles [Zhang and Arfanakis, 2018]. The correlation is presented for increasing eigenvalues displayed from top to bottom. Abbreviations used in naming the fiber bundles are presented in Table C.2.

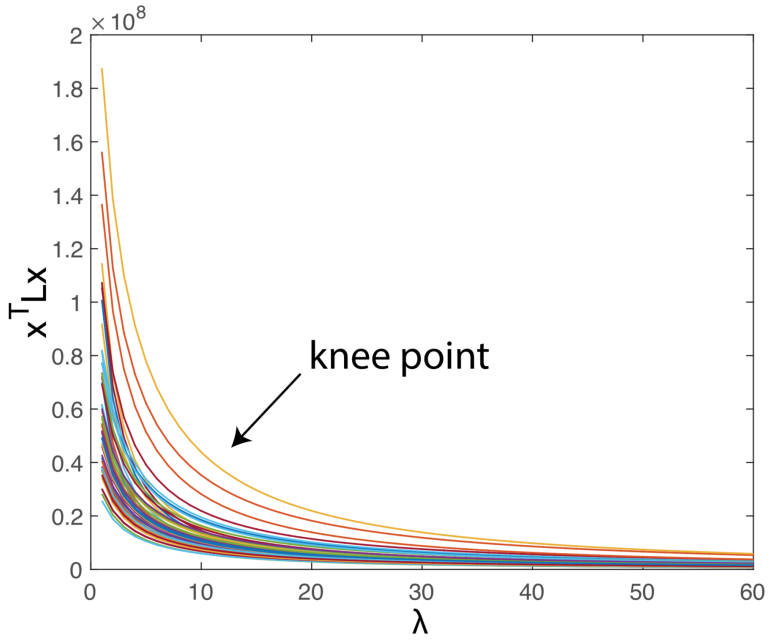

Figure C.3 – **Regularization parameter for the signal recovery framework.** For each subject, we evaluate the cost of the regularization term in the set $\lambda = \{1,2,3,\ldots M\}$, where $\lambda = 60$. The knee point of the plot is taken as the optimal λ. The range of optimal parameters used in the dataset analyzed are $\lambda^* = \{9,10,11\}$.

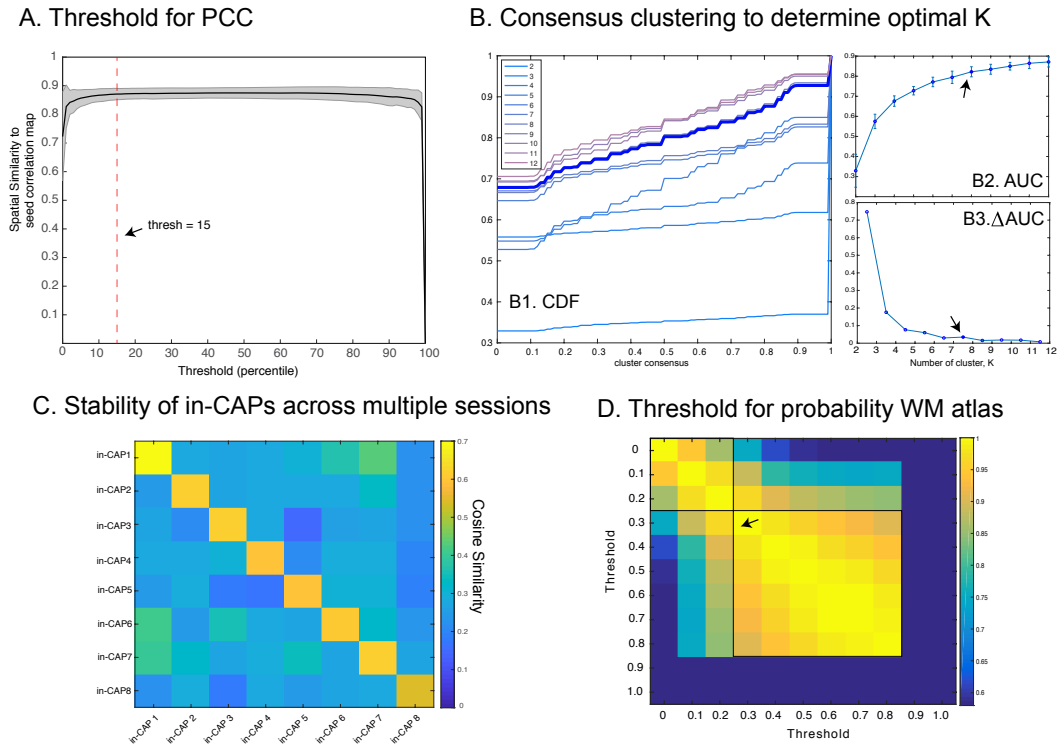

Figure C.4 – **Parameter optimisation for the extraction and analysis of in-CAPs.** Spatial similarity between the PCC-seed correlation and the averaged inpainted frames that survived different thresholds applied to the PCC. The similarity increases as we include more frames and reaches a plateau upon reaching 15%, which falls down again at about 85%. (B) Consensus clustering metrics for determining the optimal number of clusters in-CAPs. The (B1) cumulative distribution function (CDF) indicates the extent to which the consensus matrix distribution is skewed toward 0 and 1, with a flat line being the ideal shape (i.e., 0 means two frames are never clustered together while 1 means frames are always clustered together). The (B2) area under the curve (AUC) of the CDF and the (B3) change in AUC displays the optimal number of cluster K to which there is a minimal increase in the AUC (K = 8). (C) Stability of in-CAPs when running the clustering in 2 sessions, instead of 4 sessions together. Cosine similarity of the in-CAPs reveals strong diagonal relationship, and weaker off-diagonal values. (D) When doing the signal averaging within the probabilistic WM atlas, a threshold is used to determine voxels that are absolutely part of a tract. We used a threshold of 0.3 based on spatial similarity of the resulting WM averages. At this threshold, the observed signal averages show no particular difference with increasing threshold. At threshold greater than 0.9, we encounter tracts that are too small to take average on.

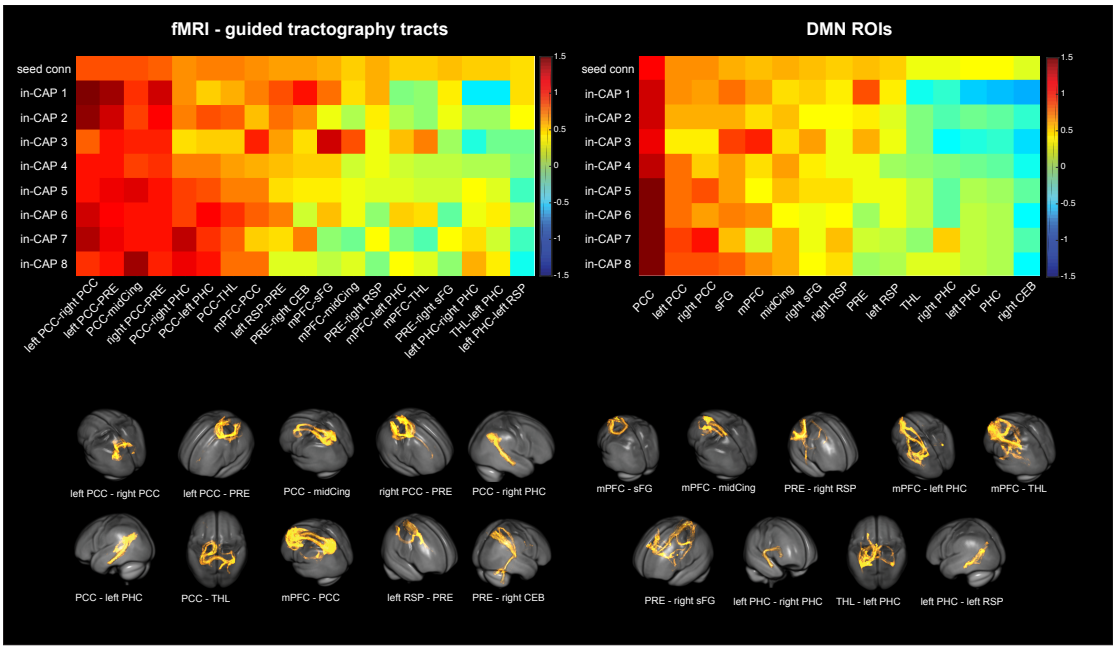

Figure C.5 – **Comparison with fMRI-guided tractography.** Average signal of the inpainted CAPs and inpainted PCC-seed connectivity within (A) the different tracts obtained from running DMN- guided probabilistic tractography [Figley et al., 2015]. (B) Average of the in-CAPs and PCC-seed connectivity within the corresponding DMN ROIs.

C.2 Supplementary Tables

Table C.1 – DMN sub-component MNI coordinates adapted from Andrews-Hannah [Andrews-Hanna et al., 2010] classification

Sub-component (left region)	Abbreviation	X	Y	Z
1. Anterior medial prefrontal cortex	amPFC	-6	52	-2
2. Posterior cingulate cortex	PCC	-8	-56	26
3. Dorsal medial prefrontal cortex	dmPFC	0	52	26
4. Temporal parietal junction	TPJ	-54	-54	28
5. Lateral temporal cortex	LTC	-60	-24	-18
6. Temporal pole	TempP	-50	14	-40
7. Ventral medial prefrontal cortex	vmPFC	0	26	-18
8. Posterior inferior parietal lobule	pIPL	-44	-74	32
9. Retrosplenial cortex	Rsp	-14	-52	8
10. Parahippocampal cortex	PHC	-28	-40	-12
11. Hippocampal formation	HF	-22	-20	-26

Table C.2 – Major white matter bundles evaluated in the study.

WM structure	Abbreviation
1. Forceps major	Fmajor
2. Forceps minor	Fminor
3. Left / right cingulum	L/Rcing
4. Left / right superior longitudinal fasciculus	L/Rslf
5. Left / right inferior longitudinal fasciculus	L/Rilf
6. Left / right cortico-spinal tract	L/Rcst
7. Left / right inferior fronto-occipital fasciculus	L/Rifo
8. Fornix	Fornix
9. Left / right hippocampal cingulum	L/Rcing2
10. Left / right uncinate fasciculus	L/Runc

D Supplementary Material for Chapter 6

D.0.1 Supplementary Methods

Computation of the Euclidean distance map

Briefly, we downsampled the preprocessed functional data, originally from 1.25 mm, diffusion-spaced, to 3mm voxel resolution. Then voxel-to-voxel FC map for all GM voxels was calculated. For each voxel, we obtained their highest 2% FC values. The mean Euclidean distance of all remaining voxels that survived the thresholding are then computed with respect to the voxel of interest. This was done in 4 resting-state sessions and the resulting euclidean distance map is averaged across 50 subjects.

