

**INFERENCE OF STRUCTURAL BRAIN NETWORKS AND MODELING OF
CORTICAL MULTI-SENSORY INTEGRATION**

A Dissertation
Presented to
The Academic Faculty

By

Kamal Shadi

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Department of Computer Science

Georgia Institute of Technology

December 2019

Copyright © Kamal Shadi 2019

**INFERENCE OF STRUCTURAL BRAIN NETWORKS AND MODELING OF
CORTICAL MULTI-SENSORY INTEGRATION**

Committee members:

Prof. Constantine Dovrolis
School of Computer Science
Georgia Institute of Technology

Prof. Ellen W. Zegura
School of Computer Science
Georgia Institute of Technology

Prof. Shella Keilholz
Department of Biomedical Engineering
*Georgia Institute of Technology and
Emory University*

Prof. Zsolt Kira
School of Interactive Computing
Georgia Institute of Technology

Prof. Eva L. Dyer
Department of Biomedical Engineering
*Georgia Institute of Technology and
Emory University*

Date Approved: November 8, 2019

To my biggest support system and my number one fans: My family

ACKNOWLEDGEMENTS

Completion of this doctoral dissertation was possible with the support of several people. I would like to express my sincere gratitude to all of them. I am deeply grateful to Professor Constantine Dovrolis, my supervisor, who guided me and encouraged me to carry on through these years and has contributed to this thesis with a major impact. Thank you as well for guiding me, often with big doses of patience, through this journey.

I am also very grateful to my committee members, in particular professor Eva Dyer who had tremendous impact in shaping chapter 4 of my dissertation, and Professor Shella Keilholz who helped with the initial ideation of a section in chapter 3. I would also like to thank professor Ellen W. Zegura and professor Zsolt Kira for their valuable feedbacks and comments for the betterment of this dissertation.

I want to express my deep gratitude to my colleague, Manoj Chandrasekaran, who helped me with many of the experiments in chapter 3.

I want to thank many of my friends who their presence was very important in a process that is sometimes hindered by rough patches of life specially when you are an immigrant far from your family. Thank you my friends: Pouya Asrar, Neda Abolhassani, Amirhossein Salahshour, Amir Rostami, Amir Yazdanbakhsh, Payam Siyari, Kaeser Sabrin.

I owe a lot to my parents and family, who encouraged and helped me at every stage of my personal and academic life, and longed to see this achievement come true. My mother, Ezzat Jahandar, and My father, Hossein Shadi, I cannot find words to express my extremely deep feelings and gratitude for both of you.

TABLE OF CONTENTS

Acknowledgments	iv
List of Tables	ix
List of Figures	x
Chapter 1: Introduction and background	1
1.1 The problem of inferring human brain anatomical connectome	2
1.2 Case studies for inferring human brain anatomical connectivity: Major depression disorder network and cortical network	3
1.3 The dynamics of multi-sensory integration over anatomical connectivity	5
1.4 Thesis outline	6
Chapter 2: A symmetry-based method to infer structural brain networks from probabilistic tractography data	8
2.1 Introduction	8
2.2 Methods	11
2.2.1 Connectivity threshold τ	14
2.2.2 Network inference as an optimization problem	15
2.2.3 Threshold-based network inference with post-symmetrization	19
2.2.4 Performance metrics	19

2.2.5	Edge ranking and confidence metric in MANIA	20
2.2.6	Group analysis using MANIA	22
2.2.7	Synthetically generated networks	24
2.3	Results	27
2.3.1	Evaluation with synthetic data	27
2.4	Discussion	32
Chapter 3: Case Studies: Application of MANIA on experimental datasets . . .		36
3.1	Case Study A: corticolimbic circuitry implicated in major depression disorders	37
3.1.1	Data set	37
3.1.2	Results	39
3.1.3	Discussion	41
3.2	Case Study B: human corticocortical anatomical network	43
3.2.1	Data set	43
3.2.2	Node delineation: ROIs	43
3.2.3	Network inference without an arbitrary threshold or density: MANIA	45
3.2.4	Results	46
3.2.5	7T validation	47
3.2.6	Hemispheric asymmetries	55
3.2.7	Discussion on HCP study - case study B	57
Chapter 4: Multi-sensory integration in the mouse cortical connectome using a network diffusion model		60
4.1	Introduction	60

4.2	Methods and Data	64
4.2.1	Structural network	64
4.2.2	The ALT model and activation cascades	67
4.2.3	Analysis of activation cascades	69
4.2.4	Comparison of modeling results with functional data	70
4.3	Results	72
4.3.1	Sensory-specific activation cascades in the mouse cortex	72
4.3.2	Model validation	76
4.3.3	Similarity of sensory cascades	78
4.3.4	Core ROIs and hourglass architecture	79
4.3.5	Location of τ -core nodes in activation cascades	82
4.3.6	Robustness of τ -core nodes	85
4.3.7	Which anatomical connections are more important in sensory cascades?	88
4.3.8	Comparison with static network analysis methods	89
4.4	Discussion	90
	Chapter 5: Closing remarks	95
5.1	Conclusions	95
5.2	Limitations	98
5.2.1	Lack of ground truth to quantify the accuracy of tractography methods	98
5.2.2	Undirected nature of diffusion data	98
5.2.3	Partial validation of the ALT model	99
5.3	Extensions	99

Appendices	102
Appendix A: List of HCP subjects used in chapter 3	103
Appendix B: Supplementary Information for chapter 4	104
B-1 Sources of sensory cascades	104
B-2 Connectome filtering	104
B-3 The activation cascade of each sensory source	104
B-4 Analysis of “disagreement cases” between VSD data and ALT modeling results	105
B-5 Similarity between activation cascades with single and complete linkage	107
B-6 Activation cascades when two sensory sources are activated simultaneously	107
B-7 Comparison with other network analysis metrics	107
References	138

LIST OF TABLES

3.1	The 18 corticolimbic ROIs we consider in the case-study.	42
3.2	Top-three nodes in rank-aggregated network based on four node-centrality metrics	43
3.3	Connections with hemispheric lateralization. The first two connection are mostly present in left hemisphere and the rest of the connection are mostly present in right hemisphere while absent in the opposite hemisphere	56
3.4	Meta information about ROIs in table 3.3	56
5.1	Summary of the contributions in this thesis by chapter and type - group 1 is related to the anatomical analysis while group 2 is related to the MSI dynamic	97
A.1	List of the subjects used in our study	103
B-1	<i>Sources of sensory cascades</i>	104
B-2	Top-five nodes according to different centrality metrics	108
B-3	Rich-club nodes, Rombach core nodes nodes, and hourglass core nodes ($\tau=90\%$): the overlap between the first two sets with the hourglass core nodes is highlighted in red.	109

LIST OF FIGURES

2.1	Running tractography with streamlines from a seed voxel q in the i th ROI to the j th target ROI. Two streamlines hit the target ROI, therefore $T_i(q, j) = \frac{2}{3}$.	13
2.2	Top: Network density ρ as a function of the connectivity threshold τ (plotted for one subject in our DTI dataset). Bottom: Network asymmetry (red) and normalized network asymmetry (blue) as functions of the network density ρ . The optimal density ρ^* is the largest value that minimizes the normalized network asymmetry.	17
2.3	As we decrease the connectivity threshold, each edge first appears at a certain value of the network density. If this density is larger than ρ^* , the corresponding edge is <i>not</i> present in the MANIA network.	21
2.4	Probabilistic error model (Z) for tractography-generated connection probabilities using the maximum entropy distribution (with one degree of freedom).	26
2.5	The Jaccard similarity difference ($\Delta J = J_{SYM} - J_{NO-SYM}$) with and without post-symmetrization for five threshold values and for MANIA. Each box plot is generated from 1000 experiments; in each experiment we generate a random network with density between 0 and 1, while the noise parameters μ_1 and μ_2 are uniformly distributed between 0 and 0.3. The red line corresponds to the median, the box boundaries to the 25 th and 75 th percentiles, while the dashed lines show the 10 th and 90 th percentiles. In all cases, $\Delta J \lessdot 0$ (one-sided Mann-Whitney U test p-values are shown next to each box plot) meaning that <i>post-symmetrization helps to improve the accuracy of network inference</i> .	28
2.6	False positive rate and false negative rate of MANIA as a function of μ_1 and μ_2 for sparse networks ($\rho_G = 0.1$), medium density networks ($\rho_G = 0.5$) and dense networks ($\rho_G = 0.9$). Each square is the average of 1000 independent simulations.	29

2.7	The Jaccard similarity difference ($\Delta J = J_{\tau_{opt}} - J_{MANIA}$) between MANIA and the optimal threshold-based scheme as a function of μ_1 and μ_2 for sparse networks ($\rho_G = 0.1$), medium density networks ($\rho_G = 0.5$) and dense networks ($\rho_G = 0.9$). Each square is the average of 1000 independent simulations.	30
2.8	The Jaccard similarity difference ($\Delta J = J_{MANIA} - J_{\tau_0}$) between MANIA and five given threshold values. The accuracy comparisons are made after post-symmetrization. Each box plot is generated from 1000 experiments; in each experiment we generate a random network with density between 0 and 1, while the noise parameters μ_1 and μ_2 are uniformly distributed between 0 and 0.3. The red line corresponds to the median, the box boundaries to the 25 th and 75 th percentiles, while the dashed lines show the 10 th and 90 th percentiles. In all cases, $\Delta J \leq 0$ (one-sided Mann-Whitney U test p-values are shown next to each box plot) meaning that <i>MANIA is more accurate than inferring the network based on a fixed threshold.</i>	31
3.1	The rank-aggregated network, based on DTI data from 28 subjects, between the 18 ROIs in Table 3.1. Every edge in the connected component has been detected in both directions, i.e., MANIA identifies a completely symmetric network (no post-symmetrization is needed). The density of the connected component is 19%. The color of each edge represents the fraction of subjects that have the corresponding edge in their individual MANIA-based networks.	39
3.2	a) 180 ROIs from HCP's multi-modal cortical parcellation (HCP-MMP) projected on the S1200 group average cortical surface (inflated left hemisphere). This visualization is generated by the data from [23] b) The size (surface area) distribution for the 180 ROIs	44
3.3	The results of applying MANIA on three sample subjects for right and left hemisphere network inference - the dashed lines show the densities chosen by MANIA for the corresponding networks. The density is calculated excluding the adjacent pairs.	48
3.4	The results of applying MANIA on three sample subjects for inter-hemispheric network inference - the dashed lines show the densities chosen by MANIA for the corresponding networks	49
3.5	The distributions of MANIA results across 37 subjects - The distributions are calculated by kernel-density estimate using Gaussian kernels [94] . . .	50

3.6	Connection probability vs axonal distance for subject 146735 in connected and none-adjacent ROIs - The graphs are presented for left, right and inter-hemispheric networks separately and MANIA thresholds are shown in the graphs. In the left hemisphere 12% of the connections are non-adjacent and 9% of connections are none-adjacent in the right hemisphere. Note that the probtrackx was run to include only streamlines extending at least 20mm (distthresh=20).	50
3.7	Sample source to target connection probability (<i>log – scale</i>) vs distance for subject 14673 both in 3T and 7T - a) An ROI pair in which 7T signal is attenuated compared to 3T b) An ROI pair with the opposite effect (amplification in 7T compared to 3T).	52
3.8	True positive and true negative sizes in subject 146735 of our benchmark dataset built using 7T data for different δ_1 and δ_2 parameters	52
3.9	Measuring the accuracy of MANIA with $\delta_1 = 0.4$, $\delta_2 = 0.01$ Top) False positive rate and true positive rate across different density thresholds - MANIA operating points are marked by red lines Bottom) F-measure and ROC for different density thresholds - MANIA operating points are marked by red line and red star	54
3.10	The F-measure across several schemes to build the benchmark dataset out of 7T dataset - the study is done for three subjects marked in the title of each figure.	55

4.1	<i>High-level illustration of approach and main result:</i>	Top-Left: An illustration of a single sensory cascade, originating at the primary visual cortex (VISp). The yellow edges show anatomical connections that participate in the cascade. The ROIs at the upper layer reside on the cortical surface while the ROIs at the lower layer are deeper in the brain. The edges between the two layers are dashed. To simplify the visualization, we only include 20 ROIs at this cascade (the complete cascade includes 67 ROIs). The cascade forms a Directed Acyclic Graph (DAG), and it is produced by applying the Asynchronous Linear Threshold (ALT) model on the cerebral cortex portion of the mouse connectome – starting the cascade at VISp. Bottom and Right: The two most important core ROIs at the hourglass waist – Claustrum (CLA) and Posterior Parietal (PTLp) cortex – jointly cover 40% of all activation paths in the ten sensory cascades we consider. Each cascade is represented by a different color. We include two circular disks (pie charts), one for each of these two core ROIs. The proportion of activation paths in each cascade is shown by the corresponding color. The protruded portion of each circular section represents those paths that traverse the corresponding core ROI. For example, the activation paths in the auditory cascade account for about 8.5% of the total number of paths – and about 36% of those paths traverse CLA.	61
4.2	<i>The location of the ten primary sensory regions:</i>	Three coronal slices from the Allen Mouse Brain Atlas with the cerebral cortex regions tinted by green and the source regions identified by an arrow. The somatosensory region includes six different sub-regions for lower limbs, upper limbs, trunk, mouth, nose and whiskers. The remaining four sensory sources are: visual (VISp), auditory (AUDp), gustatory (GU) and olfactory (MOB).	66
4.3	<i>Illustration of ALT model and τ-core analysis:</i>	(a) A toy example of a 5-node network on which we run the ALT model. Each edge is marked with a communication delay, followed by a weight. The activation threshold is $\theta=1$. The black edges represent the underlying structural network while the red unidirectional edges represent the activation cascade as it unfolds over time. (b) The activation cascade directed acyclic graph for the previous toy example. The source of the cascade is n_1 . (c) A toy example with three activation cascades (the sources are nodes u , v and y). The total number of source-target paths is 11 (4 at the left, 3 at the middle, and 4 at the right). Node w has the highest path centrality ($P(w)=9/11$). If $\tau \leq 9/11$, the τ -core consists of only that node.	68