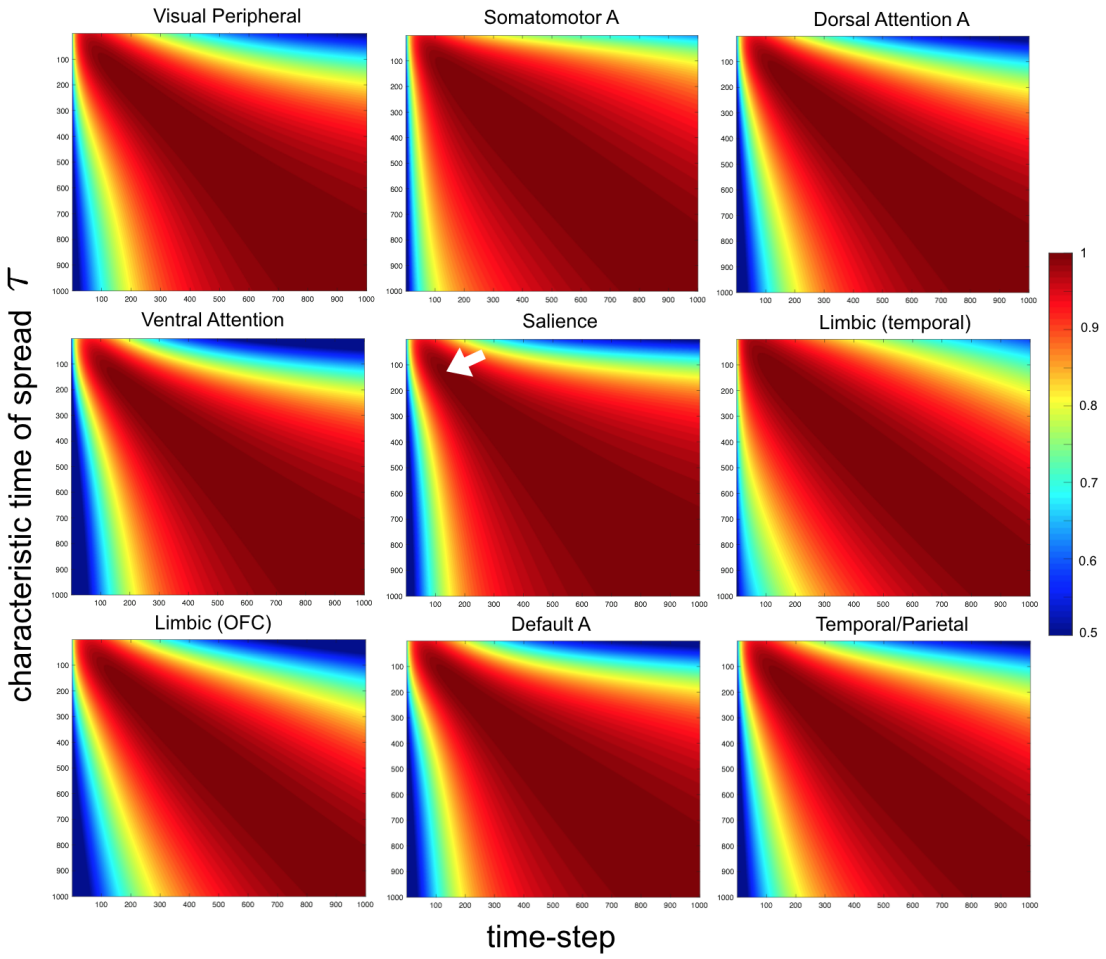

Figure D.1 – **One-to-one mapping of diffusion rate and diffusion time-steps.** Cosine similarity between diffused signals using different spectral graph diffusion rate τ and vertex diffusion time-steps corresponding to nine representative networks. The white arrow indicates the point when the inpainted volumes associated to diffusion rate τ and steady state time-step t_{final} matches. Different starting points can be observed in each network.

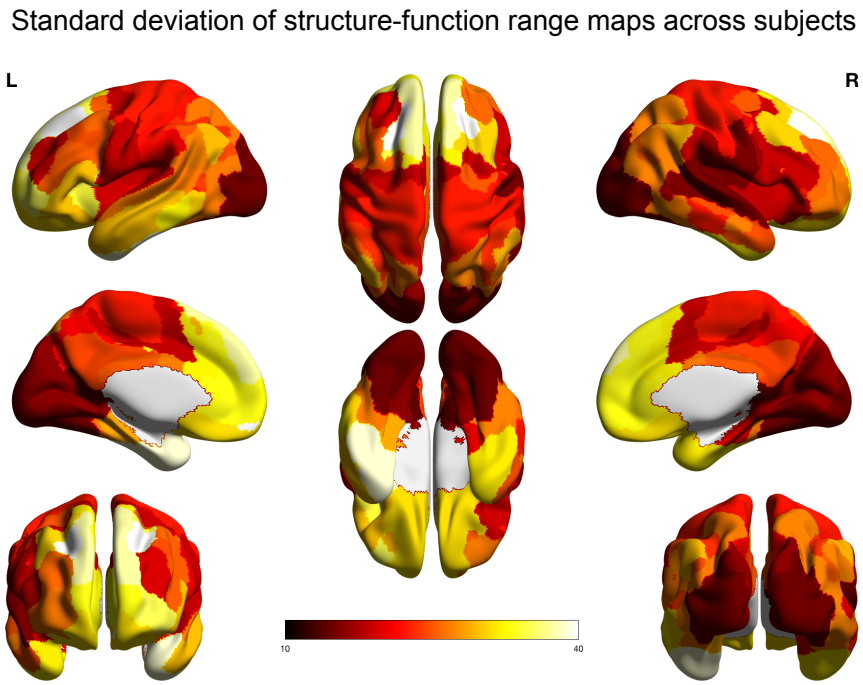

Figure D.2 – **Variability across subjects of structure-function range map.** Standard deviation of characteristic time of spread projected in terms of distance across 50 subjects evaluated. The topological variability across individuals scales with the structure-function range map.

Bibliography

[Abdelnour et al., 2018] Abdelnour, F., Dayan, M., Devinsky, O., Thesen, T., and Raj, A. (2018). Functional brain connectivity is predictable from anatomic network's Laplacian eigenstructure. *NeuroImage*, 172(February):728–739.

[Abdelnour et al., 2015] Abdelnour, F., Mueller, S., and Raj, A. (2015). Relating Cortical Atrophy in Temporal Lobe Epilepsy with Graph Diffusion-Based Network Models. *PLoS Computational Biology*, 11(10):1–24.

[Abdelnour et al., 2014] Abdelnour, F., Voss, H. U., and Raj, A. (2014). Network diffusion accurately models the relationship between structural and functional brain connectivity networks. *NeuroImage*, 90:335–347.

[Abramian et al., 2020] Abramian, D., Larsson, M., Eklund, A., and Behjat, H. (2020). Improved functional MRI activation mapping in white matter through diffusion-adapted spatial filtering. In *2020 IEEE 17th International Symposium on Biomedical Imaging*, pages 539–543. IEEE.

[Achard and Bullmore, 2007] Achard, S. and Bullmore, E. (2007). Efficiency and Cost of Economical Brain Functional Networks. *PLoS Computational Biology*, 3(2):e17.

[Achard et al., 2006] Achard, S., Salvador, R., Whitcher, B., Suckling, J., and Bullmore, E. (2006). A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *Journal of Neuroscience*, 26(1):63–72.

[Alexander et al., 2007] Alexander, A. L., Lee, J. E., Lazar, M., and Field, A. S. (2007). Diffusion tensor imaging of the brain. *Neurotherapeutics*, 4(3):316–329.

[Alexander-Bloch et al., 2013] Alexander-Bloch, A. F., Vértes, P. E., Stidd, R., Lalonde, F., Clasen, L., Rapoport, J., Giedd, J., Bullmore, E. T., and Gogtay, N. (2013). The Anatomical Distance of Functional Connections Predicts Brain Network Topology in Health and Schizophrenia. *Cerebral Cortex*, 23(1):127–138.

[Allen et al., 2014] Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., and Calhoun, V. D. (2014). Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cerebral Cortex*, 24(3):663–676.

[Alvarez and Emory, 2006] Alvarez, J. and Emory, E. (2006). Executive function and the frontal lobes: A meta-analytic review. *Neuropsychology Review*, 16(1):17–42.

[Amico and Goñi, 2018] Amico, E. and Goñi, J. (2018). Mapping hybrid functional-structural connectivity traits in the human connectome. *Network Neuroscience*, 2(3):306–322.

[Andrews-Hanna, 2012] Andrews-Hanna, J. R. (2012). The Brain's Default Network and Its Adaptive Role in Internal Mentation. *The Neuroscientist*, 18(3):251–270.

[Andrews-Hanna et al., 2010] Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., and Buckner, R. L. (2010). Functional-Anatomic Fractionation of the Brain's Default Network. *Neuron*, 65(4):550–562.

[Andrews-Hanna et al., 2007] Andrews-Hanna, J. R., Snyder, A. Z., Vincent, J. L., Lustig, C., Head, D., Raichle, M., and Buckner, R. L. (2007). Disruption of Large-Scale Brain Systems in Advanced Aging. *Neuron*, 56(5):924–935.

[Anzai et al., 2007] Anzai, A., Peng, X., and Van Essen, D. C. (2007). Neurons in monkey visual area V2 encode combinations of orientations. *Nature Neuroscience*, 10(10):1313–1321.

[Arsalidou et al., 2018] Arsalidou, M., Pawliw-Levac, M., Sadeghi, M., and Pascual-Leone, J. (2018). Brain areas associated with numbers and calculations in children: Meta-analyses of fMRI studies. *Developmental Cognitive Neuroscience*, 30(January 2017):239–250.

[Ashburner, 2007] Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95 – 113.

[Atasoy et al., 2016] Atasoy, S., Donnelly, I., and Pearson, J. (2016). Human brain networks function in connectome-specific harmonic waves. *Nature Communications*, 7(1):10340.

[Atasoy et al., 2017] Atasoy, S., Roseman, L., Kaelen, M., Kringelbach, M. L., Deco, G., and Carhart-Harris, R. L. (2017). Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under LSD. *Scientific Reports*, 7(1):17661.

[Bajada et al., 2019] Bajada, C. J., Schreiber, J., and Caspers, S. (2019). Fiber length profiling: A novel approach to structural brain organization. *NeuroImage*, 186(November 2018):164–173.

[Barch et al., 2013] 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.

[Basser et al., 1994] Basser, P., Mattiello, J., and Lebihan, D. (1994). Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo. *Journal of Magnetic Resonance, Series B*, 103(3):247–254.

[Basser and Pierpaoli, 1996] Basser, P. J. and Pierpaoli, C. (1996). Microstructural and Physiological Features of Tissues Elucidated by Quantitative-Diffusion-Tensor MRI. *Journal of Magnetic Resonance, Series B*, 111(3):209–219.

[Bassett and Bullmore, 2006] Bassett, D. S. and Bullmore, E. (2006). Small-world brain networks. *Neuroscientist*, 12(6):512–523.

[Bassett and Bullmore, 2017] Bassett, D. S. and Bullmore, E. T. (2017). Small-World Brain Networks Revisited. *Neuroscientist*, 23(5):499–516.

[Bassett et al., 2009] Bassett, D. S., Bullmore, E. T., Meyer-Lindenberg, A., Apud, J. A., Weinberger, D. R., and Coppola, R. (2009). Cognitive fitness of cost-efficient brain functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 106(28):11747–11752.

[Bassett et al., 2010] Bassett, D. S., Greenfield, D. L., Meyer-Lindenberg, A., Weinberger, D. R., Moore, S. W., and Bullmore, E. T. (2010). Efficient physical embedding of topologically complex information processing networks in brains and computer circuits. *PLoS Computational Biology*, 6(4).

[Bastiani et al., 2017] Bastiani, M., Cottaar, M., Dikranian, K., Ghosh, A., Zhang, H., Alexander, D. C., Behrens, T. E., Jbabdi, S., and Sotiropoulos, S. N. (2017). Improved tractography using asymmetric fibre orientation distributions. *NeuroImage*, 158:205–218.

[Bastien et al., 2000] Bastien, C. H., Ladouceur, C., and Campbell, K. B. (2000). EEG characteristics prior to and following the evoked K-Complex. *Canadian Journal of Experimental Psychology*, 54(4):255–265.

[Baumgartner et al., 1991] Baumgartner, C., Doppelbauer, A., Deecke, L., Barth, D. S., Zeitlhofer, J., Lindinger, G., and Sutherling, W. W. (1991). Neuromagnetic investigation of somatotopy of human hand somatosensory cortex. *Experimental Brain Research*, 87(3):641–648.

[Becker et al., 2018] Becker, C. O., Pequito, S., Pappas, G. J., Miller, M. B., Grafton, S. T., Bassett, D. S., and Preciado, V. M. (2018). Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific Reports*, 8(1):1411.

[Beckmann et al., 2005] Beckmann, C. F., DeLuca, M., Devlin, J. T., and Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1001–1013.

[Behjat and Larsson, 2020] Behjat, H. and Larsson, M. (2020). Spectral characterization of functional MRI data on voxel-resolution cortical graphs. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 558–562. IEEE.

[Behjat et al., 2015] Behjat, H., Leonardi, N., Sörnmo, L., and Van De Ville, D. (2015). Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping. *NeuroImage*, 123:185–199.

[Behjat et al., 2016] Behjat, H., Richter, U., Van De Ville, D., and Sörnmo, L. (2016). Signal-adapted tight frames on graphs. *IEEE Trans. Signal Process.*, 64(22):6017–6029.

[Behjat and Van De Ville, 2019] Behjat, H. and Van De Ville, D. (2019). Spectral design of signal-adapted tight frames on graphs. In *Vertex-Frequency Analysis of Graph Signals*, pages 177–206. Springer.

[Betzel et al., 2016a] Betzel, R. F., Avena-Koenigsberger, A., Goñi, J., He, Y., de Reus, M. A., Griffa, A., Vértes, P. E., Mišic, B., Thiran, J. P., Hagmann, P., van den Heuvel, M., Zuo, X. N., Bullmore, E. T., and Sporns, O. (2016a). Generative models of the human connectome. *NeuroImage*, 124:1054–1064.

[Betzel and Bassett, 2018] Betzel, R. F. and Bassett, D. S. (2018). Specificity and robustness of long-distance connections in weighted, interareal connectomes. *Proceedings of the National Academy of Sciences*, 115(21):E4880–E4889.

[Betzel et al., 2016b] Betzel, R. F., Fukushima, M., He, Y., Zuo, X.-N., and Sporns, O. (2016b). Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *NeuroImage*, 127:287–297.

[Betzel et al., 2016c] Betzel, R. F., Gu, S., Medaglia, J. D., Pasqualetti, F., and Bassett, D. S. (2016c). Optimally controlling the human connectome: The role of network topology. *Scientific Reports*, 6:1–14.

[Biswal et al., 1995] Biswal, B. F., Yetkin, Z., 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*, 34:537–541.

[Block et al., 1946] Block, F., Hansen, W. W., and Packard, M. (1946). The Nuclear Induction Experiment. *Physical Review*, 70(5):474–485.

[Bolton et al., 2018] Bolton, T. A., Farouj, Y., Obertino, S., and Van De Ville, D. (2018). Graph Slepians to strike a balance between local and global network interactions: Application to functional brain imaging. In *2018 IEEE 15th International Symposium on Biomedical Imaging*, pages 1239–1243. IEEE.

[Bolton et al., 2020] Bolton, T. A., Morgenroth, E., Preti, M. G., and Van De Ville, D. (2020). Tapping into Multi-Faceted Human Behavior and Psychopathology Using fMRI Brain Dynamics. *Trends in Neurosciences*, xx(xx):1–14.

[Boly et al., 2012] Boly, M., Perlbarg, V., Marrelec, G., Schabus, M., Laureys, S., Doyon, J., Pelegrini-Issac, M., Maquet, P., and Benali, H. (2012). Hierarchical clustering of brain activity during human nonrapid eye movement sleep. *Proceedings of the National Academy of Sciences*, 109(15):5856–5861.

[Boly et al., 2008] Boly, M., Phillips, C., Tshibanda, L., Vanhaudenhuyse, A., Schabus, M., Dang?Vu, T., Moonen, G., Hustinx, R., Maquet, P., and Laureys, S. (2008). Intrinsic Brain

Activity in Altered States of Consciousness How Conscious Is the Default Mode of Brain Function? *Annals of the New York Academy of Sciences*, 129:119–129.

[Buades et al., 2010] Buades, A., Coll, B., and Morel, J.-m. (2010). A review of image denoising algorithms , with a new one. *SIAM Journal on Multiscale Modeling and Simulation: A SIAM Interdisciplinary Journal*, 4(2):490–530.

[Buckner et al., 2008] Buckner, R. L., Andrews-Hanna, J. R., and Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124:1–38.

[Bufalari et al., 2007] Bufalari, I., Aprile, T., Avenanti, A., Di Russo, F., and Aglioti, S. M. (2007). Empathy for pain and touch in the human somatosensory cortex. *Cerebral Cortex*, 17(11):2553–2561.

[Bullmore and Sporns, 2009] Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198.

[Bullmore and Sporns, 2012] Bullmore, E. and Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336–349.

[Cabral et al., 2011] Cabral, J., Hugues, E., Sporns, O., and Deco, G. (2011). Role of local network oscillations in resting-state functional connectivity. *NeuroImage*, 57(1):130–139.

[Calamante et al., 2017] Calamante, F., Smith, R. E., Liang, X., Zalesky, A., and Connelly, A. (2017). Track-weighted dynamic functional connectivity (TW-dFC): a new method to study time-resolved functional connectivity. *Brain Structure and Function*, 222(8):3761–3774.

[Candès et al., 2006] Candès, E. J., Romberg, J. K., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223.

[Canto et al., 2017] Canto, C. B., Onuki, Y., Bruinsma, B., van der Werf, Y. D., and De Zeeuw, C. I. (2017). The Sleeping Cerebellum. *Trends in Neurosciences*, 40(5):309–323.

[Catani et al., 2002] Catani, M., Howard, R. J., Pajevic, S., and Jones, D. K. (2002). Virtual in Vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage*, 17(1):77–94.

[Cercignani et al., 2000] Cercignani, M., Iannucci, G., Rocca, M., Comi, G., Horsfield, M., and Filippi, M. (2000). Pathologic damage in ms assessed by diffusion-weighted and magnetization transfer mri. *Neurology*, 54(5):1139–1144.

[Chamberland et al., 2015] Chamberland, M., Bernier, M., Fortin, D., Whittingstall, K., and Descoteaux, M. (2015). 3D interactive tractography-informed resting-state fMRI connectivity. *Frontiers in Neuroscience*, 9(JUL):1–15.

[Chamberland et al., 2017] Chamberland, M., Girard, G., Bernier, M., Fortin, D., Descoteaux, M., and Whittingstall, K. (2017). On the Origin of Individual Functional Connectivity Variability: The Role of White Matter Architecture. *Brain Connectivity*, 7(8):491–503.

[Chambolle, 2004] Chambolle, A. (2004). An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20:89–97.

[Chang and Glover, 2010] Chang, C. and Glover, G. H. (2010). Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage*, 50(1):81–98.

[Chen et al., 2015a] Chen, S., Sandryhaila, A., Moura, J. M., and Kovačević, J. (2015a). Signal recovery on graphs: Variation minimization. *IEEE Transactions on Signal Processing*, 63(17):4609–4624.

[Chen et al., 2015b] Chen, S., Sandryhaila, A., Moura, J. M. F., and Kovacevic, J. (2015b). Signal Recovery on Graphs: Variation Minimization. *IEEE Transactions on Signal Processing*, 63(17):4609–4624.

[Cheung et al., 2018] Cheung, G., Magli, E., Tanaka, Y., and Ng, M. K. (2018). Graph Spectral Image Processing. *Proceedings of the IEEE*, 106(5):907–930.

[Chung, 1997] Chung, F. (1997). *Spectral graph theory*. AMS, Providence, RI.

[Cirelli and Tononi, 2008] Cirelli, C. and Tononi, G. (2008). Is Sleep Essential? *PLoS Biology*, 6(8):e216.

[Ciuciu et al., 2012] Ciuciu, P., Varoquaux, G., Abry, P., Sadaghiani, S., and Kleinschmidt, A. (2012). Scale-free and multifractal time dynamics of fmri signals during rest and task. *Frontiers in Physiology*, 3 JUN(June):1–18.

[Conturo et al., 1999] Conturo, T. E., Lori, N. F., Cull, T. S., Akbudak, E., Snyder, A. Z., Shimony, J. S., Mckinstry, R. C., Burton, H., and Raichle, M. E. (1999). Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 96(18):10422–10427.

[Currie et al., 2013] Currie, S., Hoggard, N., Craven, I. J., Hadjivassiliou, M., and Wilkinson, I. D. (2013). Understanding MRI: Basic MR physics for physicians. *Postgraduate Medical Journal*, 89(1050):209–223.

[Damadian et al., 1974] Damadian, R., Zaner, K., Hor, D., and DiMaio, T. (1974). Human tumors detected by nuclear magnetic resonance. *Proceedings of the National Academy of Sciences of the United States of America*, 71(4):1471–1473.

[Damoiseaux and Greicius, 2009] Damoiseaux, J. S. and Greicius, M. D. (2009). Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain Structure and Function*, 213(6):525–533.

[Dang-vu et al., 2010] Dang-vu, T. T., Schabus, M., Desseilles, M., Sterpenich, V., and Bonjean, M. (2010). Functional Neuroimaging Insights into the Physiology of Human Sleep. *Sleep*, 33:1589–1603.

[Deco et al., 2011] Deco, G., Jirsa, V. K., and McIntosh, A. R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12(1):43–56.

[Deco et al., 2013] Deco, G., Ponce-Alvarez, A., Mantini, D., Romani, G. L., Hagmann, P., and Corbetta, M. (2013). Resting-State Functional Connectivity Emerges from Structurally and Dynamically Shaped Slow Linear Fluctuations. *Journal of Neuroscience*, 33(27):11239–11252.

[Deco et al., 2012] Deco, G., Senden, M., and Jirsa, V. (2012). How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in Computational Neuroscience*, 6(September):1–7.

[Deco et al., 2017] Deco, G., Tagliazucchi, E., Laufs, H., Sanjuán, A., and Kringelbach, M. L. (2017). Novel Intrinsic Ignition Method Measuring Local-Global Integration Characterizes Wakefulness and Deep Sleep. *eneuro*, 4(5):ENEURO.0106–17.2017.

[DelRosso and Hoque, 2014] DelRosso, L. M. and Hoque, R. (2014). The Cerebellum and Sleep. *Neurologic Clinics*, 32(4):893–900.

[Dempster, 1972] Dempster, A. P. (1972). Covariance selection. *Biometrics*, pages 157–175.

[Deslauriers-Gauthier et al., 2019] Deslauriers-Gauthier, S., Lina, J.-M., Butler, R., Whittingstall, K., Gilbert, G., Bernier, P.-M., Deriche, R., and Descoteaux, M. (2019). White matter information flow mapping from diffusion MRI and EEG. *NeuroImage*, 201(February):116017.

[Ding et al., 2018] Ding, Z., Huang, Y., Bailey, S. K., Gao, Y., Cutting, L. E., Rogers, B. P., Newton, A. T., and Gore, J. C. (2018). Detection of synchronous brain activity in white matter tracts at rest and under functional loading. *Proceedings of the National Academy of Sciences*, 115(3):595–600.

[Dolan, 2002] Dolan, R. J. (2002). Neuroscience and psychology: Emotion, cognition, and behavior.

[Dong et al., 2016] Dong, X., Thanou, D., Frossard, P., and Vandergheynst, P. (2016). Learning laplacian matrix in smooth graph signal representations. *IEEE Transactions on Signal Processing*, 64(23):6160–6173.

[Dong et al., 2019] Dong, X., Thanou, D., Rabbat, M., and Frossard, P. (2019). Learning graphs from data: A signal representation perspective. *IEEE Signal Processing Magazine*, 36(3):44–63.

[Eavani et al., 2013] Eavani, H., Satterthwaite, T. D., Gur, R. E., Gur, R. C., and Davatzikos, C. (2013). Unsupervised Learning of Functional Network Dynamics in Resting State fMRI. *Inf Process Med Imaging*, 23:426–437.

[Elton and Gao, 2015] Elton, A. and Gao, W. (2015). Task-related modulation of functional connectivity variability and its behavioral correlations. *Human Brain Mapping*, 36(8):3260–3272.

[Ercsey-Ravasz et al., 2013] Ercsey-Ravasz, M., Markov, N. T., Lamy, C., Van Essen, D. C., Knoblauch, K., Toroczkai, Z., and Kennedy, H. (2013). A Predictive Network Model of Cerebral Cortical Connectivity Based on a Distance Rule. *Neuron*, 80(1):184–197.

[Erhardt et al., 2011] Erhardt, E. B., Rachakonda, S., Bedrick, E. J., Allen, E. A., Adali, T., and Calhoun, V. D. (2011). Comparison of multi-subject ICA methods for analysis of fMRI data. *Human Brain Mapping*, 32(12):2075–2095.

[Essen et al., 2013] Essen, D. V., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., and Ugurbil, K. (2013). The WU-Minn human connectome project: An overview. *Neuroimage*, pages 62–79.