4.4	<i>VSD data processing pipeline:</i> (a) <i>Lower:</i> The Allen Reference Atlas (ARA). <i>Upper-left:</i> A sample VSD image covering most of the left cortical surface five frames after visual stimulation. <i>Upper-right:</i> The ROIs at the left ARA cortical surface mapped to the native cortical surface of an animal. (b) The activation time of a pixel is defined as the frame of maximum post-stimulus VSD signal at that pixel. (c) The activation time of an ROI is defined as the activation time of <i>most</i> pixels in that ROI. (d) The output of this pipeline is an activation time for each ROI, depicted here with a grey-scale map (black for the first ROI activation and white for the last).	71
4.5	<i>Effect of parameter θ on cascade size, and similarity between the ten cascades:</i> (a) Each row of the heat map shows the fraction of activated nodes after the stimulation of a single source, for different values of the threshold θ . The selected threshold is marked with the dashed vertical line. (b) Similarity between the ten sensory cascades using the average-linkage hierarchical clustering method.	73
4.6	<i>The visual activation cascade, accoring to ALT:</i> The source for this cascade is the primary visual cortex (VISp). The red edges form the activation cascade, while the underlying grey edges show anatomical connections that do not participate in this cascade – those connections may be present in other sensory cascades or they may play a role in feedback (or second-order) interactions that are not captured by the “first ripple” scope of the ALT model. To help with the visualization, we place the nodes in eight layers, so that cascade edges only point from a layer to a higher layer (never to the same or lower layer). The vertical position of each node is slightly “jittered” to avoid cluttering due to anatomical connections between nodes of the same layer.	75
4.7	<i>Comparison between model-based and experimental temporal ordering of ROI activations:</i> a) The y-axis shows the percentage of (X, Y) ROI pairs that show temporal agreement (green), temporal disagreement (red), and insufficient temporal resolution (blue) between the activation order of X and Y in the modeling results and the mouse experiments. The plot shows results for five animals and for five sensory stimulations (a touch at the whiskers, forelimb, and hindlimb, as well as an auditory and a visual stimulation). b) The same comparison but here we have randomized the ROIs that are active during each frame, preserving the number of ROI activations in each frame.	77
4.8	<i>Path centrality and τ-core analysis:</i> (a) Path Centrality (PC) histogram for the 67 regions in N_c , considering all source-target paths across the ten activation cascades. (b) Cumulative path coverage by the top- X core nodes for $X=1 \cdots 67$. Nine regions are sufficient to cover $\tau = 90\%$ of all paths. . .	81

4.9	<i>Illustration of hourglass architecture:</i> A hypothetical network with feedforward, feedback and lateral connections between regions at different levels of the cortical hierarchy. Input information is provided at sensory-specific modules (left), while high-level cognitive tasks are performed by association regions at the other end of the hierarchy (right). The “hourglass feature refers to the fact that the high-dimensional input information is first integrated through through a relatively small number of highly central intermediate-level regions, before it is re-used at high-level cortical regions and tasks.	82
4.10	<i>Location-related metrics:</i> In both matrices, a column represents one of the ten activation cascades, originating at the node shown at the top of the column. (a) Each row represents the <i>source-distance</i> of the corresponding node from the source of that column’s cascade. White denotes a distance of one hop, while black denotes the maximum distance for that cascade. Rows are ordered in terms of the average distance (in number of hops) of the corresponding node from the sources of all activation cascades (excluding the MOB cascade, which is very different). (b) Each row represents the <i>influence</i> of the corresponding node, i.e., the number of nodes that are reachable from that node in the activation cascade that the column represents. White denotes an influence of one (only that node), while black denotes an influence that covers all network nodes. Rows are ordered in terms of the average influence of the corresponding node across all activation cascades (excluding the MOB cascade).	83
4.11	<i>Robustness results:</i> The effect of different connectome randomization methods on the core size. Light blue shade marks the 5% to 95% values among 100 randomization runs, while the solid blue line is the median of these runs. The red line represents the τ -core size for the original connectome. The dotted green line marks the τ -core size for $\tau=90\%$. The table at the bottom shows the fraction of random networks that include each of the eight τ -core nodes.	87
4.12	<i>Connections that appear in sensory cascades:</i> (a) Conditional probability that a connection e_l of physical length l appears in an activation cascade given that $l > l_0$. (b) Conditional probability that a connection e_w of weight w appears in an activation cascade given that $w > w_0$	88
B-1	<i>Connectome filtering. Left:</i> Network density versus edge p-value. Right: Connection density (edge weight) versus edge p-value	105
B-2	<i>Visual cascade</i> (source: VISp)	110
B-3	<i>Auditory cascade</i> (source: AUDp)	111

B-4	<i>Gustatory cascade</i> (source: GU)	112
B-5	<i>Upper-limb somatosensory cascade</i> (source: SSp-ul)	113
B-6	<i>Lower-limb somatosensory cascade</i> (source: SSp-ll)	114
B-7	<i>Whiskers somatosensory cascade</i> (source: SSp-bfd)	115
B-8	<i>Trunk somatosensory cascade</i> (source: SSp-tr)	116
B-9	<i>Mouth somatosensory cascade</i> (source: SSp-m)	117
B-10	<i>Nose somatosensory cascade</i> (source:SSp-n)	118
B-11	<i>Olfactory cascade</i> (source: MOB)	119
B-12	(a)-(e): The number of disagreements between VSD and ALT that involve each ROI, over the 21 ROIs that appear in the VSD data. The outlier ROIs have frequencies that exceed the red dashed line. (f) Disagreement ROI pairs that appear in all five animal datasets - each row corresponds to a different stimulus. ROIs at the boundary of the VSD visible cortical surface are marked by a star.	120
B-13	<i>Similarity between activation cascades.</i>	121
B-14	<i>Simultaneous activation of two source nodes:</i> (a) Path Centrality (PC) histogram for the 67 cortical nodes, considering all source-target paths across the $10 \times 9/2 = 45$ activation cascades. (b) Cumulative path coverage by the top- X core nodes for $X=1 \dots 67$. (c) Eight nodes are enough to cover $\tau = 90\%$ of all paths.	121
B-15	<i>Path centrality compared to other centrality metrics.</i> Each point corresponds to one of the 67 nodes in the network. The four star nodes constitute the hourglass core for $\tau=70\%$. All nodes are color-coded based on the broader brain region they belong to (isocortex, hippocampal formation, cortical subplate, olfactory areas).	122

Summary

Recent advances in neuroimaging have enabled major progress in the field of brain connectomics, i.e., constructing maps of connections between brain regions at different scales. Diffusion MRI (dMRI) and probabilistic tractography algorithms are state of the art methods to map the structural connectome of the brain non-invasively and in vivo. Although probabilistic tractography can detect many major connections in the brain, it also reports some spurious ones. We propose and evaluate a method, referred to as MANIA (Minimum Asymmetry Network Inference Algorithm) that can infer the structural brain network that interconnects a given set of Regions of Interest (ROIs) from probabilistic tractography data in a threshold-free manner. Given that diffusion MRI is unable to detect the direction of each connection, we formulate the network inference process as an optimization problem that minimizes the (appropriately normalized) asymmetry of the observed network.

The most fundamental property of the human connectome, its density, is still elusive and debated. MANIA is well-positioned to address this open question because it does not depend on an arbitrary weight threshold. We use MANIA to infer the human cortico-cortical connectome from the data published by Human Connectome Project (HCP). MANIA reports connectomes that are highly consistent across individuals at a density of approximately 3.2%. We validate the accuracy of these connectomes by comparing the connections inferred using MANIA at 3T MRI acquisitions with 7T high-resolution MRI acquisitions of the same subjects.

Having a structural network is instrumental in analyzing communication dynamics and information processing in the brain. The last research problem, we focus on relates to multi-sensory integration in the cortex. We model this process on the mouse cortical connectome (provided by the Allen Institute) by employing an Asynchronous Linear Threshold (ALT) diffusion model on that connectome. The ALT model captures how evoked activity that originates at a primary sensory region of the cortex ripples through other cortical regions. We validate the ALT model using Voltage Sensitive Dye (VSD) imaging data. Our results

show that a small number of cortical regions (including the Claustrum) integrate almost all sensory information streams, suggesting that the cortex uses an hourglass architecture to integrate and compress multi-sensory information.

CHAPTER 1

INTRODUCTION AND BACKGROUND

Complex systems comprised of a large number of components have intricate inter-connected structural and functional couplings [1]. A new approach to analyze these de facto “giant graphs” has emerged and evolved over the last two decades under the unifying umbrella of network science [2]. Network science provides powerful tools for use across different disciplines, from biology [3] to social networks [4] to the study of climate [5], etc. One of the most important frontiers (if not the most important) in which network scientists are actively pushing the envelopes is neuroscience, whose ultimate goal is to understand the human brain at both the structural and functional level [6].

The study of the brain as a complex network is called connectomics [7], and involves the mapping of neuronal inter-connections at different spatial and temporal scales and across multiple modalities and species. It is now established in the field of neuroscience that studying the brain by isolating and examining each of its individual parts in isolation has served its purpose but cannot explain the mechanism by which the brain receives, stores, retrieves, processes and integrates information [8]. Connectomics is believed to be the path forward in understanding the brain, given that the interplay of anatomical underpinning and neuronal dynamical processes is taken into consideration [9].

Analyzing the brain by focusing on its connectome (i.e., connectomics) has already proven to be a pivotal force in our understanding of various mental disorders [10]. Most of the implications of mental disorders such as the Alzheimers disease, schizophrenia, depression, and so on, are network-wide and not localized. Aside from brain diseases, the importance of the connectome has also been shown in other cognitive and psychological domains. For instance, intelligence, personality and various other psychometrics of a person are correlated with his or her connectome [11]. What is more, understanding the

brain by focusing on its connectome is also invaluable for those working on technological advances, since obtaining and analyzing intelligent connectomes, particularly those of the human brain, allows for the best blueprint and roadmap on our path to better AI and computer architecture.

1.1 The problem of inferring human brain anatomical connectome

Connectome maps can be generated in different species and at varying resolutions, from the neuronal level (microscale) to the large, population-level macroscale (approx. 1mm cubes and larger). While microscale analysis provides details of connections between all neurons, macroscale network neuroscience conceptualizes brain function as emerging from the collective action of neuronal populations, each clustered together to form homogeneous functional modules. These spatially contiguous and functionally homogeneous modules are known as Regions of Interest (ROI). The micro- and macroscale analyses both provide valuable and sometimes correlated insights [12].

To this day, *C. elegans* is the only organism that has had its complete connectome mapped [13]. *C. elegans* nervous system comprises only 302 neurons and 7,000 connections, compared to the 86 billion neurons and trillions of connections in the human brain. This maybe paints a pessimistic view of the current state of the technology when mapping more complicated brains, but the technology is developing quickly, and whole-brain connectomes are becoming plausible at the macroscale across a number of different species [14].

One of the most rapidly advancing technologies serving connectomics is Magnetic Resonance Imaging (MRI). The last two decades have seen tremendous advances in MRI which now offers the most promising technology for characterizing the connectivity of human brain non-invasively and at the macroscale. MRI technology has been used widely in both functional and structural connectome projects, such as the Human Connectome Project (HCP) [15] in the USA, and the Human Brain Project (HBP) in the European

Union [16]. Diffusion MRI (dMRI) and tractography are often bundled together as the preferred pipeline for inferring the brain structural wiring (connectome) at the macroscale. Starting from the observational data obtained using dMRI, tractography is employed to quantify inter-ROI connection probabilities and, sometimes, their strengths.

In this thesis, we propose and describe a new approach, dubbed the Minimum Asymmetry Network Inference Algorithm (MANIA), to measure connectivity threshold. It is extracted from probabilistic tractography algorithm. The main tenet is based on the fact that tractography does not provide any information about the polarity of fiber tracts. As such, we developed a method that aims to minimize the apparent (artifactual) asymmetry observed in structural connectivity matrices.

The proposed method leverages the symmetric nature of dMRI to improve the accuracy in determining an adjacency matrix from a set of probabilistic tractography results. Leveraging our method eliminates the need for subjective thresholding in the inferring of brain connectomes. Having subjective thresholds in network inference to remove spurious edges results in networks with differing densities or differing number spurious edges in the brain graphs of different subjects. This, consequently, confounds and biases the cross-comparisons between both individual connectomes and connectomes across studies. MANIA solves this issue by finding a threshold tuned for optimally of network inference for each individual dMRI dataset.

1.2 Case studies for inferring human brain anatomical connectivity: Major depression disorder network and cortical network

In this thesis, we apply MANIA to infer anatomical brain networks in two case studies. First, we apply MANIA to Diffusion Tensor Imaging (DTI) data collected in an earlier study: “Brain aging in humans, chimpanzees (*Pan troglodytes*), and rhesus macaques (*Macaca mulatta*): magnetic resonance imaging studies of macro- and microstructural changes” [17]. Twenty-eight healthy, right-handed females between the ages of 18 and

22 were selected from this study to form a network of 18 ROIs that are known to play a significant and important role in a number of mental disorders. These 18 ROIs were curated from multiple studies [18, 19, 20, 21].

Interestingly, inferred MANIA network showed that the connection between the sub-cortical limbic system and the frontal lobe of the cortex (the connection between BA25 and the Nucleus Accumbens) is a major anatomical pathway. The same connection pathway is used in deep brain stimulation (DBS) of patients with major depression disorder. In DBS, electrodes are used to deliver a small amount of electrical current to this region of the brain, which is overactive in people with depression. The region serves as a junction box, so adjusting its activity affects the entire depression circuitry.

We also applied MANIA to a more recent dataset from HCP [15]. We used 37 female subjects, aged between 31-35 years old, to infer cortico-cortical networks ¹. In doing so, instead of relying on more crude, and classical parcellations, e.g. that of [22], we opted to use multi-modal and high resolution parcellation, details of which was recently published by authors in [23] and which is known as **HCP-MMP** parcellation. Thus, our published cortical connectome is one of the first connectomes available on HCP data with HCP-MMP resolution. Leveraging these 37 cortical connectomes, we show that the human cortical network, as inferred by probabilistic tractography, has a density of approximately 3.2%.

The previous studies were conducted on data acquired by 3T MRI scanners. The HCP announced in late 2016 that a first release of MRI data from 7T scanners is available. We used three sample subject of this release to build a cortical network in 7T, and compared the network to that derived using the 3T technology (on the same subjects). We had multiple observations from this case study. Notably, we showed that using the connections and absence of connections, implied strongly by 7T data, as a proxy for ground-truth, MANIA accuracy is close to that of an optimum estimator which tries to maximize the F-measure.

¹We use cortico-cortical network and cortical network interchangeably in this thesis

1.3 The dynamics of multi-sensory integration over anatomical connectivity

The study of the brain neither ends, nor is complete at the structural level. While the anatomy and the physical wiring of the brain is the site in which functional behaviour emerges, the connectome by itself however is not sufficient to analyze communication dynamics and information processing. Thus, new approaches are needed that combine structural connectivity with models of network dynamics, such as diffusion and communication.

Understanding the path from structural constraints to functional behaviour is a new trend that needs better computational models [9] at this stage. The link from a structural to functional connectome is not direct, and goes through the neuronal dynamics at both the micro- and macroscale level [24]. For instance considering global communication dynamics, it has been shown that there are hubs with many anatomical connections that exhibit a topologically central position in the overall network [25]. Demystifying the communication dynamics in the brain is pivotal to understanding how multiple sensory streams from the environment are merged in the brain of animals for a unified perception of the world around them.

In this thesis, we make steps towards understanding multi-sensory information processing using a network diffusion modeling approach. To make these steps, we first searched for a suitable anatomical brain model. Unfortunately, the human brain networks inferred by DTI and tractography are not directed. Thus, we focused our search on the brains of other mammals. In parallel to advances in MRI used in macroscale analysis, technologies for mapping brain networks at the mesoscale (i.e. between macro and micro) and the microscale have advanced (and continue to advance) in non-human mammals. These technologies are invasive and not suitable for study of the human brain; thus they are actively used to study the brains of other animals. Tracer studies are among the most widely adopted; tracer chemicals are injected in the brain and then tracked using microscopy. They also provide directional connectome, which are necessary to study communication dynam-

ics. Of the connectomes available using tracer studies, we opted to use the connectome by [26]. We did so because of its public availability, resolution and the reputation for accuracy it has within the community.

We use the Asynchronous Linear Threshold (ALT) diffusion model to trace activity propagation across a network on the cortical mouse connectome of [26]. We do so to uncover the topological features facilitating Multi-Sensory Integration (MSI). In the simulation, primary sensory areas are stimulated and the activation flows of all subsequent nodes (all nodes were participating in all the cascades) were used to construct a Directed Acyclic Graph (DAG), referred by Activation DAG (ADAG). In order to validate the ALT model, we use Voltage Sensitive Dye (VSD) imaging data from [27]. VSD data were collected from mice while five different sensory stimuli were introduced to evoke responses in the visual, somatosensory (upper limb, lower limb and whisker), and auditory cortices. This analysis corroborated the predictive power of ALT in modeling the diffusion of cortical activity from specific source nodes.

The main result of this analysis is that there are a small number of regions with high centrality forming a bottleneck in MSI dynamic. Thus, our study shows that the mouse cortical network has an hourglass architecture as the main topological signature by which the sensory streams are integrated at early stages to form a unified perception. These core nodes include key association areas, such as the claustrum and posterior parietal associative areas, suggesting that these areas play critical roles in MSI.

1.4 Thesis outline

Our methodological contributions and findings are discussed in the thesis as follows. In chapter 2, we present a new technique (published in [28]) to convert data derived from dMRI and tractography to actual networks. The method, referred to as MANIA (Minimum Asymmetry Network Inference Algorithm), is capable of inferring the brain connectome without any parameter involved in the inference pipeline, leading to reproducible and ob-

jective networks. In chapter 3, we apply our method to two experimental datasets. First, we reconstruct the network between 18 ROIs playing significant roles in several mental disorders. Then, we apply MANIA to 37 female subjects from HCP. We further validate our inference method using high resolution MRI data available by HCP. Our results re-iterated the effectiveness of the method using more recent MRI acquisition technologies. Chapter 4 departs from structural network inference to its application on communication dynamics. Specifically, we investigate how multiple sensory information streams travel on structural underpinning. By employing Asynchronous Linear Threshold (ALT) diffusion model on the cortical structural of mice, we present evidence of hourglass-like topological attributes in multi-sensory integration in the mouse cortex. Finally, we conclude this thesis in chapter 5.

CHAPTER 2

A SYMMETRY-BASED METHOD TO INFER STRUCTURAL BRAIN NETWORKS FROM PROBABILISTIC TRACTOGRAPHY DATA

Recent progress in diffusion MRI and tractography algorithms as well as the launch of the Human Connectome Project (HCP)¹ have provided brain research with an abundance of structural connectivity data. In this work, we describe and evaluate a method that can infer the structural brain network that interconnects a given set of Regions of Interest (ROIs) from probabilistic tractography data. The proposed method, referred to as *Minimum Asymmetry Network Inference Algorithm (MANIA)*, does not determine the connectivity between two ROIs based on an arbitrary connectivity threshold. Instead, we exploit a basic limitation of the tractography process: the observed streamlines from a source to a target do not provide any information about the polarity of the underlying white matter, and so if there are some fibers connecting two voxels (or two ROIs) X and Y, tractography should be able in principle to follow this connection in both directions, from X to Y and from Y to X. We leverage this limitation to formulate the network inference process as an optimization problem that minimizes the (appropriately normalized) asymmetry of the observed network. We evaluate the proposed method on a noise model that randomly corrupts the observed connectivity of synthetic networks.

2.1 Introduction

Diffusion MRI has opened a new window at the meso-scale structure of the living brain [29]. Clinicians and researchers can now observe and measure the properties of white matter in a non-invasive manner, analyzing the location and density of neuronal fibers at a spatial granularity of 1-2mm isotropic voxels [30]. Such structural information is important

¹www.humanconnectome.org/

in deciphering how the brain works [31, 32], and it also creates new ways to understand and potentially diagnose [33, 34] or even treat [35] various brain diseases [36, 37].

Processing diffusion MRI data using tractography algorithms is the next step forward: instead of analyzing the properties of white matter at the level of individual voxels, tractography aims to detect individual bundles of neuronal fibers originating or passing through a given “seed” voxel [38]. Additionally, given a seed voxel and a target ROI, it is now possible to examine the likelihood that some white matter fibers connect the two (referred to as “probabilistic tractography”), and to track the shape of these connections [39]. In this chapter, we propose a method to further process the noisy connectivity information provided by probabilistic tractography in order to estimate an interconnection network between a given set of grey matter ROIs.

Diffusion MRI data, jointly with deterministic [40] or probabilistic tractography methods [39] have been used successfully during the last decade to infer the structure of the human brain between hundreds of ROIs [41]. Various structural properties of these networks have been discovered for the healthy brain and for various psychiatric diseases [42, 43]. When combined with fMRI and behavioral or genomic analysis, these non-trivial topological properties provide new insights about the role of individual ROIs in specific networks and the way in which these distinct ROIs exchange information to produce integrated function [44].

A major challenge in this research effort is that the inferred brain networks, as well as their topological properties, are often sensitive to the parameters of the tractography process [45]. In probabilistic tractography, the most critical of those parameters is the *connectivity threshold* τ that determines whether the tractography-generated streamlines from a given seed voxel to a target ROI occur with sufficiently large probability to indicate the presence of an actual connection [46]. If τ is too low the resulting network includes connections that do not exist in reality and the converse happens if τ is too high. Even a small number of spurious or miss-detected edges can adversely effect the properties of the

inferred networks [47]. Further, the optimal value of τ , i.e., the threshold that would result in the most accurate reconstruction of the underlying “ground truth” network, may vary between different subjects [48] and image acquisitions parameters [49].

The problem of selecting an appropriate connectivity threshold in either deterministic or probabilistic tractography is not new. One approach has been to select the largest possible threshold (i.e., fewest possible edges) so that the final inferred network remains connected across the majority of subjects [50]. Depending on the selected ROIs, this approach can lead to many miss-detections (if those ROIs are densely interconnected) or false alarms (if the ROIs are not directly connected). Another approach has been to analyze the tractography data with a wide range of threshold values, hoping that certain qualitative properties are robust and independent of the exact threshold. Li et al. investigated how the connectivity thresholds affects network density and therefore network efficiency metrics [46]. Duda et al. have shown the importance of the connectivity threshold for various network metrics such as clustering coefficient and characteristic path length [45].

This work focuses on the following problem: *how to infer the structural network between a given set of grey matter ROIs in a reliable way that does not require an arbitrary choice of the connectivity threshold?* The proposed method, referred to as *Minimum Asymmetry Network Inference Algorithm (MANIA)*, exploits a fundamental limitation of diffusion MRI imaging and of the tractography process: diffusion MRI can estimate the orientation of fibers in each voxel but it cannot infer the polarity (afferent versus efferent) of those fibers [51, 52]. Similarly, a tractography algorithm can combine those per-voxel orientations to “stitch together” expected connections but it does not provide any information about the direction of those connections [53]. Given this limitation, MANIA expects that the presence of an actual connection from voxel X to voxel Y (in that direction) will be detected by the tractography process as a symmetric connection between X and Y. Similarly, if there is no connection between X and Y, the tractography process should not detect a connection in either direction. Based on this principle, MANIA formulates the network

inference problem as an optimization over the range of connectivity threshold values: it selects the value of τ that minimizes the asymmetry of the resulting network. The network asymmetry is normalized relative to the asymmetry that would be expected due to chance alone in a random network of the same density.

MANIA can work in tandem with all probabilistic tractography methods, such as FSLs probtrackx [39], PiCo [54], and fDF-PROBA [55]. It can be also combined with deterministic tractography methods, such as FACT [40], but only if a large number of streamlines (in the thousands) are generated from randomly placed seeds within each voxel.

We expect that the given set of ROIs primarily reside in grey matter. Dilating a grey matter ROI so that it includes some white matter voxels may result in connectivity errors, especially with cortical ROIs, because of the dense white matter systems just beneath the cortical sheet [56]. The selection of ROIs and the estimation of their boundaries is an important issue that is further discussed in Section 4

We evaluate the accuracy of MANIA based on synthetically generated data in which the ground-truth network is known. We also compare MANIA with an ideal threshold-based method in which the optimal connectivity threshold is assumed to be known. Further, we show how to associate a confidence level with each edge, and how to apply MANIA in a group of subjects. In next chapter, as a case-study, we apply MANIA on diffusion MRI data from two different experimental datasets.

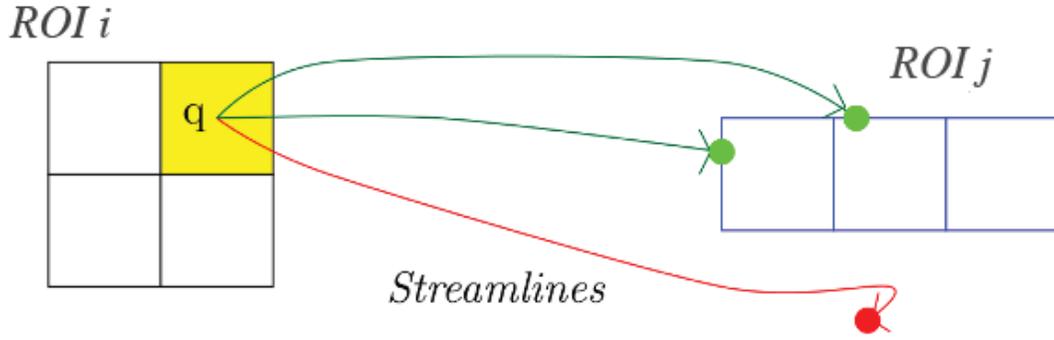
2.2 Methods

We apply MANIA in Diffusion Tensor Imaging (DTI) data collected by an earlier study: “*Brain aging in humans, chimpanzees (*Pan troglodytes*), and rhesus macaques (*Macaca mulatta*): magnetic resonance imaging studies of macro- and microstructural changes*” [17]. Twenty eight healthy right-handed females between the ages of 18 and 22 (mean: 20.2), without a history of psychiatric disorder, were selected from that study. All subjects gave written informed consent, and the study was approved by the Emory

University Institutional Review Board. Diffusion-weighted images were acquired using a Siemens 3T with a 12-channel parallel imaging phase-array coil. Foam cushions were used to minimize head motion. Diffusion MRI data were collected with a diffusion weighted SE-EPI sequence (Generalized Autocalibrating Partially Parallel Acquisitions [GRAPPA] factor of 2). A dual spin-echo technique combined with bipolar gradients was used to minimize eddy-current effects. The parameters used for diffusion data acquisition were as follows: diffusion-weighting gradients applied in 60 directions with a b value of 1000 s/mm²; TR/TE of 8500/95 ms; FOV of 216× 256 mm²; matrix size of 108× 128; resolution of 2× 2× 2 mm³; and 64 slices with no gap, covering the whole brain. Averages of 2 sets of diffusion-weighted images with phase-encoding directions of opposite polarity (left-right) were acquired to correct for susceptibility distortion. For each average of diffusion-weighted images, 4 images without diffusion weighting (b=0 s/mm²) were also acquired with matching imaging parameters. The total diffusion MRI scan time was approximately 20 minutes. T1-weighted images were acquired with a 3D MPRAGE sequence (GRAPPA factor of 2) for all participants. The scan protocol, optimized at 3T, used a TR/TI/TE of 2600/900/3.02 ms, flip angle of 8°, volume of view of 224× 256× 176 mm³, matrix of 224× 256× 176, and resolution of 1× 1× 1 mm³. Total T1 scan time was approximately 4 minutes.

The resulting DTI data were processed using the FMRIBs Diffusion Toolbox (FDT) provided by FSL (FMRIB 4 Software Library) [57]. The FDT probabilistic tractography parameters were set to their default values (number of streamlines=5000, maximum number of steps=2000, loop check: set, curvature threshold=0.2, step length=0.5mm, no distance bias correction).

We applied MANIA on 18 corticolimbic ROIs. All ROIs are localized in Montreal Neurological Institute (MNI) standard space using the Automated Anatomical Labeling (AAL) [58] of the WFU PickAtlas toolbox [59]. The ROI acronym as well as the number of voxels in each ROI are shown in Table ???. The shape of these ROIs are not dilated and



$$T_i(q, j) = \frac{2}{2+1} = \frac{2}{3}$$

Figure 2.1: Running tractography with streamlines from a seed voxel q in the i th ROI to the j th target ROI. Two streamlines hit the target ROI, therefore $T_i(q, j) = \frac{2}{3}$.

we adhere to the standard masks provided in the WFU PickAtlas toolbox. We chose these ROIs because they are known to play a significant role in various psychiatric disorders such as MDD, PTSD, OCD, anxiety and addiction [60, 61] [35, 21] . This sample of ROIs includes cortical, subcortical and limbic regions. Specifically, the cortical ROIs are BA6, BA9, BA10, BA40, BA46, BA47, the limbic are BA24, Th, BS, and the sub-cortical are BA11, BA25, BA32, Acb, Amg, Hp, Ht, Ins, Pc.

MANIA inputs

The proposed network inference method requires the following inputs:

1. A set of N ROIs that represent the nodes of the structural brain network. The i th ROI is a spatially connected cluster of v_i voxels ($i = 1 \dots N$). The selection of ROIs is important [62, 63] but outside the scope of MANIA. MANIA attempts to find the anatomic network between the given ROIs independent of whether the latter are defined by an expert neuroanatomist or by a data-driven method. For instance, ROIs may correspond to different Brodmann areas or other anatomical atlases [64] [65]. Or, it could be that the spatial extent of ROIs results from the analysis of fMRI

data [66, 67, 68, 69]. The results of tractography depend on the selection of ROIs, including their size, shape, grey/white matter composition, and their distance to other ROIs.

2. The results of the tractography process between the previous N ROIs. We assume that the tractography results are structured as voxel-to-ROI matrices (i.e., streamlines are generated from each voxel towards each target ROI), instead of voxel-to-voxel or ROI-to-ROI matrices. Specifically, we represent the output of tractography with N matrices $T_i (i = 1 \dots N)$, defined as follows. The i th ROI corresponds to a matrix T_i with v_i rows (i.e., the number of voxels in that ROI) and N columns. The element (j, k) of matrix T_i represents *the fraction of tractography-generated streamlines that originate from the seed voxel j of the i th ROI and reach any voxel of the k th ROI*. The same number of streamlines is generated for every seed to target pair (j, k) . The i th column of matrix T_i is set to zero, meaning that we do not consider edges from an ROI back to itself (even if such fibers exist). Figure 2.1 illustrates this notation. Since we have N ROIs, there will be N input matrices, one for each source ROI.