[Farouj et al., 2017] Farouj, Y., Karahanoglu, F. I., and Van De Ville, D. (2017). Regularized spatiotemporal deconvolution of fMRI data using gray-matter constrained total variation. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 472–475. IEEE.

[Figley et al., 2015] Figley, T. D., Bhullar, N., Courtney, S. M., and Figley, C. R. (2015). Probabilistic atlases of default mode, executive control and salience network white matter tracts: an fMRI-guided diffusion tensor imaging and tractography study. *Frontiers in Human Neuroscience*, 9(NOVEMBER):1–20.

[Finn et al., 2015] 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.

[Fox et al., 2005] Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678.

[Fox and Lancaster, 2002] Fox, P. T. and Lancaster, J. L. (2002). Mapping context and content: The BrainMap model. *Nature Reviews Neuroscience*, 3(4):319–321.

[Fransson and Marrelec, 2008] Fransson, P. and Marrelec, G. (2008). The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. *NeuroImage*, 42(3):1178–1184.

[Fraser et al., 2012] Fraser, L. M., Stevens, M., Beyea, S. D., and D'Arcy, R. C. N. (2012). White versus gray matter: fMRI hemodynamic responses show similar characteristics, but differ in peak amplitude. *BMC Neuroscience*, 13(1):91.

[Frau-Pascual et al., 2019] Frau-Pascual, A., Fogarty, M., Fischl, B., Yendiki, A., Aganj, I., Initiative, A. D. N., et al. (2019). Quantification of structural brain connectivity via a conductance model. *NeuroImage*, 189:485–496.

[Friman et al., 2006] Friman, O., Farneback, G., and Westin, C.-F. (2006). A Bayesian approach for stochastic white matter tractography. *IEEE Transactions on Medical Imaging*, 25(8):965–978.

[Friston et al., 1999] Friston, K., Zarahn, E., Josephs, O., Henson, R., and Dale, A. (1999). Stochastic Designs in Event-Related fMRI. *NeuroImage*, 10(5):607–619.

[Friston, 1994] Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, pages 56–78.

[Friston, 2011] Friston, K. J. (2011). Functional and Effective Connectivity: A Review. *Brain Connectivity*, 1(1):13–36.

[Friston et al., 1994] Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., and Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2(4):189–210.

[Fukushima et al., 2015] Fukushima, M., Yamashita, O., Knösche, T. R., and Sato, M.-a. (2015). MEG source reconstruction based on identification of directed source interactions on whole-brain anatomical networks. *NeuroImage*, 105:408–427.

[Galán, 2008] Galán, R. F. (2008). On How Network Architecture Determines the Dominant Patterns of Spontaneous Neural Activity. *PLoS ONE*, 3(5):e2148.

[Gao et al., 2012] Gao, Z., van Beugen, B. J., and De Zeeuw, C. I. (2012). Distributed synergistic plasticity and cerebellar learning. *Nature Reviews Neuroscience*, 13(9):619–635.

[Gawryluk et al., 2014] Gawryluk, J. R., Mazerolle, E. L., and D'Arcy, R. C. N. (2014). Does functional MRI detect activation in white matter? A review of emerging evidence, issues, and future directions. *Frontiers in Neuroscience*, 8(8 JUL):1–12.

[Gent et al., 2018] Gent, T. C., Bandarabadi, M., Herrera, C. G., and Adamantidis, A. R. (2018). Thalamic dual control of sleep and wakefulness. *Nature Neuroscience*, 21(7):974–984.

[Giorgio et al., 2017] Giorgio, A., Zhang, J., Stromillo, M. L., Rossi, F., Battaglini, M., Nichelli, L., Mortilla, M., Portaccio, E., Hakiki, B., Amato, M. P., and De Stefano, N. (2017). Pronounced Structural and Functional Damage in Early Adult Pediatric-Onset Multiple Sclerosis with No or Minimal Clinical Disability. *Frontiers in Neurology*, 8(NOV).

[Glass et al., 2008] Glass, H. C., Shaw, G. M., Ma, C., and Sherr, E. H. (2008). Agenesis of the corpus callosum in California 1983-2003: A population-based study. *American Journal of Medical Genetics, Part A*, 146(19):2495–2500.

[Glasser et al., 2013] 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.

[Goni et al., 2014] Goni, J., van den Heuvel, M. P., Avena-Koenigsberger, A., Velez de Mendizabal, N., Betzel, R. F., Griffa, A., Hagmann, P., Corominas-Murtra, B., Thiran, J.-P., and Sporns, O. (2014). Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences*, 111(2):833–838.

[Goodall, 1991] Goodall, C. (1991). Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society*, 53(2):285–339.

[Greicius et al., 2003] 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.

[Greicius et al., 2004] Greicius, M. D., Srivastava, G., Reiss, A. L., and Menon, V. (2004). Default-mode network activity distinguishes Alzheimer's disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences*, 101(13):4637–4642.

[Greicius et al., 2009] Greicius, M. D., Supekar, K., Menon, V., and Dougherty, R. F. (2009). Resting-State Functional Connectivity Reflects Structural Connectivity in the Default Mode Network. *Cerebral Cortex*, 19(1):72–78.

[Gu et al., 2015] Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., Medaglia, J. D., Vettel, J. M., Miller, M. B., Grafton, S. T., and Bassett, D. S. (2015). Controllability of structural brain networks. *Nature Communications*, 6(1):8414.

[Hagmann et al., 2008] Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and Sporns, O. (2008). Mapping the Structural Core of Human Cerebral Cortex. *PLoS Biology*, 6(7):e159.

[Hahn, 1950] Hahn, E. L. (1950). Spin echoes. *Phys. Rev.*, 80:580–594.

[Haimovici et al., 2017] Haimovici, A., Tagliazucchi, E., Balenzuela, P., and Laufs, H. (2017). On wakefulness fluctuations as a source of BOLD functional connectivity dynamics. *Scientific Reports*, 7(1):1–13.

[Hale et al., 2016] Hale, J. R., White, T. P., Mayhew, S. D., Wilson, R. S., Rollings, D. T., Khalsa, S., Arvanitis, T. N., and Bagshaw, A. P. (2016). Altered thalamocortical and intra-thalamic functional connectivity during light sleep compared with wake. *NeuroImage*, 125:657–667.

[Hanlon et al., 2012] Hanlon, F. M., Houck, J. M., Klimaj, S. D., Caprihan, A., Mayer, A. R., Weisend, M. P., Bustillo, J. R., Hamilton, D. A., and Tesche, C. D. (2012). Frontotemporal anatomical connectivity and working-relational memory performance predict everyday functioning in schizophrenia. *Psychophysiology*, 49(10):1340–1352.

[Hariri and Mazziotta, 2000] Hariri, Ahmad, B. S. and Mazziotta, J. (2000). Modulating emotional responses. *NeuroReport*, 11(1):43–48.

[Hasanzadeh et al., 2017] Hasanzadeh, A., Liu, X., Duffield, N., and Narayanan, K. R. (2017). Piecewise Stationary Modeling of Random Processes Over Graphs With an Application to Traffic Prediction. (January 2018).

[He et al., 2010] He, B. J., Zempel, J. M., Snyder, A. Z., and Raichle, M. E. (2010). The temporal structures and functional significance of scale-free brain activity. *Neuron*, 66(3):353–369.

[He et al., 2018] He, K., Stankovic, L., Liao, J., and Stankovic, V. (2018). Non-Intrusive Load Disaggregation Using Graph Signal Processing. *IEEE Transactions on Smart Grid*, 9(3):1739–1747.

[Hegdé and Van Essen, 2000] Hegdé, J. and Van Essen, D. C. (2000). Selectivity for complex shapes in primate visual area V2. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 20(5):1–6.

[Hiroki et al., 2005] Hiroki, M., Uema, T., Kajimura, N., Ogawa, K., Nishikawa, M., Kato, M., Watanabe, T., Nakajima, T., Takano, H., Imabayashi, E., Ohnishi, T., Takayama, Y., Matsuda, H., Uchiyama, M., Okawa, M., Takahashi, K., and Fukuyama, H. (2005). Cerebral white matter blood flow is constant during human non-rapid eye movement sleep: A positron emission tomographic study. *Journal of Applied Physiology*, 98(5):1846–1854.

[Honey et al., 2007] Honey, C. J., Kotter, R., Breakspear, M., and Sporns, O. (2007). Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences*, 104(24):10240–10245.

[Honey et al., 2009] Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040.

[Horn et al., 2014] Horn, A., Ostwald, D., Reisert, M., and Blankenburg, F. (2014). The structural–functional connectome and the default mode network of the human brain. *NeuroImage*, 102(P1):142–151.

[Horovitz et al., 2009] Horovitz, S. G., Braun, A. R., Carr, W. S., Picchioni, D., Balkin, T. J., Fukunaga, M., and Duyn, J. H. (2009). Decoupling of the brain ' s default mode network during deep sleep. *Proceedings of the National Academy of Sciences of the United States of America*.

[Huang et al., 2011] Huang, H., Prince, J. L., Mishra, V., Carass, A., Landman, B., Park, D. C., Tamminga, C., King, R., Miller, M. I., van Zijl, P. C., and Mori, S. (2011). A framework on surface-based connectivity quantification for the human brain. *Journal of Neuroscience Methods*, 197(2):324–332.

[Huang et al., 2018a] Huang, W., Bolton, T. A. W., Medaglia, J. D., Bassett, D. S., Ribeiro, A., and Van De Ville, D. (2018a). A Graph Signal Processing Perspective on Functional Brain Imaging. *Proceedings of the IEEE*, 106(5):868–885.

[Huang et al., 2018b] Huang, Y., Bailey, S. K., Wang, P., Cutting, L. E., Gore, J. C., and Ding, Z. (2018b). Voxel-wise detection of functional networks in white matter. *NeuroImage*, 183:544–552.

[Iturria-Medina et al., 2007] Iturria-Medina, Y., Canales-Rodríguez, E., Melie-García, L., Valdés-Hernández, P., Martínez-Montes, E., Alemán-Gómez, Y., and Sánchez-Bornot, J. (2007). Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage*, 36(3):645–660.

[Jablonski, 2017] Jablonski, I. (2017). Graph Signal Processing in Applications to Sensor Networks, Smart Grids, and Smart Cities. *IEEE Sensors Journal*, 17(23):7659–7666.

[Jahnke et al., 2012] Jahnke, K., von Wegner, F., Morzelewski, A., Borisov, S., Maischein, M., Steinmetz, H., and Laufs, H. (2012). To wake or not to wake? The two-sided nature of the human K-complex. *NeuroImage*, 59(2):1631–1638.

[Jain et al., 2014] Jain, R. K., Moura, J. M., and Kontokosta, C. E. (2014). Big Data + Big Cities: Graph Signals of Urban Air Pollution [Exploratory SP]. *IEEE Signal Processing Magazine*, 31(5):130–136.

[Jbabdi et al., 2015] Jbabdi, S., Sotiropoulos, S. N., Haber, S. N., Van Essen, D. C., and Behrens, T. E. (2015). Measuring macroscopic brain connections in vivo. *Nature Neuroscience*, 18(11):1546–1555.

[Jellison et al., 2004] Jellison, B., Field, A., Medow, J., Lazar, M., Salamat, M., and Alexander, A. (2004). Diffusion tensor imaging of cerebral white matter: A pictorial review of physics, fiber tract anatomy, and tumor imaging patterns. *American Journal of Neuroradiology*, 25(3):356–369. cited By 457.

[Jobst et al., 2017] Jobst, B. M., Hindriks, R., Laufs, H., Tagliazucchi, E., Hahn, G., Ponce-Alvarez, A., Stevner, A. B., Kringelbach, M. L., and Deco, G. (2017). Increased Stability and Breakdown of Brain Effective Connectivity during Slow-Wave Sleep: Mechanistic Insights from Whole-Brain Computational Modelling. *Scientific Reports*, 7(1):1–16.

[Jones et al., 2013] Jones, D. K., Knösche, T. R., and Turner, R. (2013). White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI. *NeuroImage*, 73:239–254.