2.2.1 Connectivity threshold τ

How can we decide whether voxel j of the i th ROI connects to the k th ROI given the fraction $T_i(j, k)$? The simplest approach is to examine if $T_i(j, k)$ is larger than a given “connectivity threshold” τ ($0 \leq \tau \leq 1$). In MANIA, τ is not a given threshold but an optimization variable, as described in the next section.

As in prior work, we assume that the i th ROI is connected to the k th ROI as long as at least one voxel of the former is connected to the latter [70, 53]. This assumption is not central to MANIA however, and it can be easily replaced with a stronger connectivity constraint.

2.2.2 Network inference as an optimization problem

For a given value of τ , we can identify the voxels of the i th ROI that connect to every other ROI. If this process is repeated for every i , we can construct a directed network in which the i th ROI is connected to the j th ROI if there is at least one voxel in the former that is connected to the latter for that value of τ . This network can be represented with an adjacency matrix² G_τ , as follows:

$$G_\tau(i, k) = \begin{cases} 1 & \text{if } T_i(j, k) > \tau \text{ for at least one voxel } j \\ 0 & \text{otherwise} \end{cases} \quad (2.1)$$

So, the element (i, k) of this matrix is equal to one if there is a (directed) edge from the i th ROI to the k th ROI. The diagonal entries of G_τ are set to zero because we do not consider streamlines from a source ROI back to itself.

We define the *asymmetry* $\phi(G)$ of a directed network G as the fraction of edges that are present in only one direction,

$$\phi(G) = \frac{\sum_{i=1}^N \sum_{k=1}^N G(i, k) (1 - G(k, i))}{\sum_{i=1}^N \sum_{k=1}^N G(i, k)} \quad (2.2)$$

The asymmetry of a network G depends on its *density* $\rho(G)$, defined as the fraction of connected node-pairs,

$$\rho(G) = \frac{\sum_{i=1}^N \sum_{k=1}^N G(i, k)}{N(N-1)} \quad (2.3)$$

The more edges a directed network has, the more likely it becomes that a pair of nodes will be connected in both directions, i.e., the higher the density, the lower the asymmetry.

More formally, consider a directed network with K directed edges and N nodes.

² In graph theory, an $N \times N$ adjacency matrix represents a directed and unweighted graph with N nodes as follows: if there is an edge from node i to node j the (i, j) element of the adjacency matrix is 1; otherwise it is 0. The graph (and the corresponding adjacency matrix) are referred to as “weighted” if each edge is associated with a weight, which typically represents the strength of the edge.

The density is $\rho = K/[N(N-1)]$ ($0 < \rho < 1$). To construct such a network randomly, denoted by G_ρ , we simply connect $N(N-1)\rho$ randomly selected but distinct pairs of nodes with directed edges. The expected number of edges that exist in only one direction is $N(N-1)\rho(1-\rho)$ and so the expected value of the asymmetry of G_ρ is:

$$\bar{\phi}(G_\rho) = \frac{N(N-1)\rho(1-\rho)}{N(N-1)\rho} = 1 - \rho \quad (2.4)$$

To quantify the actual asymmetry of an observed network G_τ , we normalize $\phi(G_\tau)$ by the asymmetry that is expected simply due to chance given the density of this network. So, we define the *normalized asymmetry* of G_τ as

$$\Phi(G_\tau) = \frac{\phi(G_\tau)}{\bar{\phi}(G_\tau)} \quad (2.5)$$

which is well defined as long as $\rho(G_\tau) < 1$.

MANIA is based on the following premise: *the inferred directed network should be as symmetric as possible*. The reason is that the tractography process is unable to infer the actual direction (polarity) of the underlying neural fibers. So, if there are some fibers connecting two voxels X and Y, tractography should be able, in principle, to follow this connection in both directions, from X to Y and from Y to X. We do not claim that two connected ROIs are always attached with both afferent and efferent fibers; instead, we argue that tractography is not able to discover the polarity of those fibers and so the corresponding connection should be traceable in both directions.

The presence of some streamlines from some voxels in ROI X to ROI Y does not necessarily mean however that the network inference method will detect an edge both from X to Y and from Y to X; this also depends on the parameter τ . Given that we aim to minimize the asymmetry of the inferred network, MANIA aims to select the value of τ that leads to the lowest possible asymmetry.

The corresponding optimization problem can be stated as follows: determine the

DTI data - subject 17

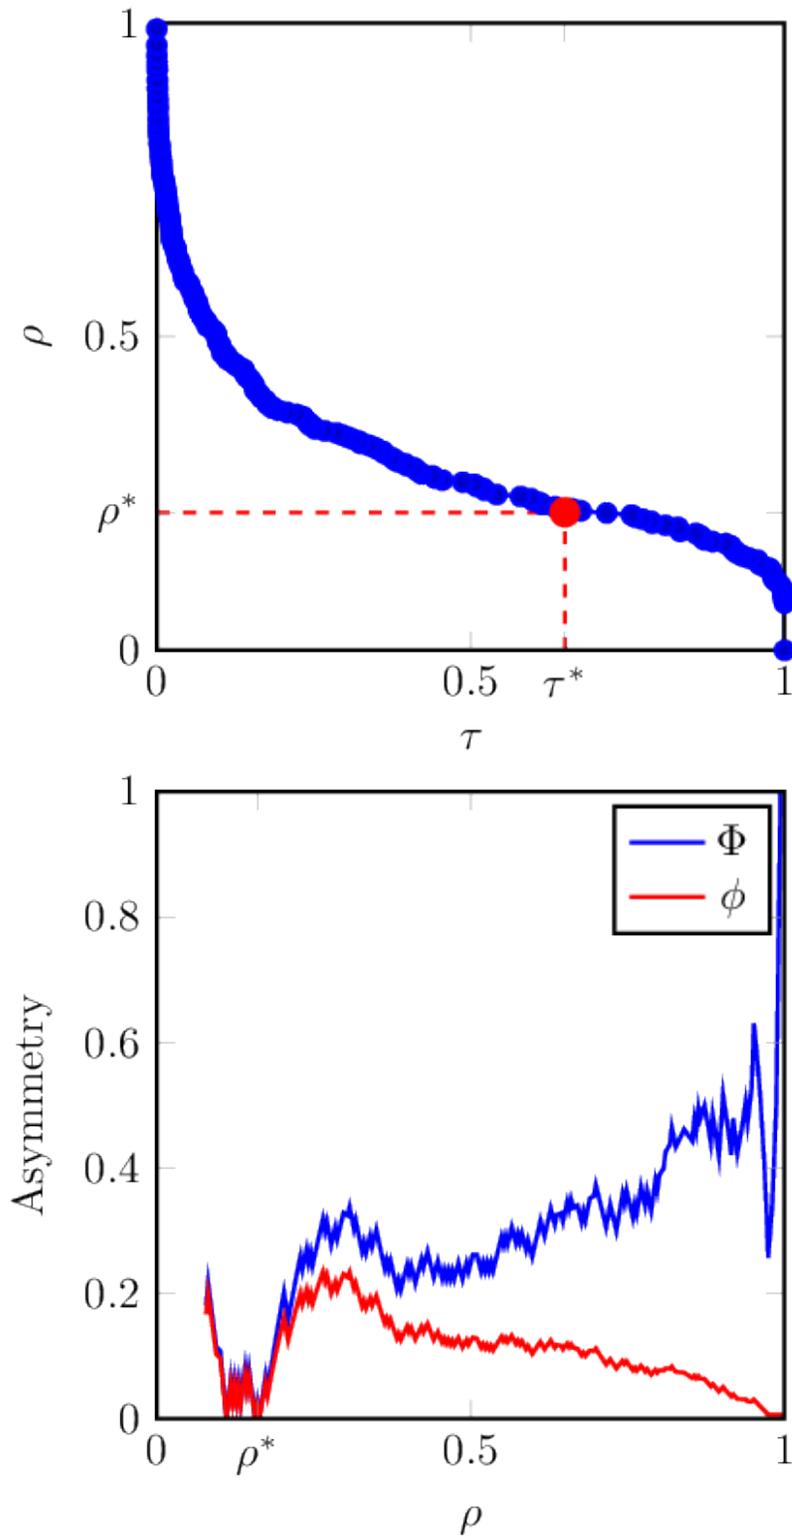


Figure 2.2: Top: Network density ρ as a function of the connectivity threshold τ (plotted for one subject in our DTI dataset). Bottom: Network asymmetry (red) and normalized network asymmetry (blue) as functions of the network density ρ . The optimal density ρ^* is the largest value that minimizes the normalized network asymmetry.

adjacency matrix $\widehat{G} = G_{\tau^*}$, where τ^* is a value of the connectivity threshold that minimizes the normalized asymmetry of G_τ across all possible values of τ ,

$$\tau^* = \arg \min_{0 < \tau < 1} \Phi(G_\tau) \quad (2.6)$$

So, MANIA is based on the premise that there is an ideal value (or range of values) of the connectivity threshold that can correctly classify every directed pair of ROIs as either “connection exists” or “connection does not exist”. When such a threshold exists, it will result in a completely symmetric network (because a perfectly accurate tractography-based network cannot be asymmetric). On the other hand, if such an ideal threshold does not exist (for instance, it may be that two connected ROIs are too far from each other and tractography cannot “see” their connection, or that it is impossible for streamlines to cross the white matter/grey matter boundary of a certain ROI in one direction but not in the opposite), then MANIA aims to at least minimize the normalized asymmetry metric, even if the resulting network will not be completely symmetric.

If there is more than one value of τ that results in the same minimum of the normalized asymmetry (potentially zero), MANIA reports the network with the largest density. The rationale behind this tie-breaker is to avoid trivial solutions that include only a subset of the actual network edges. The previous optimization problem can be solved numerically by scanning the range of τ values with a given resolution.³¹ The density of the resulting network is denoted by

$$\rho^* = \rho(G_{\tau^*}) \quad (2.7)$$

As an illustration of the previous method, Figure 2.2 shows how the network asymmetry (both ϕ and Φ) varies with the density ρ as well as the relation between ρ and τ for the dataset that corresponds to one of the subjects in our case-study.

³¹Since we set the number of streamlines to 5000, the minimum resolution is $\tau = \frac{1}{15000} = 0.0002$.

2.2.3 Threshold-based network inference with post-symmetrization

A common network inference method is to rely on a given connectivity threshold τ , as shown in Equation 2.1. This threshold is sometimes chosen to achieve a certain network density or to ensure that the network is connected [50]).

Given that tractography cannot detect the direction of inferred edges, the resulting network can be then “post-symmetrized” as follows. Consider two network nodes i and j and suppose that the fraction of streamlines from i to j is denoted by $T_{i,j}$. Suppose that $T_{i,j} > \tau$ but $T_{j,i} < \tau$. We can resolve the conflicting evidence between the two directions of this edge by comparing the ratio $(T_{i,j} - \tau) / (1 - \tau)$ that reflects our confidence that the edge from i to j exists, with the ratio $(\tau - T_{j,i}) / \tau$ that reflects our confidence that the edge from j to i does *not* exist.

This post-symmetrization step is different than MANIA in several ways. First, post-symmetrization relies on an arbitrary connectivity threshold to make all edges symmetric, while MANIA selects the threshold value that minimizes the asymmetry metric Φ . Second, post-symmetrization considers each pair of nodes individually, while MANIA considers the entire network, normalizing the observed asymmetry by the expected asymmetry of a random network of equal density. Third, post-symmetrization always results in a symmetric network, while MANIA may not do so if there is no value of the connectivity threshold value that would result in perfect symmetry. Optionally, post-symmetrization can also be applied on the output of MANIA, if the resulting MANIA network is not completely symmetric.

2.2.4 Performance metrics

MANIA can be viewed as a binary classifier: each possible directed edge is classified as *present* or *absent*. We evaluate MANIA based on the following standard metrics for binary classification: the *false positive rate* (or false alarm) p_f , and the *false negative rate* (or miss detection) p_m . The former is defined as the fraction of absent edges that are

incorrectly classified as present, while the latter is defined as the fraction of present edges that are incorrectly classified as absent.

Also, the *Jaccard similarity* between the sets of edges $E(G)$ and $E(\hat{G})$ of the actual network G and the MANIA network \hat{G} , respectively, is defined as

$$J(G, \hat{G}) = \frac{|E(G) \cap E(\hat{G})|}{|E(G) \cup E(\hat{G})|} \quad (2.8)$$

and it varies between zero (no common edges) and one (identical networks).

2.7 Optimal threshold-based network inference

We can also compare the network that results from MANIA with the network that would result if we knew the optimal value τ^{opt} of the connectivity threshold, i.e., the value of τ that maximizes the Jaccard similarity between the inferred network G_τ and the ground truth network G :

$$\tau^{opt} = \arg \max_{0 < \tau < 1} J(G_\tau, G) \quad (2.9)$$

Even though it is not possible to know this optimal threshold value when analyzing real tractography data, we can easily compute its value (or range of values) in experiments with synthetically generated networks, where the ground-truth network G is known.

2.2.5 Edge ranking and confidence metric in MANIA

The output of MANIA is an unweighted directed network. We can quantify the level of confidence we have in each edge with the following edge ranking scheme.

Let ρ_α be the minimum network density at which edge α is present. If the edge α is from a source X to a target Y , the lower ρ_α is, the higher the fraction of streamlines from X to Y . Consequently, we can rank edges so that we are more confident in the presence of edge α than of edge β if $\rho_\alpha < \rho_\beta$ (see Figure 2.3).

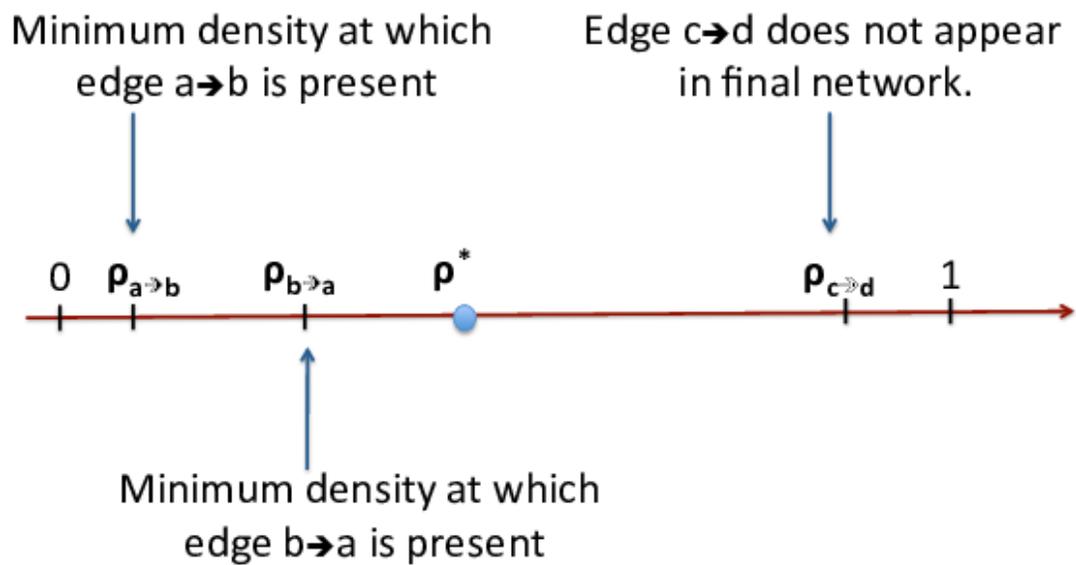


Figure 2.3: As we decrease the connectivity threshold, each edge first appears at a certain value of the network density. If this density is larger than ρ^* , the corresponding edge is *not* present in the MANIA network.

We define a *confidence metric* for an edge α that is present in the MANIA network (i.e., $\rho_\alpha < \rho^*$) as follows

$$C(\alpha) = \frac{\rho^* - \rho_\alpha}{\rho^*} \quad (2.10)$$

$C(\alpha)$ varies from 0 (the edge is only marginally present) to 1 (highest confidence that the edge is present). Similarly, if edge α is absent from the MANIA network (i.e., $\rho_\alpha > \rho^*$), its confidence metric is defined as

$$C(\alpha) = \frac{\rho^* - \rho_\alpha}{1 - \rho^*} \quad (2.11)$$

and $C(\alpha)$ varies from 0 (the edge is only marginally absent) to -1 (highest confidence that the edge is absent).

We also define a confidence metric for a pair of nodes (X, Y) , as the arithmetic mean of the confidence metric of the two directed edges between X and Y ,

$$C(X, Y) = \frac{C(X \rightarrow Y) + C(Y \rightarrow X)}{2} \quad (2.12)$$

Note that one of the two edges may be present while the other may be absent. In that case, the confidence of the corresponding node-pair will be less than the confidence of the present edge.

Note that this edge confidence metric is not related to connection “strength” or “quality”, and the resulting network is still meant to be interpreted as an unweighted graph.

2.2.6 Group analysis using MANIA

If the objective is to create a single “average network” based on data from several subjects, the question is how to best aggregate the tractography data from the given group. One approach is to average the diffusion MRI data, after transforming them in a standard space. Another approach is to average the fraction of streamlines from a given seed voxel to

a given target ROI, across all subjects. These approaches are sensitive to outliers, variations in the diffusion MRI process across subjects, tractography errors and mapping/warping into a standard space. A third approach could be to construct an individual network for each subject, perhaps using MANIA, and then construct an aggregate network only keeping those edges that appear in a large fraction of subjects. This approach requires a group-level threshold to represent the minimum fraction of subjects that should have a connection. For instance, de Reus et al. have proposed a statistically rigorous method to compute such a threshold [62].

Here, we propose a different group analysis method, referred to as *group-MANIA*, that is based on the aggregation of edge-rankings across subjects. The rank-based nature of this method makes it robust to outliers.

As in the previous section, the edges of a subject can be ranked based on the minimum network density at which an edge first appears. We are more confident in the presence of edge α than of edge β if $\rho_\alpha < \rho_\beta$ (see Figure 3).

Suppose that we compute an edge rank vector R_m for each subject m , so that the lowest rank $R_m(1)$ corresponds to the edge for which we are most confident. The number of possible (not necessarily present) directed edges is the same for all subjects: $N(N-1)$ where N is the number of network nodes.

Given a group of size M , we have M distinct rank vectors R_m , $m = 1 \dots M$. The computational problem of *rank aggregation* [71] is to compute an optimal permutation R of the $N(N-1)$ possible edges that captures as well as possible the ordering relations in the M input rank vectors. Specifically, the *Kemeny distance* between two rank vectors R_1 and R_2 is defined as

$$\sum_i \sum_j \delta_{i,j}(R_1, R_2) \tag{2.13}$$

where $\delta_{i,j}(R_1, R_2) = 1$ if R_1 and R_2 disagree in the relative position of elements i and j , and zero otherwise. Rank aggregation aims to compute a vector \hat{R} that minimizes

the cumulative Kemeny distance between \widehat{R} and all input rank vectors R_m . It is an NP-hard problem, and so it is typically solved heuristically. We use the Quicksort algorithm [72] since it has been shown to provide a good approximation of the optimum solution. QuickSort selects a random edge as pivot at each recursive step, while the remaining edges are separated in a left and right list. The left list includes edges that have a lower rank than the pivot in the majority of the subjects; similarly for the right list. The algorithm proceeds recursively in the left and right lists until all edges are ordered.

After computing the optimal aggregate rank vector, we apply MANIA on R to compute the network with the minimum normalized asymmetry (as in the case of a single subject). Note however that the input to MANIA in this case is an ordered list of edges R rather than the set of connectivity matrices T (see section 2.2). Group-MANIA forms a network with the *first* $K = N(N - 1)\rho$ edges in \widehat{R} , and it computes the normalized asymmetry of that network. It then repeats this step, for all values of K , to identify the network with the minimum value of Φ . We refer to the resulting network as the *rank-aggregated network*.

2.2.7 Synthetically generated networks

To evaluate the accuracy and sensitivity of MANIA in a reliable manner we need to rely on synthetic networks rather than actual DTI and tractography data. The benefit of these computational experiments is that we can test MANIA under a wide range of noise conditions and for arbitrary network densities. Unfortunately there are no good statistical models for the noise in DTI and tractography data [51]. We evaluate MANIA based on a simple noise model that is based on the theory of *maximum entropy distributions*, as described next.

For simplicity, each ROI of the synthetically generated networks is simply a voxel. Modeling multi-voxel ROIs in these simulation experiments would not add any new insights. Suppose that the directed network between N nodes is represented by the $N \times N$

adjacency matrix G . Let $T_{i,j}$ be the fraction of streamlines that originate from node i and terminate at node j . Ideally, in the absence of any noise in the DTI data and without any errors in the tractography process, it should be that

$$T_{i,j} = \begin{cases} 1 & \text{if } G_{i,j} = 1 \text{ or } G_{j,i} = 1 \\ 0 & \text{if } G_{i,j} = 0 \text{ and } G_{j,i} = 0 \end{cases} \quad (2.14)$$

So, if there is an edge between nodes i and j in either direction, the fraction of streamlines from node i to node j should be 100% ; otherwise, it should be zero.

In practice, there is significant noise in DTI data and the tractography process can be error-prone, especially when the ROIs are in grey matter and/or when neural fibers cross each other, split or merge, or fan out as they approach their targets. Consequently, the tractography output may show that some streamlines do not reach from node i to node j even when the two nodes are connected, or that some streamlines get from i to j even when there is no connection between the two nodes. We model these errors probabilistically, as follows:

$$T_{i,j} = \begin{cases} 1 - Z_1 & \text{if } G_{i,j} = 1 \text{ or } G_{j,i} = 1 \\ Z_2 & \text{if } G_{i,j} = 0 \text{ and } G_{j,i} = 0 \end{cases} \quad (2.15)$$

where Z_1 and Z_2 are two (generally different) random variables with $[0,1]$ support. If their probability mass is concentrated close to 0, the results of the tractography process are not significantly affected by noise. On the other hand, if these two random variables are uniformly distributed in $[0,1]$, the tractography results are completely random and any network inference process is hopeless.

We model the two random variables Z_1 and Z_2 with the *Maximum Entropy distribution with one-degree of freedom*. In this case, this distribution is the truncated exponential

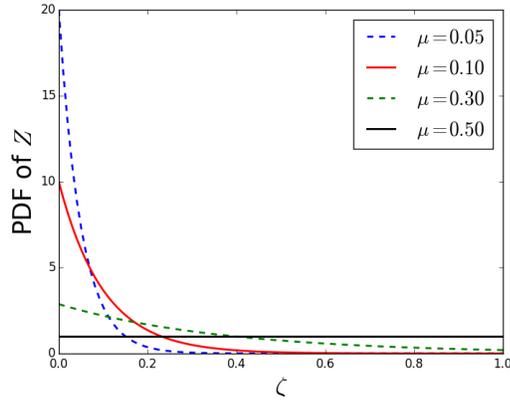


Figure 2.4: Probabilistic error model (Z) for tractography-generated connection probabilities using the maximum entropy distribution (with one degree of freedom).

distribution with support $[0,1]$,

$$f_Z(\zeta) = \begin{cases} \frac{\alpha}{1-e^{-\alpha}} e^{-\alpha\zeta} & \zeta \in [0, 1] \\ 0 & \text{otherwise,} \end{cases} \quad (2.16)$$

where $\alpha > 0$ is a parameter that determines the mean and variance of the distribution.

Instead of controlling α , we control the intensity of noise through the mean of Z ,

$$\mu = E[Z] = \frac{1 - (1 + \alpha) e^{-\alpha}}{\alpha(1 - e^{-\alpha})} \quad (2.17)$$

The two distributions Z_1 and Z_2 follow this statistical model with means μ_1 and μ_2 respectively. Figure 2.4 shows the previous distribution for four values of μ . Note that the distribution Z becomes almost “flat” (close to the uniform distribution) when its mean is higher than 0.3, meaning that tractography would be extremely inaccurate when the noise intensity exceeds that level. In the rest of this chapter we limit the range of μ_1 and μ_2 between 0 and 0.3.

2.3 Results

2.3.1 Evaluation with synthetic data

We evaluate MANIA based on computational experiments with synthetic data and random networks. The “ground-truth” networks G are constructed as follows. Suppose that G has N nodes and density ρ . We place $\lfloor \rho \frac{N(N-1)}{2} \rfloor$ undirected edges between randomly selected but distinct pairs of nodes. Note that G is symmetric by construction because the tractography process cannot infer the true polarity of the underlying neural fibers. Given G , we then create the tractography matrix T that represents the “noisy” fraction of streamlines between any pair of nodes, as shown in equation 2.15. Note that the fraction of streamlines from a node X to a node Y is typically different than the fraction of streamlines from Y to X . In the following experiments, N is set to 50 nodes, and each experiment is repeated for 1000 networks G .

We first examine the effect of post-symmetrization on the accuracy of both threshold-based network inference and MANIA. In the former, the connections are determined based on a given threshold τ_0 , as discussed in Section 2.5. We denote the Jaccard similarity between the inferred network and the ground-truth network with J_{SYM} when post-symmetrization is performed, and with J_{NO-SYM} otherwise. Figure 2.5 shows the difference $\Delta J = J_{SYM} - J_{NO-SYM}$ for several choices of τ_0 as well as for MANIA. Each box-plot is generated from 1000 experiments; in each experiment we generate a random network with density between 0 and 1, while the noise parameters μ_1 and μ_2 are uniformly distributed between 0 and 0.3. The red line corresponds to the median, the box boundaries correspond to the 25th and 75th percentiles, while the dashed lines show the 10th and 90th percentiles. In all cases, $\Delta J > 0$ (one-sided Mann-Whitney U test p-values shown next to each box plot), meaning that *post-symmetrization helps to improve the accuracy of network inference*. This is true for both MANIA and threshold-based inference, even though the improvement is larger for the latter. Because of the positive effect of post-symmetrization, in

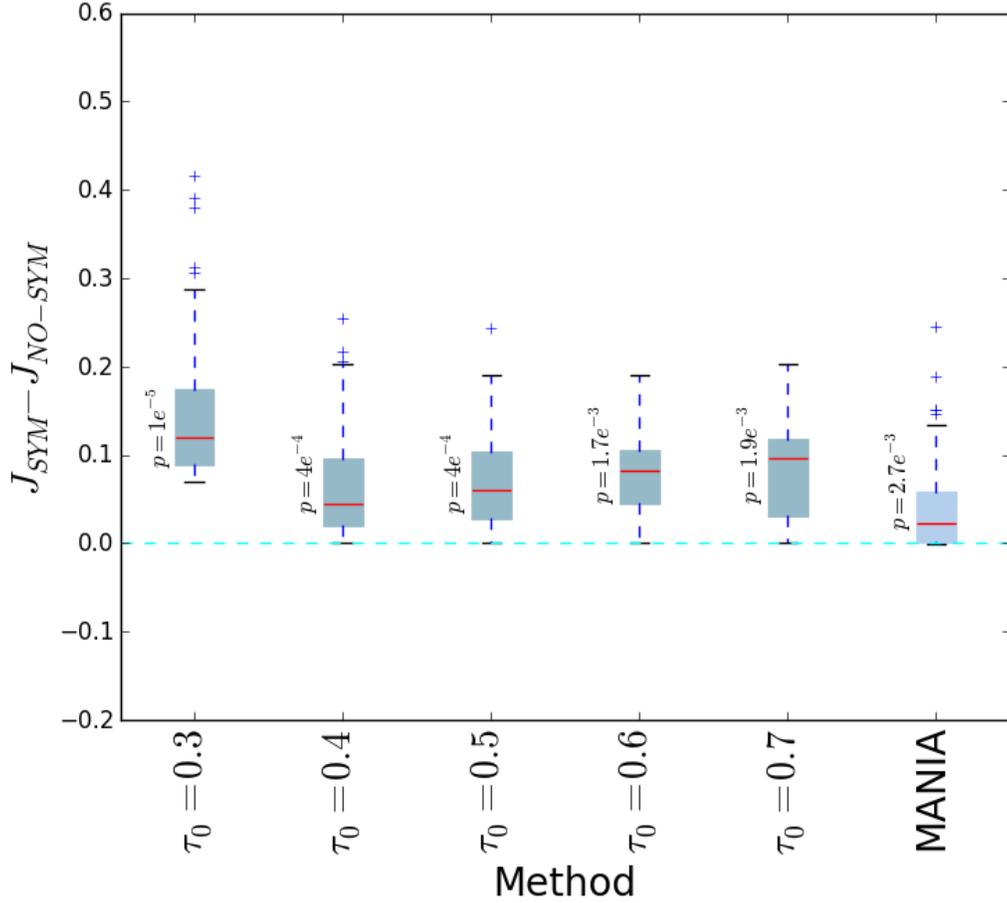


Figure 2.5: The Jaccard similarity difference ($\Delta J = J_{SYM} - J_{NO-SYM}$) with and without post-symmetrization for five threshold values and for MANIA. Each box plot is generated from 1000 experiments; in each experiment we generate a random network with density between 0 and 1, while the noise parameters μ_1 and μ_2 are uniformly distributed between 0 and 0.3. The red line corresponds to the median, the box boundaries to the 25th and 75th percentiles, while the dashed lines show the 10th and 90th percentiles. In all cases, $\Delta J \geq 0$ (one-sided Mann-Whitney U test p-values are shown next to each box plot) meaning that *post-symmetrization helps to improve the accuracy of network inference*.