[Jung et al., 2019] Jung, K. I., Park, M. H., Park, B., Kim, S. Y., Kim, Y. O., Kim, B. N., Park, S., and Song, C. H. (2019). Cerebellar Gray Matter Volume, Executive Function, and Insomnia: Gender Differences in Adolescents. *Scientific Reports*, 9(1):1–9.

[Karahanoglu and Van De Ville, 2015] Karahanoglu, F. I. and Van De Ville, D. (2015). Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nature Communications*, 6(1):7751.

[Karahanoğlu et al., 2013] Karahanoğlu, F. I., Caballero-Gaudes, C., Lazeyras, F., and Van De Ville, D. (2013). Total activation: fMRI deconvolution through spatio-temporal regularization. *NeuroImage*, 73:121–134.

[Kaufmann et al., 2006] Kaufmann, C., Wehrle, R., Wetter, T. C., Holsboer, F., Auer, D. P., Pollmächer, T., and Czisch, M. (2006). Brain activation and hypothalamic functional connectivity during human non-rapid eye movement sleep: An EEG/fMRI study. *Brain*, 129(3):655–667.

[Kendall, 1989] Kendall, D. G. (1989). A Survey of the Statistical Theory of Shape. *Statistical Science*, 4(2):87–120.

[Kinany et al., 2020] Kinany, N., Pirondini, E., Micera, S., and Van De Ville, D. (2020). Dynamic Functional Connectivity of Resting-State Spinal Cord fMRI Reveals Fine-Grained Intrinsic Architecture. *Neuron*, 108(3):424–435.e4.

[Koch et al., 2002] Koch, M. A., Norris, D. G., and Hund-Georgiadis, M. (2002). An Investigation of Functional and Anatomical Connectivity Using Magnetic Resonance Imaging. *NeuroImage*, 16(1):241–250.

[Kondor and Lafferty, 2002] Kondor, R. I. and Lafferty, J. D. (2002). Diffusion Kernels on Graphs and Other Discrete Input Spaces. *Proceedings of the Nineteenth International Conference on Machine Learning*.

[Krohn and Ostwald, 2017] Krohn, S. and Ostwald, D. (2017). Computing integrated information. *Neuroscience of Consciousness*, 2017(1).

[Kropf et al., 2019] Kropf, E., Syan, S. K., Minuzzi, L., and Frey, B. N. (2019). From anatomy to function: the role of the somatosensory cortex in emotional regulation. *Revista brasileira de psiquiatria (Sao Paulo, Brazil : 1999)*, 41(3):261–269.

[Kucyi and Davis, 2014] Kucyi, A. and Davis, K. D. (2014). Dynamic functional connectivity of the default mode network tracks daydreaming. *NeuroImage*, 100:471–480.

[Kuhn,] Kuhn, H. W.

[Kumar et al., 1975] Kumar, A., Welti, D., and Ernst, R. R. (1975). NMR Fourier zeugmatography. *Journal of Magnetic Resonance (1969)*, 18(1):69–83.

[Kung et al., 2019] Kung, Y.-c., Wu, C. W., Shuo, C.-w. L., Sharon, C., Biswal, B., Lane, C.-y. Z. L. T. J., and Lin, C.-p. (2019). Instability of brain connectivity during nonrapid eye movement sleep reflects altered properties of information integration. *Human Brain Mapping*, (September 2018):3192–3202.

[Larson-Prior et al., 2009] Larson-Prior, L. J., Zempel, J. M., Nolan, T. S., Prior, F. W., Snyder, A. Z., and Raichle, M. E. (2009). Cortical network functional connectivity in the descent to sleep. *Proceedings of the National Academy of Sciences*, 106(11):4489–4494.

[Lauterbur, 1973] Lauterbur, P. C. (1973). Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance. *Nature*, 242:190–191.

[Le Bihan, 2003] Le Bihan, D. (2003). Looking into the functional architecture of the brain with diffusion MRI. *Nature Reviews Neuroscience*, 4(6):469–480.

[Leech et al., 2011] Leech, R., Kamourieh, S., Beckmann, C. F., and Sharp, D. J. (2011). Fractionating the Default Mode Network: Distinct Contributions of the Ventral and Dorsal Posterior Cingulate Cortex to Cognitive Control. *Journal of Neuroscience*, 31(9):3217–3224.

[Leonardi et al., 2013] Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-m., Schluep, M., Vuilleumier, P., and Van De Ville, D. (2013). Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest. *NeuroImage*, 83:937–950.

[Leonardi and Van De Ville, 2015] Leonardi, N. and Van De Ville, D. (2015). On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage*, 104:430–436.

[Lerner et al., 2014] Lerner, A., Mogensen, M. A., Kim, P. E., Shiroishi, M. S., Hwang, D. H., and Law, M. (2014). Clinical Applications of Diffusion Tensor Imaging. *World Neurosurgery*, 82(1-2):96–109.

[Lewis et al., 2009] Lewis, J. D., Theilmann, R. J., Sereno, M. I., and Townsend, J. (2009). The relation between connection length and degree of connectivity in young adults: A DTI analysis. *Cerebral Cortex*, 19(3):554–562.

[Li et al., 2019] Li, M., Newton, A. T., Anderson, A. W., Ding, Z., and Gore, J. C. (2019). Characterization of the hemodynamic response function in white matter tracts for event-related fMRI. *Nature Communications*, 10(1):1140.

[Liao et al., 2014] Liao, W., Zhang, Z., Mantini, D., Xu, Q., Ji, G. J., Zhang, H., Wang, J., Wang, Z., Chen, G., Tian, L., Jiao, Q., Zang, Y. F., and Lu, G. (2014). Dynamical intrinsic functional architecture of the brain during absence seizures. *Brain Structure and Function*, 219(6):2001–2015.

[Liégeois et al., 2020] Liégeois, R., Santos, A., Matta, V., Van De Ville, D., and Sayed, A. H. (2020). Revisiting correlation-based functional connectivity and its relationship with structural connectivity. *Network Neuroscience*, pages 1–25.

[Liégeois et al., 2016] Liégeois, R., Ziegler, E., Phillips, C., Geurts, P., Gómez, F., Bahri, M. A., Yeo, B. T. T., Soddu, A., Vanhaudenhuyse, A., Laureys, S., and Sepulchre, R. (2016). Cerebral functional connectivity periodically (de)synchronizes with anatomical constraints. *Brain Structure and Function*, 221(6):2985–2997.

[Liu and Duyn, 2013] Liu, X. and Duyn, J. H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, 110(11):4392–4397.

[Logothetis et al., 2001] Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843):150–157.

[Logothetis and Wandell, 2004] Logothetis, N. K. and Wandell, B. A. (2004). Interpreting the BOLD Signal. *Annual Review of Physiology*, 66(1):735–769.

[Madhyastha and Grabowski, 2014] Madhyastha, T. M. and Grabowski, T. J. (2014). Age-related differences in the dynamic architecture of intrinsic networks. *Brain connectivity*, 4(4):231–241.

[Maghsadhagh et al., 2019] Maghsadhagh, S., Eklund, A., and Behjat, H. (2019). Graph spectral characterization of brain cortical morphology. In *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc.*, pages 458–462.

[Maier-Hein et al., 2017] Maier-Hein, K. H., Neher, P. F., Houde, J.-C., Côté, M.-A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.-C., Lin, Y.-C., Ji, Q., Reddick, W. E., Glass, J. O., Chen, D. Q., Feng, Y., Gao, C., Wu, Y., Ma, J., He, R., Li, Q., Westin, C.-F., Deslauriers-Gauthier, S., González, J. O. O., Paquette, M., St-Jean, S., Girard, G., Rheault, F., Sidhu, J., Tax, C. M. W., Guo, F., Mesri, H. Y., Dávid, S., Froeling, M., Heemskerk, A. M., Leemans, A., Boré, A., Pinsard, B., Bedetti, C., Desrosiers, M., Brambati, S., Doyon, J., Sarica, A., Vasta, R., Cerasa, A., Quattrone, A., Yeatman, J., Khan, A. R., Hodges, W., Alexander, S., Romascano, D., Barakovic, M., Auría, A., Esteban, O., Lemkaddem, A., Thiran, J.-P., Cetingul, H. E., Odry, B. L., Mailhe, B., Nadar, M. S., Pizzagalli, F., Prasad, G., Villalon-Reina, J. E., Galvis, J., Thompson, P. M., Requejo, F. D. S., Laguna, P. L., Lacerda, L. M., Barrett, R., Dell'Acqua, F., Catani, M., Petit, L., Caruyer, E., Daducci, A., Dyrby, T. B., Holland-Letz, T., Hilgetag, C. C., Stieltjes, B., and Descoteaux, M. (2017). The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8(1):1349.

[Mansfield, 1977] Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3):L55–L58.

[Margulies et al., 2016] Margulies, D. S., Ghoshb, S. S., Goulasd, A., Falkiewicza, M., Huntenburga, J. M., Langsf, G., Bezginh, G., Eickhoffi, S. B., Castellanosk, F. X., Petridesm, M., Jefferiesn, E., and Smallwoodn, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *PNAS*, 113(44):12574–12579.

[Markov et al., 2011] Markov, N. T., Misery, P., Falchier, A., Lamy, C., Vezoli, J., Quilodran, R., Gariel, M. A., Giroud, P., Ercsey-Ravasz, M., Pilaz, L. J., Huissoud, C., Barone, P., Dehay, C., Toroczkai, Z., Van Essen, D. C., Kennedy, H., and Knoblauch, K. (2011). Weight consistency specifies regularities of macaque cortical networks. *Cerebral Cortex*, 21(6):1254–1272.

[Mars et al., 2018] Mars, R. B., Passingham, R. E., and Jbabdi, S. (2018). Connectivity Fingerprints: From Areal Descriptions to Abstract Spaces. *Trends in Cognitive Sciences*, 22(11):1026–1037.

[Masoumi et al., 2020] Masoumi, M., Marcoux, M., Maignel, L., and Pomar, C. (2020). Spectralweight: a spectral graph wavelet framework for weight prediction of pork cuts. *arXiv preprint arXiv:2005.05406*.

[Maxim et al., 2005] Maxim, V., Aendur, L., Fadili, J., Suckling, J., Gould, R., Howard, R., and Bullmore, E. (2005). Fractional gaussian noise, functional mri and alzheimer's disease. *NeuroImage*, 25(1):141 – 158.

[Medaglia et al., 2018] Medaglia, J. D., Huang, W., Karuza, E. A., Kelkar, A., Thompson-Schill, S. L., Ribeiro, A., and Bassett, D. S. (2018). Functional alignment with anatomical networks is associated with cognitive flexibility. *Nature Human Behaviour*, 2(2):156–164.

[Mesulam, 1998] Mesulam, M. M. (1998). From sensation to cognition. *Brain*, 121(6):1013–1052.

[Monti et al., 2013] Monti, M. M., Lutkenhoff, E. S., Rubinov, M., Boveroux, P., Vanhaudenhuyse, A., Gosseries, O., Bruno, M. A., Noirhomme, Q., Boly, M., and Laureys, S. (2013). Dynamic Change of Global and Local Information Processing in Propofol-Induced Loss and Recovery of Consciousness. *PLoS Computational Biology*, 9(10).

[Monti et al., 2003] Monti, S., Tamayo, P., Mesirov, J., and Golub, T. (2003). Consensus clustering: A resampling-based method for class discovery and visualization of gene expression microarray data. *Machine Learning*, 52(1-2):91–118.

[Mueller et al., 2013] Mueller, S., Wang, D., Fox, M. D., Yeo, B. T., Sepulcre, J., Sabuncu, M. R., Shafee, R., Lu, J., and Liu, H. (2013). Individual Variability in Functional Connectivity Architecture of the Human Brain. *Neuron*, 77(3):586–595.

[Mugler III and Brookeman, 1991] Mugler III, J. P. and Brookeman, J. R. (1991). Rapid three-dimensional t1-weighted mr imaging with the mp-rage sequence. *Journal of Magnetic Resonance Imaging*, 1(5):561–567.