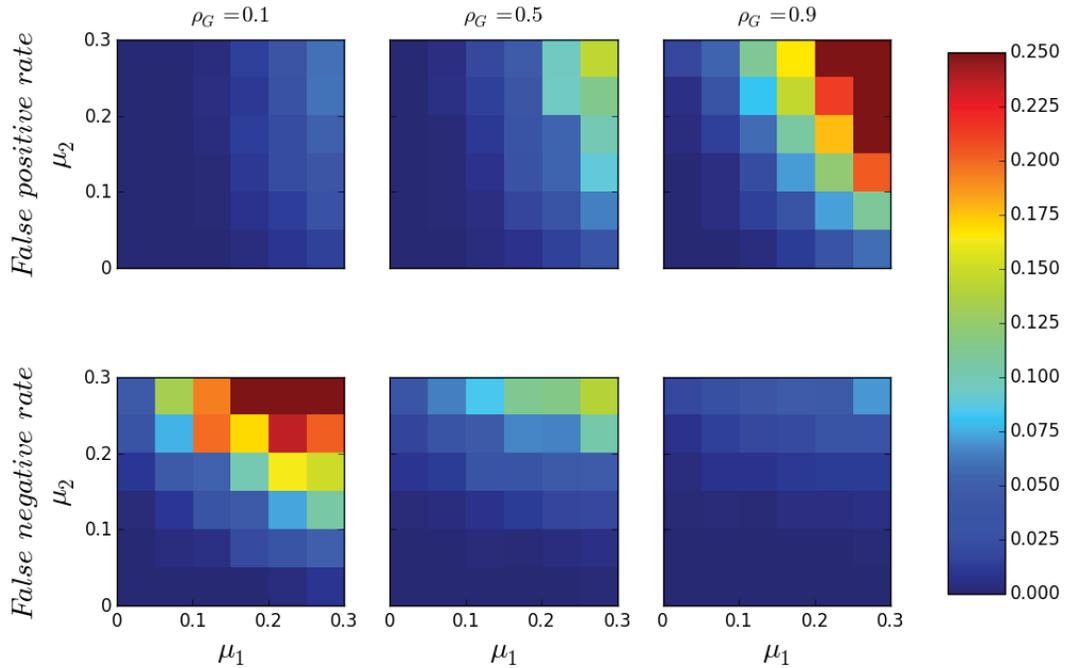


Figure 2.6: False positive rate and false negative rate of MANIA as a function of μ_1 and μ_2 for sparse networks ($\rho_G = 0.1$), medium density networks ($\rho_G = 0.5$) and dense networks ($\rho_G = 0.9$). Each square is the average of 1000 independent simulations.

the rest of the chapter, we apply it in all network inference experiments.

Figure 2.6 illustrates the performance of MANIA in the two-dimensional space defined by the noise parameters μ_1 and μ_2 for three values of the network density. Each square in these heat maps is the median across 1000 experiments. The false positive and false negative rates are close to 0 (less than 5%) in the lower-left half of each heat map (i.e., when $\mu_1 + \mu_2 < 0.3$). For higher values of the noise intensity, the accuracy of MANIA depends on the density of the underlying network. In the case of sparse networks, MANIA also infers a sparse network and most errors are false negatives, i.e., MANIA does not detect some of the few existing edges. For dense networks, MANIA also infers a dense network and most errors are false positives, i.e., MANIA detects a few extra edges that do not actually exist. In mid-range densities, the errors are more balanced between false positives and false negatives. In all cases the maximum false positive (or negative) rate

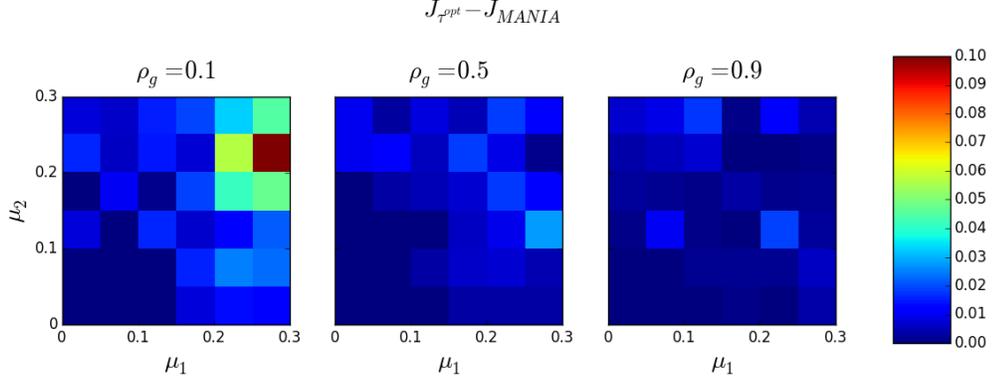


Figure 2.7: The Jaccard similarity difference ($\Delta J = J_{\tau^{opt}} - J_{MANIA}$) between MANIA and the optimal threshold-based scheme as a function of μ_1 and μ_2 for sparse networks ($\rho_G = 0.1$), medium density networks ($\rho_G = 0.5$) and dense networks ($\rho_G = 0.9$). Each square is the average of 1000 independent simulations.

when $\mu_1 = \mu_2 = 0.3$ is less than 25%. Recall from Figure 2.4 that these noise intensity levels should be considered quite high in practice.

Figure 2.7 compares the MANIA-inferred network with the network that corresponds to the optimal threshold τ^{opt} , as discussed in Section 2.7. Specifically, the heat maps of Figure 2.7 compare the Jaccard similarity J_{MANIA} between MANIA and the ground-truth network, with the Jaccard similarity $J_{\tau^{opt}}$ between the optimal threshold based network and the ground-truth network. The accuracy of MANIA is typically close to that of the optimal threshold method. Even under the highest noise intensity we consider ($\mu_1 = \mu_2 = 0.3$), J_{MANIA} is only 10% lower than $J_{\tau^{opt}}$. These results suggest that MANIA selects automatically a connectivity threshold value that results in almost optimal accuracy, across all possible such threshold values.

Finally, Figure 2.8 compares MANIA with five given threshold values τ_0 . The comparison is in terms of the Jaccard similarity difference $\Delta J = J_{MANIA} - J_{\tau_0}$. As in Figure 2.5, each box-plot is generated from 1000 experiments in which we vary the network density between 0 and 1, and the noise parameters μ_1 and μ_2 between 0 and 0.3. The median ΔJ is always positive and the distribution of ΔJ is skewed towards positive

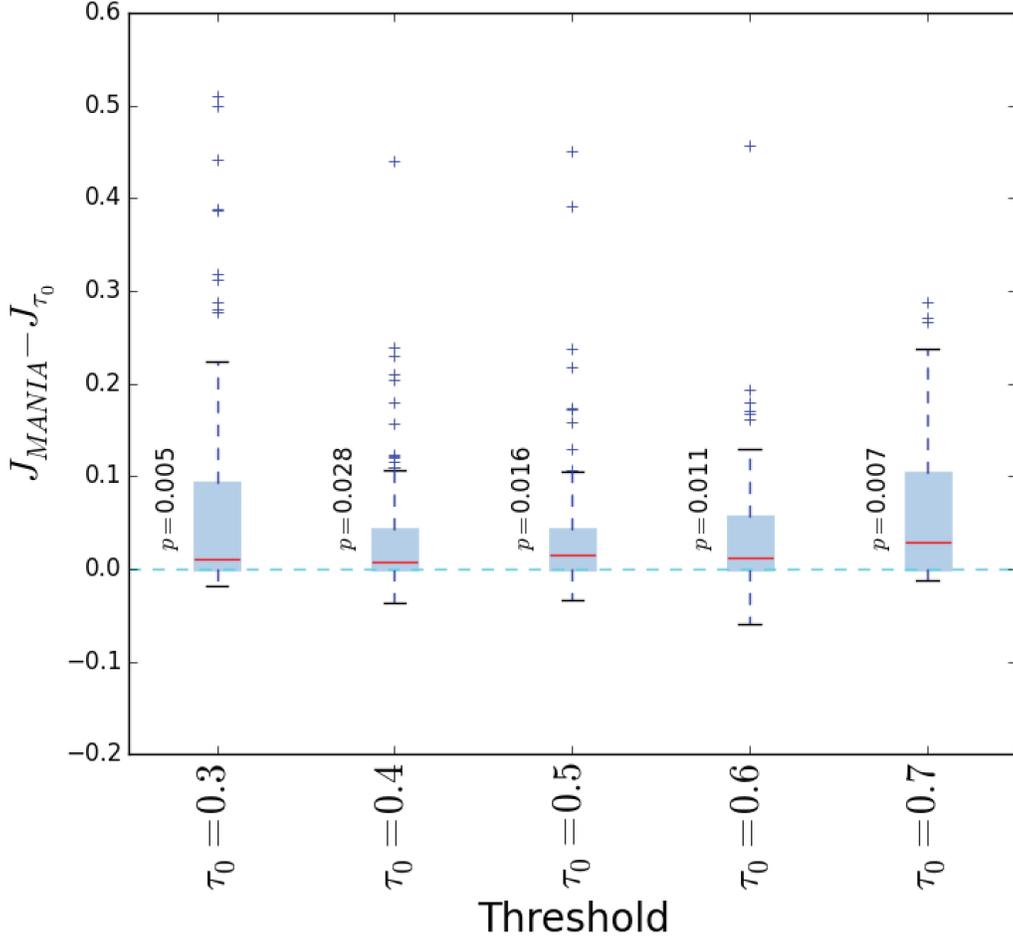


Figure 2.8: The Jaccard similarity difference ($\Delta J = J_{MANIA} - J_{\tau_0}$) between MANIA and five given threshold values. The accuracy comparisons are made after post-symmetrization. Each box plot is generated from 1000 experiments; in each experiment we generate a random network with density between 0 and 1, while the noise parameters μ_1 and μ_2 are uniformly distributed between 0 and 0.3. The red line corresponds to the median, the box boundaries to the 25th and 75th percentiles, while the dashed lines show the 10th and 90th percentiles. In all cases, $\Delta J \not\leq 0$ (one-sided Mann-Whitney U test p-values are shown next to each box plot) meaning that *MANIA is more accurate than inferring the network based on a fixed threshold.*

values (one-sided Mann-Whitney U test p-values shown next to each box plot), meaning that MANIA typically performs better than a fixed threshold scheme, independent of the selected threshold.

2.4 Discussion

Thomas et al. have recently shown that inferring long-range anatomical connections between grey matter ROIs from DWI data can be significantly inaccurate [73]. The authors also note that “(probabilistic tractography methods) are less susceptible to changes in the composition of an ROI but only if an optimized threshold can be derived and used.” More recently Reveley et al. [56] have investigated the key reasons behind the negative results of [73]. They showed that the dense system of white matter fibers residing just under the cortical sheet poses severe challenges for long-range tractography, concluding that it is “extremely difficult to determine precisely where small axonal tracts join and leave larger white matter fasciculi.”

In light of the previous results, we believe that there is a need for new network inference methods. MANIA is moving in the right direction for the following reasons:

1. The results of [73] suggest that probabilistic tractography can be more accurate than deterministic tractography in terms of sensitivity and specificity as long as its parameters are appropriately optimized. MANIA is indeed mostly applicable to probabilistic tractography, and its main focus is how to “self-configure” its connectivity threshold τ in an optimized manner, relying on what we expect to be true about the structure of the resulting solution (namely, a symmetric network).
2. The results of [56] suggest that it is risky to artificially dilate the given ROIs, which are typically mostly grey matter, so that they also include some white matter voxels. Those voxels may be part of the white matter fiber systems that reside just under the cortical sheet. In other words, if our goal is to understand the connectivity between

grey matter ROIs, we should not use tractography seeds that reside in white matter; instead, we need to seed from grey matter even if the diffusion signal is much weaker there. So, we need to expect that some connections may appear as asymmetric, which is what MANIA anticipates.

3. The results of [56] can be interpreted as follows: because it is hard for any tractography method to accurately cross the white matter/grey matter boundary (WGB), especially in the case of cortical ROIs, a network inference method should be able to deal somehow with erroneous measurements about specific connections. In other words, just because tractography failed to cross the WGB going from seed X to target Y does *not* mean that we should conclude that X and Y are not connected. And so, given that the input data about individual connections is quite noisy, we need to examine if there is any additional “hidden structure” in the inference problem that we can exploit. If this is case, we can then look for a solution that satisfies the constraints of that additional structure. In MANIA, this “hidden structure” in the inference problem is that the resulting network should be as symmetric as possible.

Of course, we do *not* claim that MANIA addresses every concern about tractography-based network inference. On the contrary, there are more open issues that need to be addressed. Two of them are further discussed next.

Even if the thresholding problem is adequately addressed with MANIA, there is another important problem in structural network inference: the *distance bias* of the tractography process [74]. It is harder to discover connections between distal regions due to the accumulation of uncertainty in long streamlines, causing false negatives for long-range connections [46]. Additionally, it is more likely to incorrectly detect connections between proximal regions, especially in the presence of crossing or turning fibers, causing false positives. The FDT toolbox provides a “distance correction” option by multiplying the number of streamlines that cross a voxel by the average length of those streamlines⁴¹ there is no ev-

⁴¹ <http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FDT/UserGuide>

idence however that this simple form of distance correction is able to improve significantly the accuracy of the network inference process [46]. A more sophisticated method is that of [75], which compares the tractography-generated connectivity probabilities with a null model that gives the corresponding connectivity probabilities with a random tracking process that is dominated by the same distance effects. We view distance correction methods as an independent processing step that can be applied prior to applying MANIA. For instance, the method by Morris et al. first creates a “null frequency of connection map”, it then filters the “experimental frequency of connection map” that is produced by a probabilistic tractography tool, resulting in the so-called “significance of connection map” (which is supposed to have fewer false positives). MANIA can be then applied on the latter, rather than on the experimental frequency of connection map. Even though it is still not clear if the Morris distance correction method is sufficient to completely address the distance correction bias [76], we anticipate that the combination of MANIA with the distance correction method of Morris et al. will improve the accuracy of the resulting networks.

A network representation consists of both nodes and edges. The clinical and research value of representing a brain as a network depends critically on the selected nodes and on the exact boundaries of the corresponding ROIs [63]. MANIA assumes that the set of given nodes is sufficiently specified, and that their spatial boundaries are accurately defined. In practice, this step of the network inference process is always an “inexact science” given that the functional role of any given ROI is at best only partially understood and the anatomical boundary of each ROI is subject-dependent [77].

A voxel-level analysis [78] avoids the selection of functionally specific ROIs but it makes it harder to associate the topological properties of the observed network, which now consists of many thousands of nodes, to any known brain circuits and their function. Again, we view this important issue as orthogonal to MANIA: improved brain parcellation methods, such as data-driven parcellations [79, 80] and decreasing voxel sizes can be

used jointly with MANIA to identify structural networks that are consistent, or that can explain well, the observed spatio-temporal correlations in resting-state or task-based fMRI analyses. This coupled exploitation of fMRI and diffusion MRI data has provided valuable insights about the underlying anatomy of the brain structures that result in the Default Mode Network [81], and they can become more common now that the HCP project provides both functional and diffusion data for hundreds of subjects [30].

CHAPTER 3

CASE STUDIES: APPLICATION OF MANIA ON EXPERIMENTAL DATASETS

As a case-study, we first (case study A) apply MANIA on diffusion MRI data from 28 healthy subjects to infer the structural network between 18 corticolimbic ROIs that are associated with various neuropsychiatric conditions such as major depressive disorder (MDD), post traumatic stress disorder (PTSD), obsessive compulsive disorder (OCD), generalized anxiety disorder (GAD) and addictive disorder (AD) [18, 19] [20, 21]. We note however that, even though these ROIs are generally associated with various psychiatric disorders and the aspects of emotional regulation putatively impacted by these disorders, the objective of this work is *not* to infer the network that is associated with any particular disorder.

In second case study (case study B), we apply MANIA to the state-of-the-art dataset provided by HCP, not only to elicit findings about the human brain, but also to provide connectome on a platform that is open, reproducible, and reputable within the neuroscience community. This will lead to a connectome on a dataset that is shared by a large community, which in turn provides reproducible results and facilitates findings based on a common benchmark. MANIA fits as a powerful tool in the HCP ecosystem, whose overarching objective is to map the connectivity gathered by non-invasive imaging techniques (MANIA only focuses on the anatomical goal). We reconstruct and present cortical anatomical connectomes for 37 HCP subjects. The connectomes are reconstructed for left and right hemispheres, and also for inter-hemispheric connections.

Also, in case study B, we employed the most recent widely accepted human brain parcellation (introduced in [23]), known as HCP-MMP. The HCP-MMP parcellation is derived from the combination of the anatomical landmarks and the functional organization of the brain as inferred by fMRI. The higher resolution and functional homogeneity of HCP-MMP parcellation come with the additional benefit of the surface-based parcels (ROIs).

Surface-based parcellation has significant effect on increasing the accuracy when mapping cortical connectivity, since volume-based parcellation cannot capture the many extreme folding structures of the cortex [82]. Also, the volume files have biases in their streamline counts (and are more prone to catch fibers that are passing by, rather than terminating).

Moreover, in this case study, we validate the accuracy of MANIA using data from a high resolution 7T MRI scanner (and compare it to that from a 3T scanner). This further proves the efficacy of MANIA, as previously discussed in chapter 2, where the emphasis was on phantom and synthetic data studies. We leverage three subjects in the HCP dataset who had both 3T and 7T MRI scans. Since 7T provides a considerable boost in signal-to-noise ratio, and spatial resolution, we build a benchmark dataset based on 7T and show that MANIA results run on 3T can predict connections that are highly anticipated by 7T acquisition, while controlling the false positive rate. This proves that MANIA is (nearly) striking an optimal point between specificity and sensitivity.

3.1 Case Study A: corticolimbic circuitry implicated in major depression disorders

3.1.1 Data set

We apply MANIA in Diffusion Tensor Imaging (DTI) data collected by an earlier study: “*Brain aging in humans, chimpanzees (*Pan troglodytes*), and rhesus macaques (*Macaca mulatta*): magnetic resonance imaging studies of macro- and microstructural changes*” [17]. Twenty eight healthy right-handed females between the ages of 18 and 22 (mean: 20.2), without a history of psychiatric disorder, were selected from that study. All subjects gave written informed consent, and the study was approved by the Emory University Institutional Review Board. Diffusion-weighted images were acquired using a Siemens 3T with a 12-channel parallel imaging phase-array coil. Foam cushions were used to minimize head motion. Diffusion MRI data were collected with a diffusion weighted SE-EPI sequence (Generalized Autocalibrating Partially Parallel Acquisitions [GRAPPA] factor of 2). A dual spin-echo technique combined with bipolar gradients was used to minimize

eddy-current effects. The parameters used for diffusion data acquisition were as follows: diffusion-weighting gradients applied in 60 directions with a b value of 1000 s/mm²; TR/TE of 8500/95 ms; FOV of 216× 256 mm²; matrix size of 108× 128; resolution of 2× 2× 2 mm³; and 64 slices with no gap, covering the whole brain. Averages of 2 sets of diffusion-weighted images with phase-encoding directions of opposite polarity (leftright) were acquired to correct for susceptibility distortion. For each average of diffusion-weighted images, 4 images without diffusion weighting (b=0 s/mm²) were also acquired with matching imaging parameters. The total diffusion MRI scan time was approximately 20 minutes. T1-weighted images were acquired with a 3D MPRAGE sequence (GRAPPA factor of 2) for all participants. The scan protocol, optimized at 3T, used a TR/TI/TE of 2600/900/3.02 ms, flip angle of 8°, volume of view of 224× 256× 176 mm³, matrix of 224× 256× 176, and resolution of 1× 1× 1 mm³. Total T1 scan time was approximately 4 minutes.

The resulting DTI data were processed using the FMRIBs Diffusion Toolbox (FDT) provided by FSL (FMRIB 4 Software Library) [57]. The FDT probabilistic tractography parameters were set to their default values (number of streamlines=5000, maximum number of steps=2000, loop check: set, curvature threshold=0.2, step length=0.5mm, no distance bias correction).

We applied MANIA on 18 corticolimbic ROIs. All ROIs are localized in Montreal Neurological Institute (MNI) standard space using the Automated Anatomical Labeling (AAL) [58] of the WFU PickAtlas toolbox [59]. The ROI acronym as well as the number of voxels in each ROI are shown in Table ?? . The shape of these ROIs are not dilated and we adhere to the standard masks provided in the WFU PickAtlas toolbox. We chose these ROIs because they are known to play a significant role in various psychiatric disorders such as MDD, PTSD, OCD, anxiety and addiction [60, 61] [35, 21] . This sample of ROIs includes cortical, subcortical and limbic regions. Specifically, the cortical ROIs are BA6, BA9, BA10, BA40, BA46, BA47, the limbic are BA24, Th, BS, and the sub-cortical are BA11, BA25, BA32, Acb, Amg, Hp, Ht, Ins, Pc.

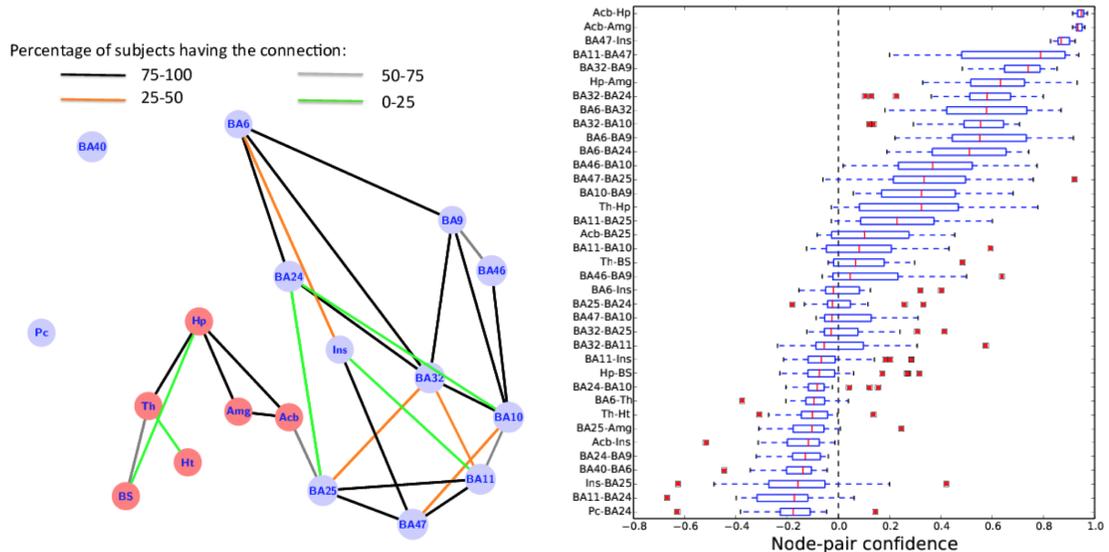


Figure 3.1: The rank-aggregated network, based on DTI data from 28 subjects, between the 18 ROIs in Table 3.1. Every edge in the connected component has been detected in both directions, i.e., MANIA identifies a completely symmetric network (no post-symmetrization is needed). The density of the connected component is 19%. The color of each edge represents the fraction of subjects that have the corresponding edge in their individual MANIA-based networks.

3.1.2 Results

We applied MANIA in the DTI data presented in Section 3.1.1, based on the 18 ROIs listed in Table 3.1. A single rank-aggregated network is constructed, using the group-MANIA, aggregating data from 28 subjects. The rank-aggregated network is shown in Figure 3.1.

Two ROIs (Pc and BA40) appear to not be *directly* connected with the other 16 ROIs; of course there may be indirect connections through other ROIs that have not been included here (we return to this point in Section 4). Every edge in the connected component of Figure 3.1 has been detected in both directions, i.e., MANIA identifies a completely symmetric network in this case (i.e., no post-symmetrization is needed). The density of the connected component (16 ROIs) is 19%. The color of each edge in Figure 3.1 represents the fraction of the 28 subjects that have the corresponding edge in their individual networks

(constructed by MANIA).

We measured the “centrality” of each node in the rank-aggregated network, based on four centrality metrics (degree, closeness, betweenness, PageRank) [83]. Different centrality metrics focus on different notions of importance. For instance, the degree centrality metric associates importance with the number of direct connections a node has; BA32 (Ventral anterior cingulate) has the largest number (six) of direct connections in this network (see Table 3.2). This may be because BA32 is spatially adjacent to both BA10 and BA25, and those ROIs are also of high degree. The betweenness centrality of a node X, on the other hand, focuses on the number of shortest paths between any pair of nodes that go through X; BA25 (subcallosal cingulate) is the most important node from this perspective because it serves as the “unique bridge” between the 6 red nodes at its left and the 9 blue nodes at its right. BA25 is also the most central node in terms of its average distance to all other nodes (closeness centrality).

Similarly, we measured the edge centrality of all connected node pairs. In terms of edge betweenness centrality, the connection between BA25 and the Nucleus Accumbens (Acb) is by far the most central in this network. It is interesting to note that this edge includes the segment of white matter that is the target of Deep Brain Stimulation (DBS) therapies for the treatment of MDD [35]. In fact, the DBS target is typically the point at which the fibers between (BA25-Acb), (BA25-BA32) and (BA25-BA24) intersect.

Figure 3.1 also shows some percentiles of the per-subject node-pair confidence metric (median, 25-75th percentiles, 10-90th percentiles, and outliers) for each node-pair that appears connected in at least one of the 28 subjects. The connections between the following node pairs appear in all subjects and have the highest confidence: Hp-Acb, Amg-Acb, BA47-Ins. On the other hand, the following connections appear only in some subjects and their confidence metric varies around zero: Th-BS, BA46-BA9, BA6-Ins. Some connections that appear in 1-2 subjects but have very low confidence are: Pc-BA24, BA11-BA24, Ins-BA25, BA40-BA6.

3.1.3 Discussion

In our 18-ROI case-study, summarized in Figure 3.1, the use of mostly large ROIs that do not necessarily correspond to distinct functional units, together with the distance bias of the tractography process, may account for the lack of certain expected connections. Two such expected connections are between Pc and BA40 [84], and between BA9 and BA40 [85]; the latter is a long-distance connection. Additionally, large cortical ROIs such as BA9 and BA40 are only imprecisely defined, which may also explain the absence of some of their connections. The limbic and subcortical ROIs, on the other hand, are more precisely defined and their connections are mostly running over shorter distances. These findings suggest that MANIA should be evaluated in the future jointly with, first, advanced distance correction methods, and second, with either more precisely defined ROIs or on a whole-brain parcellation template.

Table 3.1: The 18 corticolimbic ROIs we consider in the case-study.

ROIs	Acronym	Number of voxels
Premotor cortex	BA6	3131
Insula	Ins	1858
Ventromedial prefrontal cortex	BA10	1784
Inferior parietal cortex	BA40	1598
Dorsolateral prefrontal cortex	BA9	1422
Mid-brain and pons	BS	1406
Orbito-frontal cortex	BA11	1243
Thalamus	Th	1100
Hippocampus	Hp	932
Precuneus	Pc	861
Inferior prefrontal gyrus	BA47	851
Ventral anterior cingulate	BA32	721
Dorsal anterior cingulate cortex	BA24	593
Dorsolateral prefrontal cortex	BA46	574
Amygdala	Amg	220
Subcallosal cingulate	BA25	204
Nucleus accumbens	Acb	140
Hypothalamus	Ht	13

Table 3.2: Top-three nodes in rank-aggregated network based on four node-centrality metrics

Centrality	Top-three nodes		
Degree	BA32	BA10	BA25
Closeness	BA25	BA32	Acb
Betweenness	BA25	Acb	Hp
PageRank	BA10	BA32	BA11

3.2 Case Study B: human corticocortical anatomical network

3.2.1 Data set

This case study is carried out using data from the Human Connectome Project (HCP) which is publicly available [15]. We used HCP data in the two S900, and S1200 releases. HCP is a large-scale NIH-funded effort to map the macroscopic wiring diagram of the human brain. HCP has chosen dMRI and probabilistic tractography as the methods of choice for mapping the in-vivo anatomical connectome which makes their data highly relevant to our studies and compatible with our pipeline described in chapter 2. For our evaluation purposes, we focused on a subset of subjects from the HCP dataset comprising 37 female adults of age 31-35. Restricting the gender and age allows us to minimize the subject variability to the best of our ability so that we focus on the method evaluation and connectome characteristics while also being congruent to the study in chapter 2. The list of all subjects used in this chapter are available in Appendix A.