[Nir et al., 2013] Nir, T. M., Jahanshad, N., Villalon-Reina, J. E., Toga, A. W., Jack, C. R., Weiner, M. W., and Thompson, P. M. (2013). Effectiveness of regional DTI measures in distinguishing Alzheimer's disease, MCI, and normal aging. *NeuroImage: Clinical*, 3:180–195.

[Nofzinger et al., 2013] Nofzinger, E., Maquet, P., and Thorpy, M. J. (2013). *Neuroimaging of sleep and sleep disorders*. Cambridge, UK: Cambridge University Press.

[Ogawa et al., 1990] Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872.

[Oligschläger et al., 2017] Oligschläger, S., Huntenburg, J. M., Golchert, J., Lauckner, M. E., Bonnen, T., and Margulies, D. S. (2017). Gradients of connectivity distance are anchored in primary cortex. *Brain Structure and Function*, 222(5):2173–2182.

[Oligschläger et al., 2019] Oligschläger, S., Xu, T., Baczkowski, B. M., Falkiewicz, M., Falchier, A., Linn, G., and Margulies, D. S. (2019). Gradients of connectivity distance in the cerebral cortex of the macaque monkey. *Brain Structure and Function*, 224(2):925–935.

[O'Muircheartaigh and Jbabdi, 2018] O'Muircheartaigh, J. and Jbabdi, S. (2018). Concurrent white matter bundles and grey matter networks using independent component analysis. *NeuroImage*, 170(May 2017):296–306.

[Ortega et al., 2018] Ortega, A., Frossard, P., Kovačević, J., Moura, J. M. F., and Vandergheynst, P. (2018). Graph signal processing: Overview, challenges, and applications. *Proc. IEEE*, 106(5):808–828.

[Ouyang et al., 2016] Ouyang, M., Jeon, T., Mishra, V., Du, H., Wang, Y., Peng, Y., and Huang, H. (2016). Global and regional cortical connectivity maturation index (CCMI) of developmental human brain with quantification of short-range association tracts. *Medical Imaging 2016: Biomedical Applications in Molecular, Structural, and Functional Imaging*, 9788(March 2016):97881B.

[Owen et al., 2013] Owen, J. P., Li, Y.-O., Yang, F. G., Shetty, C., Bukshpun, P., Vora, S., Wakahiro, M., Hinkley, L. B., Nagarajan, S. S., Sherr, E. H., and Mukherjee, P. (2013). Resting-State Networks and the Functional Connectome of the Human Brain in Agenesis of the Corpus Callosum. *Brain Connectivity*, 3(6):547–562.

[Padula et al., 2017] Padula, M. C., Schaer, M., Scariati, E., Mutlu, A. K., Zöller, D., Schneider, M., and Eliez, S. (2017). Quantifying indices of short- and long-range white matter connectivity at each cortical vertex. *PLoS ONE*, 12(11):1–20.

[Peer et al., 2017] Peer, M., Nitzan, M., Bick, A. S., Levin, N., and Arzy, S. (2017). Evidence for Functional Networks within the Human Brain's White Matter. *The Journal of Neuroscience*, 37(27):6394–6407.

[Petric Maretic and Frossard, 2020] Petric Maretic, H. and Frossard, P. (2020). Graph Laplacian Mixture Model. *IEEE Transactions on Signal and Information Processing over Networks*, (c):1–1.

[Petrovic et al., 2019] Petrovic, M., Bolton, T. A. W., Preti, M. G., Liégeois, R., and Van De Ville, D. (2019). Guided graph spectral embedding: Application to the C. elegans connectome. *Network Neuroscience*, 3(3):807–826.

[Picchioni et al., 2013] Picchioni, D., Duyn, J. H., and Horovitz, S. G. (2013). Sleep and the functional connectome ? *NeuroImage*, 80:387–396.

[Picchioni et al., 2014] Picchioni, D., Pixa, M. L., Fukunaga, M., Carr, W. S., Horovitz, S. G., Braun, A. R., and Duyn, J. H. (2014). Decreased Connectivity between the Thalamus and the Neocortex during Human Nonrapid Eye Movement Sleep. *Sleep*, 37(2):387–397.

[Piccoli et al., 2015] Piccoli, T., Valente, G., Linden, D. E. J., Re, M., Esposito, F., Sack, A. T., and Salle, F. D. (2015). The Default Mode Network and the Working Memory Network Are Not Anti-Correlated during All Phases of a Working Memory Task. *PLOS ONE*, 10(4):e0123354.

[Power et al., 2012] Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., and Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*, 59(3):2142–2154.

[Preckel et al., 2019] Preckel, K., Trautwein, F. M., Paulus, F. M., Kirsch, P., Krach, S., Singer, T., and Kanske, P. (2019). Neural mechanisms of affective matching across faces and scenes. *Scientific Reports*, 9(1):1–10.

[Preti et al., 2017] Preti, M. G., Bolton, T. A., and Van De Ville, D. (2017). The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage*, 160(December):41–54.

[Preti and Van De Ville, 2019] Preti, M. G. and Van De Ville, D. (2019). Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nature Communications*, 10.

[Raj et al., 2012] Raj, A., Kuceyeski, A., and Weiner, M. (2012). A Network Diffusion Model of Disease Progression in Dementia. *Neuron*, 73(6):1204–1215.

[Robinson et al., 2016] Robinson, P., Zhao, X., Aquino, K., Griffiths, J., Sarkar, S., and Mehta-Pandejee, G. (2016). Eigenmodes of brain activity: Neural field theory predictions and comparison with experiment. *NeuroImage*, 142:79–98.

[Robinson, 2012] Robinson, P. A. (2012). Interrelating anatomical, effective, and functional brain connectivity using propagators and neural field theory. *Physical Review E*, 85(1):011912.

[Romani et al., 1982] Romani, G., Williamson, S., and Kaufman, L. (1982). Tonotopic organization of the human auditory cortex. *Science*, 216(4552):1339–1340.

[Rudin et al., 1992] Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268.

[Salvador et al., 2005] Salvador, R., Suckling, J., Coleman, M. R., Pickard, J. D., Menon, D., and Bullmore, E. (2005). Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*, 15(9):1332–2342.

[Sämann et al., 2011] Sämann, P. G., Wehrle, R., Hoehn, D., Spoormaker, V. I., Peters, H., Tully, C., Holsboer, F., and Czisch, M. (2011). Development of the Brain's Default Mode Network from Wakefulness to Slow Wave Sleep. *Cerebral Cortex*, 21(9):2082–2093.

[Sandryhaila and Moura, 2014] Sandryhaila, A. and Moura, J. M. F. (2014). Discrete Signal Processing on Graphs: Frequency Analysis. *IEEE Transactions on Signal Processing*, 62(12):3042–3054.

[Saper et al., 2010] Saper, C. B., Fuller, P. M., Pedersen, N. P., Lu, J., and Scammell, T. E. (2010). Sleep State Switching. *Neuron*, 68(6):1023–1042.

[Sepulcre et al., 2010] Sepulcre, J., Liu, H., Talukdar, T., Martincorena, I., Yeo, B. T. T., and Buckner, R. L. (2010). The Organization of Local and Distant Functional Connectivity in the Human Brain. *PLoS Computational Biology*, 6(6):e1000808.

[Shirer et al., 2012] Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., and Greicius, M. D. (2012). Decoding Subject-Driven Cognitive States with Whole-Brain Connectivity Patterns. *Cerebral Cortex*, 22(1):158–165.

[Shuman et al., 2013] Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98.

[Smith et al., 2009] 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 of the United States of America*, 106(31):13040–13045.

[Smith et al., 2012] Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., Van Essen, D. C., Feinberg, D. A., Yacoubb, E. S., and Ugurbil, K. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences of the United States of America*, 109(8):3131–3136.

[Smith et al., 2011] Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D., and Woolrich, M. W. (2011). Network modelling methods for FMRI. *NeuroImage*, 54(2):875–891.

[Smitha et al., 2017] Smitha, K., Akhil Raja, K., Arun, K., Rajesh, P., Thomas, B., Kapilamoorthy, T., and Kesavadas, C. (2017). Resting state fMRI: A review on methods in resting state connectivity analysis and resting state networks. *The Neuroradiology Journal*, 30(4):305–317.

[Spoormaker et al., 2012] Spoormaker, V. I., Gleiser, P. M., and Czisch, M. (2012). Frontoparietal connectivity and hierarchical structure of the brain's functional network during sleep. *Frontiers in Neurology*, MAY(May):1–10.

[Spoormaker et al., 2010] Spoormaker, V. I., Schroter, M. S., Gleiser, P. M., Andrade, K. C., Dresler, M., Wehrle, R., Samann, P. G., and Czisch, M. (2010). Development of a Large-Scale Functional Brain Network during Human Non-Rapid Eye Movement Sleep. *Journal of Neuroscience*, 30(34):11379–11387.

[Sporns, 2018] Sporns, O. (2018). Graph theory methods: applications in brain networks.

[Sporns and Zwi, 2004] Sporns, O. and Zwi, J. D. (2004). The Small World of the Cerebral Cortex. *Neuroinformatics*, 2(2):145–162.

[Squire et al., 2015] Squire, L. R., Genzel, L., Wixted, J. T., and Morris, R. G. (2015). Memory Consolidation. *Cold Spring Harbor Perspectives in Biology*, 7(8):a021766.

[Staempfli et al., 2006] Staempfli, P., Jaermann, T., Crelier, G., Kollias, S., Valavanis, A., and Boesiger, P. (2006). Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. *NeuroImage*, 30(1):110–120.

[Stankovic et al., 2020] Stankovic, L., Mandic, D., Dakovic, M., Brajovic, M., Scalzo, B., Li, S., and Constantinides, A. G. (2020). Graph signal processing–part iii: Machine learning on graphs, from graph topology to applications. *arXiv preprint arXiv:2001.00426*.

[Stankovic et al., 2019] Stankovic, L., Mandic, D. P., Dakovic, M., Kisil, I., Sejdic, E., and Constantinides, A. G. (2019). Understanding the Basis of Graph Signal Processing via an Intuitive Example-Driven Approach [Lecture Notes]. *IEEE Signal Processing Magazine*, 36(6):133–145.

[Stevner et al., 2019] Stevner, A. B. A., Vidaurre, D., Cabral, J., Rapuano, K., Nielsen, S. F. V., Tagliazucchi, E., Laufs, H., Vuust, P., Deco, G., Woolrich, M. W., Van Someren, E., and Kringelbach, M. L. (2019). Discovery of key whole-brain transitions and dynamics during human wakefulness and non-REM sleep. *Nature Communications*, 10(1):1035.

[Stewart, 2001] Stewart, G. W. (2001). A krylov-schur algorithm for large eigenproblems. *SIAM Journal of Matrix Analysis and Applications*, 23(3):601–614.

[Stoodley and Schmahmann, 2010] Stoodley, C. J. and Schmahmann, J. D. (2010). Evidence for topographic organization in the cerebellum of motor control versus cognitive and affective processing. *Cortex*, 46(7):831–844.

[Sui et al., 2013] Sui, J., He, H., Pearlson, G. D., Adali, T., Kiehl, K. A., Yu, Q., Clark, V. P., Castro, E., White, T., Mueller, B. A., Ho, B. C., Andreasen, N. C., and Calhoun, V. D. (2013). Three-way (N-way) fusion of brain imaging data based on mCCA+jICA and its application to discriminating schizophrenia. *NeuroImage*, 66:119–132.

[Sui et al., 2015] Sui, J., Pearlson, G. D., Du, Y., Yu, Q., Jones, T. R., Chen, J., Jiang, T., Bustillo, J., and Calhoun, V. D. (2015). In Search of Multimodal Neuroimaging Biomarkers of Cognitive Deficits in Schizophrenia. *Biological Psychiatry*, 78(11):794–804.

[Supekar et al., 2010] Supekar, K., Uddin, L. Q., Prater, K., Amin, H., Greicius, M. D., and Menon, V. (2010). Development of functional and structural connectivity within the default mode network in young children. *NeuroImage*, 52(1):290–301.

[Tagliazucchi et al., 2010] Tagliazucchi, E., Balenzuela, P., Fraiman, D., and Chialvo, D. R. (2010). Brain resting state is disrupted in chronic back pain patients. *Neuroscience Letters*, 485(1):26–31.

[Tagliazucchi et al., 2012a] Tagliazucchi, E., Balenzuela, P., Fraiman, D., and Chialvo, D. R. (2012a). Criticality in Large-Scale Brain fMRI Dynamics Unveiled by a Novel Point Process Analysis. *Frontiers in Physiology*, 3(February):1–12.

[Tagliazucchi et al., 2016] Tagliazucchi, E., Crossley, N., Bullmore, E. T., and Laufs, H. (2016). Deep sleep divides the cortex into opposite modes of anatomical-functional coupling. *Brain Structure and Function*, 221(8):4221–4234.

[Tagliazucchi and Laufs, 2014] Tagliazucchi, E. and Laufs, H. (2014). Decoding Wakefulness Levels from Typical fMRI Resting-State Data Reveals Reliable Drifts between Wakefulness and Sleep. *Neuron*, 82(3):695–708.

[Tagliazucchi and van Someren, 2017] Tagliazucchi, E. and van Someren, E. J. (2017). The large-scale functional connectivity correlates of consciousness and arousal during the healthy and pathological human sleep cycle. *NeuroImage*, 160(June):55–72.

[Tagliazucchi et al., 2012b] Tagliazucchi, E., Von Wegner, F., Morzelewski, A., Brodbeck, V., and Laufs, H. (2012b). Dynamic bold functional connectivity in humans and its electrophysiological correlates. *Frontiers in human neuroscience*, 6:339.

[Tagliazucchi et al., 2013] Tagliazucchi, E., Wegner, F. V., Morzelewski, A., Brodbeck, V., Jahnke, K., and Laufs, H. (2013). Breakdown of long-range temporal dependence in default mode and attention networks during deep sleep. *Proceedings of the National Academy of Sciences*, (31).

[Tarun et al., 2019] Tarun, A., Abramian, D., Behjat, H., and De Ville, D. V. (2019). Graph Spectral Analysis of Voxel-Wise Brain Graphs from Diffusion-Weighted MRI. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, number Isbi, pages 159–163. IEEE.

[Tarun et al., 2020a] Tarun, A., Abramian, D., Larsson, M., Behjat, H., and Van De Ville, D. (2020a). Voxel-Wise Brain Graphs from Diffusion-Weighted MRI: Spectral Analysis and Application to Functional MRI. *In Revisions, IEEE Transactions on Medical Imaging*.

[Tarun et al., 2020b] Tarun, A., Behjat, H., Bolton, T., Abramian, D., and Van De Ville, D. (2020b). Structural mediation of human brain activity revealed by white-matter interpolation of fMRI. *NeuroImage*, 213:116718.

[Tarun et al., 2020c] Tarun, A., Ricchi, I., Petric Maretic, H., Frossard, P., and Van De Ville, D. (2020c). Dynamics of Functional Network Organization Revealed by Graph Laplacian Mixture Models. *Preprint*.

[Tarun and Van De Ville, 2020] Tarun, A. and Van De Ville, D. (2020). Anatomical range of functional interactions mean- ingfully differentiates brain regions. *Preprint*.

[Tarun et al., 2020d] Tarun, A., Wainstein-Andriano, D., Sterpenich, V., Bayer, L., Perogamvros, L., Solms, M., Axmacher, N., Schwartz, S., and Van De Ville, D. (2020d). NREM sleep stages specifically alter dynamical integration of large-scale brain networks. *iScience*, pages 1–30.

[Thomas Yeo et al., 2011] 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., Fisch, 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.

[Tiemeier et al., 2010] Tiemeier, H., Lenroot, R. K., Greenstein, D. K., Tran, L., Pierson, R., and Giedd, J. N. (2010). Cerebellum development during childhood and adolescence: A longitudinal morphometric MRI study. *NeuroImage*, 49(1):63–70.

[Tononi, 2004] Tononi, G. (2004). An information integration theory of consciousness An information integration theory of consciousness. *BMC Neuroscience*, 5:42–64.

[Tononi, 2008] Tononi, G. (2008). Consciousness as integrated information: A provisional manifesto. *Biological Bulletin*, 215(3):216–242.

[Tononi et al., 2016] Tononi, G., Boly, M., Massimini, M., and Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7):450–461.

[Tononi and Koch, 2015] Tononi, G. and Koch, C. (2015). Consciousness: here, there and everywhere? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668):20140167.

[Tootell et al., 1998] Tootell, R. B., Hadjikhani, N. K., Mendola, J. D., Marrett, S., and Dale, A. M. (1998). From retinotopy to recognition: fMRI in human visual cortex. *Trends in Cognitive Sciences*, 2(5):174–183.

[Tovar-Moll et al., 2014] Tovar-Moll, F., Monteiro, M., Andrade, J., Bramati, I. E., Vianna-Barbosa, R., Marins, T., Rodrigues, E., Dantas, N., Behrens, T. E. J., de Oliveira-Souza, R., Moll, J., and Lent, R. (2014). Structural and functional brain rewiring clarifies preserved interhemispheric transfer in humans born without the corpus callosum. *Proceedings of the National Academy of Sciences*, 111(21):7843–7848.

[Tozer et al., 2012] Tozer, D. J., Chard, D. T., Bodini, B., Ciccarelli, O., Miller, D. H., Thompson, A. J., and Wheeler-Kingshott, C. A. (2012). Linking white matter tracts to associated cortical grey matter: A tract extension methodology. *NeuroImage*, 59(4):3094–3102.

[Tremblay and Borgnat, 2014] Tremblay, N. and Borgnat, P. (2014). Graph Wavelets for Multi-scale Community Mining. *IEEE Transactions on Signal Processing*, 62(20):5227–5239.

[Tuch, 2004] Tuch, D. S. (2004). Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372.

[Tuch et al., 2003] Tuch, D. S., Reese, T. G., Wiegell, M. R., and Van J. Wedeen (2003). Diffusion MRI of Complex Neural Architecture. *Neuron*, 40(5):885–895.

[Tzikas et al., 2008] Tzikas, D. G., Likas, A. C., and Galatsanos, N. P. (2008). The variational approximation for Bayesian inference. *IEEE Signal Processing Magazine*, 25(6):131–146.

[Tzourio-Mazoyer et al., 2002] 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.

[Van De Ville et al., 2010] Van De Ville, D., Britz, J., and Michel, C. M. (2010). EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. *Proceedings of the National Academy of Sciences*, 107(42):18179–18184.

[Van De Ville et al., 2017] Van De Ville, D., Demesmaeker, R., and Preti, M. G. (2017). When slepian meets fiedler: Putting a focus on the graph spectrum. *IEEE Signal Processing Letters*, 24(7):1001–1004.

[van den Heuvel and Hulshoff Pol, 2010] van den Heuvel, M. P. and Hulshoff Pol, H. E. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8):519–534.

[van den Heuvel et al., 2009] van den Heuvel, M. P., Mandl, R. C., Kahn, R. S., and Hulshoff Pol, H. E. (2009). Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Human Brain Mapping*, 30(10):3127–3141.

[van den Heuvel and Sporns, 2011] van den Heuvel, M. P. and Sporns, O. (2011). Rich-club organization of the human connectome. *Journal of Neuroscience*, 31(44):15775–15786.

[Van Dijk et al., 2010] Van Dijk, K. R. A., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., and Buckner, R. L. (2010). Intrinsic Functional Connectivity As a Tool For Human Connectomics: Theory, Properties, and Optimization. *Journal of Neurophysiology*, 103(1):297–321.

[Venkatesh et al., 2020] Venkatesh, M., Jaja, J., and Pessoa, L. (2020). Comparing functional connectivity matrices: A geometry-aware approach applied to participant identification. *NeuroImage*, 207(June 2019):116398.

[Vidaurre et al., 2017] Vidaurre, D., Smith, S. M., and Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences of the United States of America*, 114(48):12827–12832.

[Wagstyl et al., 2015] Wagstyl, K., Ronan, L., Goodyer, I. M., and Fletcher, P. C. (2015). Cortical thickness gradients in structural hierarchies. *NeuroImage*, 111:241–250.

[Wang et al., 2010] Wang, J., Zuo, X., and He, Y. (2010). Graph-based network analysis of resting-state functional MRI. *Frontiers in Systems Neuroscience*, 4(June):1–14.

[Wang et al., 2017] Wang, M. B., Owen, J. P., Mukherjee, P., and Raj, A. (2017). Brain network eigenmodes provide a robust and compact representation of the structural connectome in health and disease. *PLOS Computational Biology*, 13(6):e1005550.

[Wang et al., 2016] Wang, Y., Kang, J., Kemmer, P. B., and Guo, Y. (2016). An efficient and reliable statistical method for estimating functional connectivity in large scale brain networks using partial correlation. *Frontiers in Neuroscience*, 10(MAR).

[Wheeler-Kingshott and Cercignani, 2009] Wheeler-Kingshott, C. A. and Cercignani, M. (2009). About "axial" and "radial" diffusivities. *Magnetic Resonance in Medicine*, 61(5):1255–1260.

[Winkler et al., 1975] Winkler, M., Ortendahl, D. A., Mills, T. C., Crooks, L. E., Sheldon, P. E., Kaufman, L., and Kramer, D. M. (1975). Characteristics of partial flip angle and gradient reversal MR imaging. *Radiology*, 166(1):17–26.

[Yang et al., 2014] Yang, Z., Craddock, R. C., Margulies, D. S., Yan, C.-G., and Milham, M. P. (2014). Common intrinsic connectivity states among posteromedial cortex subdivisions: Insights from analysis of temporal dynamics. *NeuroImage*, 93(P1):124–137.

[Yarkoni et al., 2011] Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., and Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, 8(8):665–670.

[Yeh et al., 2016] Yeh, F.-C., Vettel, J. M., Singh, A., Poczos, B., Grafton, S. T., Erickson, K. I., Tseng, W.-Y. I., and Verstynen, T. D. (2016). Quantifying Differences and Similarities in Whole-Brain White Matter Architecture Using Local Connectome Fingerprints. *PLOS Computational Biology*, 12(11):e1005203.

[Yeh et al., 2010] Yeh, F. C., Wedeen, V. J., and Tseng, W. Y. I. (2010). Generalized q-sampling imaging. *IEEE transactions on medical imaging*, 29(9):1626–1635.

[Zeithamova et al., 2012] Zeithamova, D., Dominick, A. L., and Preston, A. R. (2012). Hippocampal and Ventral Medial Prefrontal Activation during Retrieval-Mediated Learning Supports Novel Inference. *Neuron*, 75(1):168–179.

[Zhang and Arfanakis, 2018] Zhang, S. and Arfanakis, K. (2018). Evaluation of standardized and study-specific diffusion tensor imaging templates of the adult human brain: Template characteristics, spatial normalization accuracy, and detection of small inter-group FA differences. *NeuroImage*, 172(July 2017):40–50.

[Zhu et al., 2014] Zhu, D., Zhang, T., Jiang, X., Hu, X., Chen, H., Yang, N., Lv, J., Han, J., Guo, L., and Liu, T. (2014). Fusing DTI and fMRI data: A survey of methods and applications. *NeuroImage*, 102:184–191.

[Zöller et al., 2019] Zöller, D., Sandini, C., Karahano?lu, F. I., Padula, M. C., Schaer, M., Eliez, S., and Van De Ville, D. (2019). Large-Scale Brain Network Dynamics Provide a Measure of Psychosis and Anxiety in 22q11.2 Deletion Syndrome. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, (October):881–892.

[Zoller et al., 2019] Zoller, D. M., Bolton, T. A. W., Karahanoglu, F. I., Eliez, S., Schaer, M., and Van De Ville, D. (2019). Robust Recovery of Temporal Overlap Between Network Activity Using Transient-Informed Spatio-Temporal Regression. *IEEE Transactions on Medical Imaging*, 38(1):291–302.

ANJALI TARUN

anjalibtarun@gmail.com, +41 78 676 53 99, Les Rives de la Morges 3, 1110, Morges, Switzerland
www.linkedin.com/in/anjali-tarun

I am an engineer working at the interface of signal processing, machine learning, and computational neuroscience with 4+ years of experience in medical data analysis. I am highly experienced in all stages of research and implementation of data science projects: from project design, and model development to interpretation and results dissemination through written publications and oral public presentations. I worked with various machine learning techniques, such as classification, regression, dimensionality reduction, graph learning, and other supervised and unsupervised learning approaches. I'm a goal-driven, flexible, and resourceful individual who maintains a positive, proactive attitude when faced with challenges and difficulty.

EDUCATION

2016 – 12.2020 (expected)

PhD in Electrical Engineering
École polytechnique fédérale de Lausanne (EPFL), Campus Biotech, Gèneve, Switzerland
Medical Image Processing Laboratory with Prof. Dimitri Van De Ville

Thesis: "Exploring dynamic functional connectivity by incorporating prior knowledge of brain structure"
Nominated for Thesis Distinction: Best Thesis Award in the Doctoral School

Courses: Signal Processing for functional brain imaging; Adaptation and Learning with Prof. Ali Sayed; Advanced Biomedical Imaging

2009 – 2016

Master of Science in Physics (First Class Honors)
National Institute of Physics, University of the Philippines Diliman

Bachelor of Science in Applied/Computational Physics (Second Class Honors)
National Institute of Physics, University of the Philippines Diliman
Academic Merit Award - Full Scholar, 100% Tuition and living allowances

WORK AND TEACHING EXPERIENCES

2016 – 2020

PhD Doctoral Assistant - Medical Image Processing Laboratory, EPFL
- Image and Signal Processing, Machine Learning, Data Science, Computational Neuroscience, Magnetic Resonance Imaging (MRI) analysis

Teaching Assistant, École polytechnique fédérale de Lausanne (EPFL)
- Image Processing I and II in ImageJ and Java, ~250 Master students
- Signal Processing for functional brain imaging in MATLAB, ~40 Master students

2012 – 2016

Researcher - **Instrumentation Physics Laboratory**, University of the Philippines Diliman
- Designed experiments to characterize the flow of macroscopic particles using MATLAB.
- Handled big datasets of (1) national exam scores of Philippine public schools and (2) natural disasters, to assess school hazard vulnerability and student learning

2014 – 2016

Full-time Faculty / Physics Instructor Level 3, University of the Philippines Diliman
- Fundamental physics courses: Classical mechanics, Electrodynamics, Thermodynamics and Optics, 60-120 student
- Designed experiments for physics laboratory classes ~ 60 Bachelor students

SKILLS AND QUALIFICATIONS

Languages	English (Fluent), Filipino (mother tongue), French (basic)
Computer Skills	MATLAB (7 years), Python (3 years), Java, Linux/Unix, bash, HTML, Visual Studio, Git, Sourcetree, Latex, MySQL, ImageJ, Adobe Illustrator, MS Office, R, Solidworks, Mathematica, neuroimaging platforms – SPM, FSL
Image and video processing	segmentation, 3D-surface reconstructions, 3D image analysis, particle tracking, multidimensional detection techniques, morphological operations
Machine Learning	Unsupervised learning: k-means clustering, Hidden Markov Models. Supervised learning: SVM, LDA, Naives Bayes, Decision trees, feature extraction and dimensionality reduction: PCA and ICA
Engineering Skills	graph signal processing, network theory, signal processing, image and video processing, neuroimaging, microprocessors, machine-language programming, data acquisition

HONORS AND AWARDS

Department of Science and Technology Merit Award
Awarded to students with outstanding academic achievements (Full Scholarship, tuition fees and allowances)

Outstanding Presentation Award
International Center for Theoretical Physics, Summer School for Complex Systems Research

Master of Science in Physics -- Summa Cum Laude – First Class Honors
Equivalent to awards given to students with the highest GPA (top 0.5%) in the number 1 university in the country

Bachelor of Science in Applied Physics -- Cum Laude – Second Class Honors
Awarded to students with distinction (top 3%) in the number 1 university in the country

LEADERSHIP AND STUDENT SUPERVISIONS

1. **Céline Provins**, *Master Thesis*, Master of Physics, EPFL
2. **Zhiwei Huang**, *Semester Project*, Master of Bioengineering, EPFL
3. **Ilaria Ricchi**, *Semester Project*, Master in Neuroengineering, EPFL
4. **Sergio Hernandez Charpak**, *Semester Project*, Master in Computational Science, EPFL
5. **Marta Boscaglia**, *Semester Project*, Master in Biomedical Engineering, EPFL
6. **Dionessa Biton**, *Bachelors Thesis*, BS Applied Physics, University of the Philippines

2019-2020 Ambassador for E3 – EPFL Excellence in Engineering program for AY 2019-2020: responsible for recruiting applicants for summer internships at the EPFL

06.2019 Workshop Organizer -- EPFL Network Science: Foundations and Perspectives, EPFL, co-organized with Daniela Zoller

LIST OF PUBLICATIONS

Publications In Revision/ Under Review / Preprint

1. **A. Tarun**, D. Abramian, H. Behjat, and D. Van De Ville, (2020) "Large – Scale Brain Graphs from Diffusion Weighted MRI: Spectral Analysis and Application to fMRI", Under Revisions for consideration in *IEEE Transactions on Medical Imaging (IEEE TMI)*
2. **A. Tarun** and D. Van De Ville, (2020) "Anatomical range of functional interactions meaningfully differentiates brain regions", *Preprint.*
3. **A. Tarun***, I. Ricchi*, H. Maretic, P. Frossard, and D. Van De Ville, (2020) "Dynamics of Functional Network Organization Revealed by Graph Laplacian Mixture Models". *Preprint.*
4. C. Sandini, D. Zoller, M. Schneider, **A. Tarun**, M. Armando, D. Van De Ville, S. Eliez, (2020) Characterization and Prediction of Clinical Pathways of Vulnerability to Psychosis through Graph Signal Processing. Revised and resubmitted to *e-Life.*
5. G. Bommarito, **A. Tarun**, et al., (2020), "Functional network dynamics in progressive multiple sclerosis". Submitted.
6. V. Siffredi, Y. Farouj, **A. Tarun**, et.al., (2020) "Large-scale brain network dynamics in callosal agenesis : subcortical contribution and laterality of networks", *Preprint.*

Published Manuscripts

1. **A. Tarun**, D. Wainstein, V. Sterpenich, L. Bayer, L. Perogamvros, N. Axmacher, S. Schwartz, D. Van De Ville (2020), "NREM sleep stages specifically alter dynamical integration of large-scale brain networks", *iScience* (in press).
2. **A. Tarun**, H. Behjat, D. Abramian and D. Van De Ville, (2020), "Structural mediation of human brain activity revealed by white matter interpolation of fMRI", *NeuroImage,* **213**, 116718
3. D. Biton, **A. Tarun**, R. Batac, (2020) "Comparing spatio-temporal networks of intermittent avalanche events: experiment, model, and empirical data", *Chaos, Solitons and Fractals* 130, 109519
4. TAW. Bolton, **A. Tarun**, V. Sterpenich, S. Schwartz, D. Van De Ville, (2018), "Interactions between large-scale functional brain networks are captured by sparse coupled HMMs", *IEEE Transactions on Medical Imaging*, **37**(1) 230-240
5. C. David, S. Monterola, E. Legara, **A. Tarun**, R. Batac, J. Osorio, (2018), "School hazard vulnerability and student learning", *International Journal of Educational Research*, **92** 20-29
6. R. Batac, A. Paguirigan Jr, **A. Tarun**, A. Longjas, (2017), "Sandpile-based model for capturing magnitude distributions and spatiotemporal clustering and separation in regional earthquakes", *Nonlinear Processes in Geophysics* **24** (2), 179-187
7. **A. Tarun,** A. Paguirigan, R. Batac, (2015) "Spatiotemporal recurrences of sandpile avalanches", *Physica A: Statistical Mechanics and Applications*, **436** 293-300

SCHOOLS, WORKSHOPS AND INVITED TALKS

Invited Talks

06.2020	Presentation to NeuroInformatics Group, University of Vienna, Austria
04.2018	Swiss Medical Image Computing Day 2018, University of Bern, Switzerland
06.2018	CIBM and Campus Biotech Day, University of Geneva, Geneva, Switzerland
03.2017	BrainHack, Zürich, Dynamic Functional Connectivity for Dummies, BrainHack Global, University of Zurich, Zurich, Switzerland

Workshops / Schools

10.2019	Fresh Ideas Talk, Sleep Modulations Technologies Conference, Ascona, Switzerland
04.2019	IEEE 16th International Symposium on Biomedical Imaging (ISBI), Venice, Italy
06.2019	Annual Meeting of the Organization for Human Brain Mapping (OHBM), Rome, Italy
06.2018	Graph Signal Processing Workshop 2018, EPFL, Lausanne, Switzerland
06.2018	Annual Meeting of the Organization for Human Brain Mapping (OHBM), Singapore, Singapore
11.2017	Computational Brain Connectivity Mapping, Institut National de Recherche en Informatique et en Automatique (inria), Juan-les-Pins, France
06.2017	Annual Meeting of the Organization for Human Brain Mapping (OHBM), Vancouver, Canada
06.2016	Complex Systems Summer School, Santa Fe Institute, Santa Fe, New Mexico, USA (highly competitive selection of students)
06.2015	Hands-on Research in Complex Systems Summer School, The Abdus Salam International Centre for Theoretical Physics, (ICTP), Trieste, Italy (Best Poster Award)
02.2015	NTU-University of Warwick Joint Winter School on Complexity, Nanyang Technological University, Singapore

FIRST AUTHOR CONFERENCES AND REFEREED PROCEEDINGS

1. **A. Tarun,** L. Perogamvros, D. Van De Ville, V. Sterpenich, S. Schwartz (2020). “Intrinsic organization of fMRI networks during sleep and its relevance to dreaming”. J. Sleep.Res. 29, 203-204
2. **A. Tarun**, D. Wainstein, V. Sterpenich, L., Bayer, L. Perogamvros, N. Axmachere, S. Schwartz, D. Van De Ville, “Temporally overlapping functional brain networks capture new features of information integration and reduced consciousness during sleep”, *OHBM 2020,* Montreal, Canada (organized online due to COVID-19)
3. I. Ricchi*, **A. Tarun***, H. Maretic, P. Frossard, and D. Van De Ville, “Inference of multiple functional brain networks using Graph Laplacian Mixture Model”, *OHBM 2020,* Montreal, Canada (organized online due to COVID-19)
4. **A. Tarun**, TAW. Bolton, V. Sterpenich, S. Schwartz, D. Van De Ville, “Dynamic interactions of spatiotemporally overlapping functional networks across sleep stages”, *OHBM 2019,* Rome, Italy
5. **A. Tarun**, D. Abramian, H. Behjat, and D. Van De Ville, “Graph spectral analysis of voxel-wise brain graphs from diffusion weighted MRI”, *IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019),* Venice, Italy
6. **A. Tarun**, and D. Van De Ville “Extrapolating functional MRI data into white matter via structurally-informed graph diffusion”, *OHBM 2018,* Singapore
7. **A. Tarun**, TAW. Bolton, V. Sterpenich, S. Schwartz, D. Van De Ville, “Spatiotemporal organization of intrinsic functional brain networks evaluated across different sleep stages”, *OHBM 2018,* Singapore
8. **A. Tarun**, Y. Farouj, I. Karahanoglu, V. Sterpenich, S. Schwartz, D. Van De Ville “Transient brain activity reveals spatiotemporal organization of functional networks during rest and sleep”, *OHBM 2017,* Vancouver, Canada
9. **A. Tarun**, J. Bantang, and R. Batac, “Analysis of avalanche dynamics and internal deformations using color-coded layers of grains”, *Annual Meeting of the Physical Society of the Republic of China 2016*, Kaohsiung, Taiwan

REFERENCE

Dimitri Van De Ville
Associate Professor, EPFL
H4 3 132.087 (Campus Biotech bâtiment H4)
Ch. des Mines 9, CH-1202 Genève
Contact: dimitri.vandeville@epfl.ch
+41 21 69 39669
http://miplab.epfl.ch/