3.2.2 Node delineation: ROIs

The first step in reconstructing an individual connectome is specifying a parcellation scheme to subdivide cortical gray matter into discrete, spatially contiguous parcels used as the graph nodes in the network inference phase. Parcellation can be thought as dimensionality reduction in connectome studies, reducing the computational time, increasing the signal-to-noise ratio without eliminating valuable information about the interactions between different brain regions. Also, studying the connectome at macroscale parcellation resolution

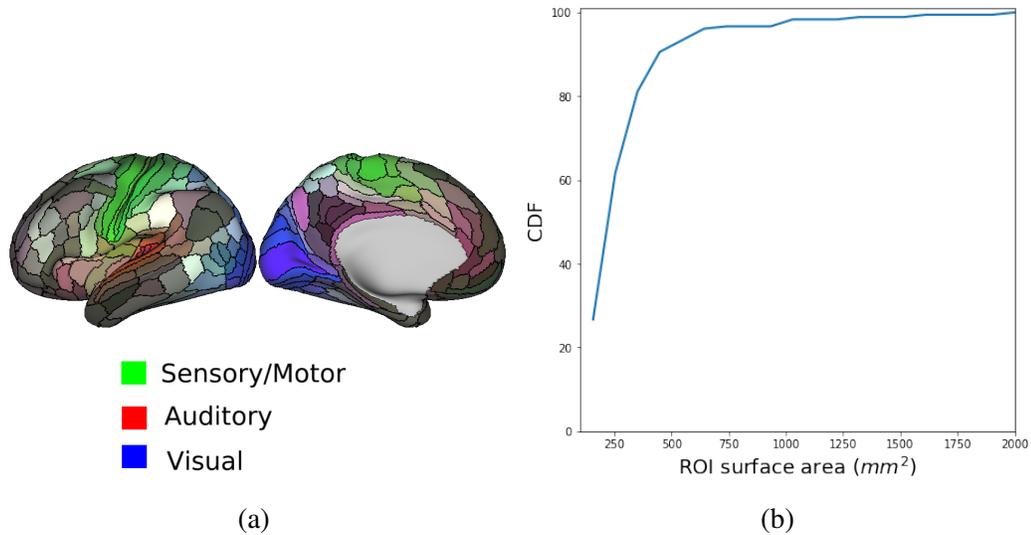


Figure 3.2: a) 180 ROIs from HCP’s multi-modal cortical parcellation (HCP-MMP) projected on the S1200 group average cortical surface (inflated left hemisphere). This visualization is generated by the data from [23] b) The size (surface area) distribution for the 180 ROIs

allows us to unravel the mechanisms through which these interconnections give rise to complex cognitive processes at a high level.

Parcellation is the first step in the pipeline of connectome inference. As a result, parcellation is one of the most critical steps in connectivity inference. Cytoarchitecture or anatomical landmarks have long been used for this task. For example, Brodmann areas have been used to define regions of interest (ROIs) for network analysis [86]. Also, in the previous chapter, we applied MANIA on 18 corticolimbic ROIs localized in Montreal Neurological Institute (MNI) standard space using the Automated Anatomical Labeling (AAL) [58] of the WFU PickAtlas toolbox [59] which is again a parcellation based on anatomical landmarks. The aforementioned parcellations were volumetric, however most recent studies, specifically ones focusing on cortex, are surface-based [87, 23], i.e, the surface of the cortex is divided into regions rather than its volume. Surface-based analysis is believed to yield better results as the human cortex is extremely convoluted [88]. Thus, surface representations offer a better description of the boundaries at any resolution for the cortical

areas. Voxel-based representation of the white-gray matter boundary cannot necessarily follow the highly convoluted boundary of the cortical surface and thus increases the gyral bias [89], consequently decreasing the accuracy of any tractography method.

Besides, anatomical parcellations, surface- or volume-based, some studies choose to use uniform parcellation to control the network nodes for similar size. The nodes in these studies are inherently not linked to either anatomical or functional specificity of the brain directly [90]. However, we chose to stay relevant to the anatomy and functional homogeneity of underlying data, using multi-modal parcellation that employs both. Also, limiting our study to cortex, we looked only for surface-based parcellations. Having all these constraints in the study, we parcellated subjects' brain based on HCP-MMP atlas [23] to 180 regions per hemisphere (see figure 3.2). The HCP-MMP atlas is shown in figure 3.2 with the distribution of surface areas of its regions.

3.2.3 Network inference without an arbitrary threshold or density: MANIA

The network reconstruction process entailed using MANIA on dMRI images for each of the 37 female individuals with the HCP-MMP parcellation. We ran probtrackx as described in [28] (also chapter 2) with three differences. First, Streamline seeding in this work is from white-gray matter boundary, encouraged by [91]. We mapped surface boundary between white-matter and gray matter on a cortical surface mesh with 32000 vertices. Each of these vertices were also assigned to one of the HCP-MMP ROIs. Then, for each vertex in the ROI we assigned a randomly seed point on that ROI surface and initiated 5000 streamlines. Streamlines reaching from the source to the target ROI were used to form matrix T of equation 2.1. Second, since MMP parcellation consists of 1050 adjacent ROIs, we excluded the adjacent pairs from the process of choosing the threshold in MANIA, i.e., MANIA was run to minimize the asymmetry between non-adjacent pairs. The 1050 adjacent ROIs in the MMP parcellation are already strongly and symmetrically connected given they share physical border. Lastly, as opposed to the first case studies that *distthresh* parameter in

probtrackx was set to zero, the *distthresh* parameter in this study was set to 20mm, i.e., the network results are based on streamlines that are at least 20mm long.

3.2.4 Results

Individual networks

The network reconstruction procedure was applied separately to inter and intra-hemispheric connections. First, we looked at the density at which the *normalized asymmetry ratio* (NAR) is minimized for each subject. The first column of figure 3.3 shows the NAR vs density plots for three sample subjects and for the two hemispheric modes (right hemisphere, left hemisphere). Likewise, figure 3.4 shows the NAR vs density plots for three sample subjects in the inter-hemispheric mode. Note that as noted in section 3.2.3 the adjacent ROIs are excluded in minimizing the NAR, hence the much smaller density range in the intra-hemispheric modes of figure 3.3 compared to the density range in figure 3.4. We found that NAR is consistently minimized (at the value of approximately 0.08) at a density close to 3.2% for inter-hemispheric and 0.4% for intra-hemispheric modes, regardless of the subject. Including the adjacent pairs to the network, the intra-hemispheric densities also reach to 3.2%. This is the density that is also reported by deterministic tractography algorithms when mapping brain connections [92]. The second column of figure 3.3 depicts the relation between the density and threshold. To reach to the optimum density of around 3.2%, MANIA used thresholds ranging from 35% to 60% across the subjects. The distribution of densities, minimum NAR and MANIA thresholds are all plotted in figure 3.5.

The optimum threshold returned by MANIA (approximately 45%) is neither very strict nor very relaxed, and it is located in the plateau of the density-threshold curve (around 2250 streamlines out of 5000 are required for a connection), as seen in figure 3.3 and figure 3.4. Note that to reach to larger connectome densities (such as 10%, which is frequently set as the arbitrary connectome density in the literature [93]), one needs to apply considerably smaller thresholds (approximately 2% in our dataset), which consequently increases the

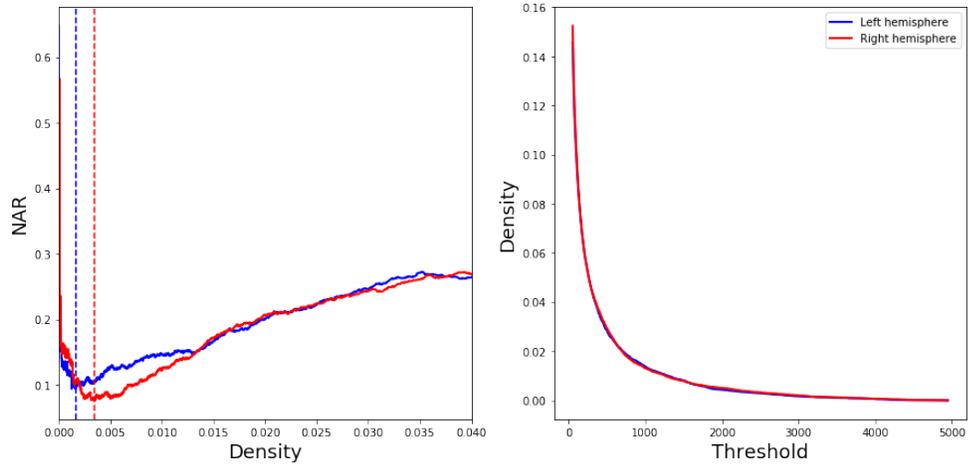
false positive connections of the connectome. Even thresholds close to zero will result in a network not denser than 20%, since there are no streamlines running between many of the source-target pairs (see the asymptotic behaviour of the density in the second column of the figures).

The distribution of the NAR, threshold and density values are shown in figure 3.5. Note that densities in this figure is considered for whole network including all pairs adjacent or non-adjacent.

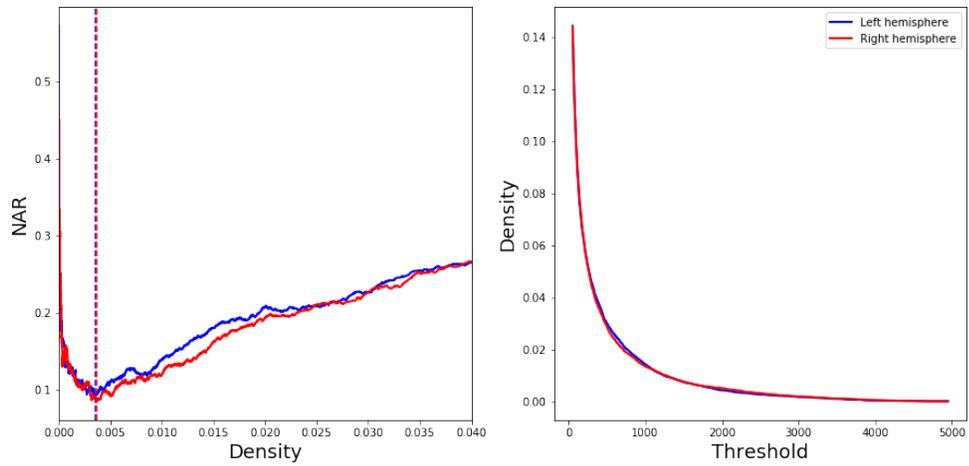
Figure 3.6 shows the axonal distance and connection probability scatter plot for a sample subject. The x-axis value is the average length of the connected streamlines between the corresponding two ROIs, and the y-axis value would be the corresponding fraction of connected streamlines. The axonal distance and connection probability relationship is depicted for left, right and inter-hemispheric connections separately in figure 3.6. The density of the networks are 3.45%,3.33% and 2.98% for left, and right hemisphere and for inter-hemispheric respectively. The axonal distance (the average streamline length connecting two ROIs) for detected ROI pairs varies between 0 to 100mm. The median axonal distance in all these three networks is approximately 19mm .Although, connection probability shows a decreasing trend by axonal distance, as evident in the figure, for each distance the connection probability still varies significantly and axonal distance could not determine the connection strength.

3.2.5 7T validation

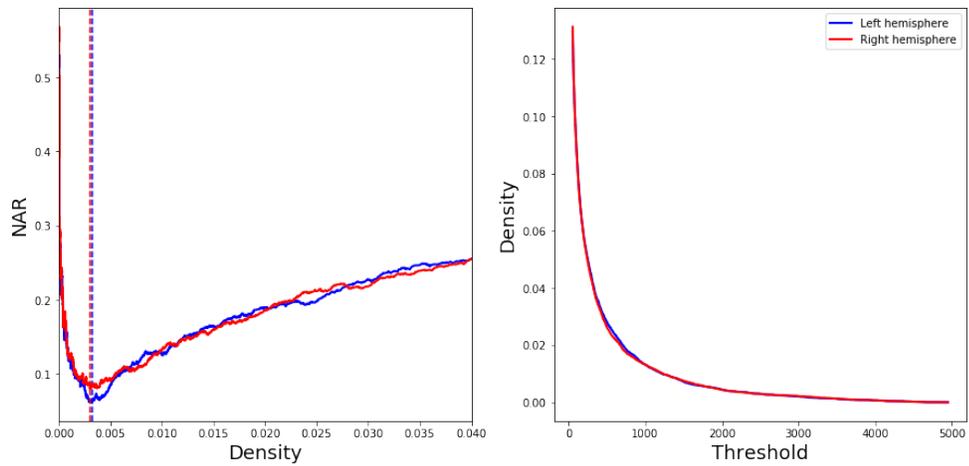
Validation of MANIA accuracy using high resolution 7T MRI scans The absence of a ground truth limits our ability to assess the accuracy and precision of connectomes built using tractography. Hence, the accuracy of MANIA was quantified using the phantom and synthetic dataset discussed in [28]. In this section, we further quantify the accuracy of MANIA by exploiting the enhanced signal-to-noise ratio (SNR), contrast, and resolution attainable through the use of data derived from high-field 7T MRI as part of the HCP [95].



(a) Subject 126426

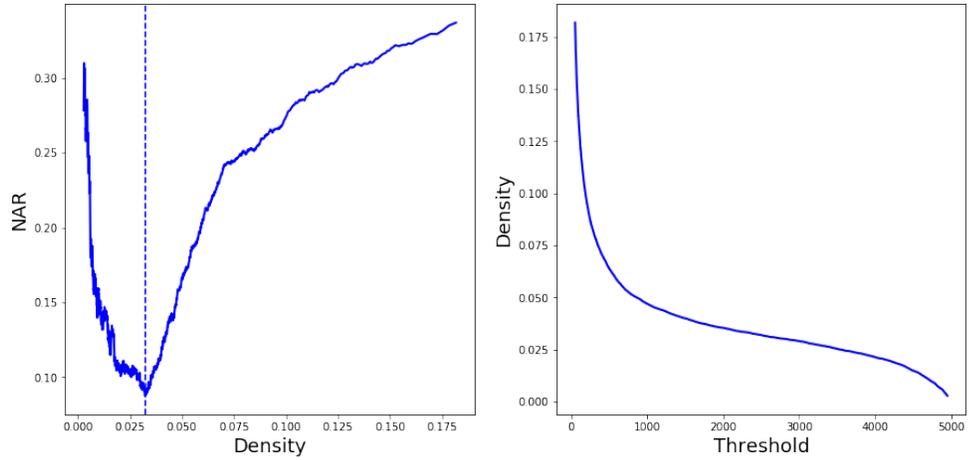


(b) Subject 146735

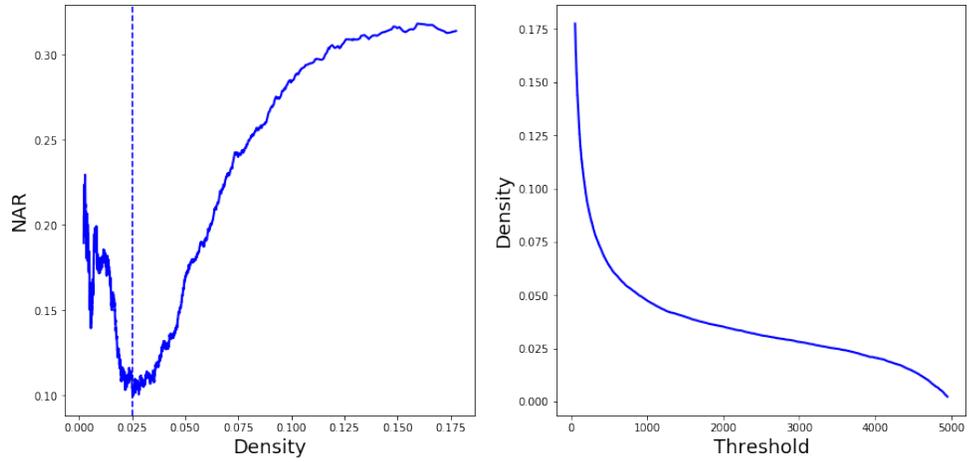


(c) Subject 135124

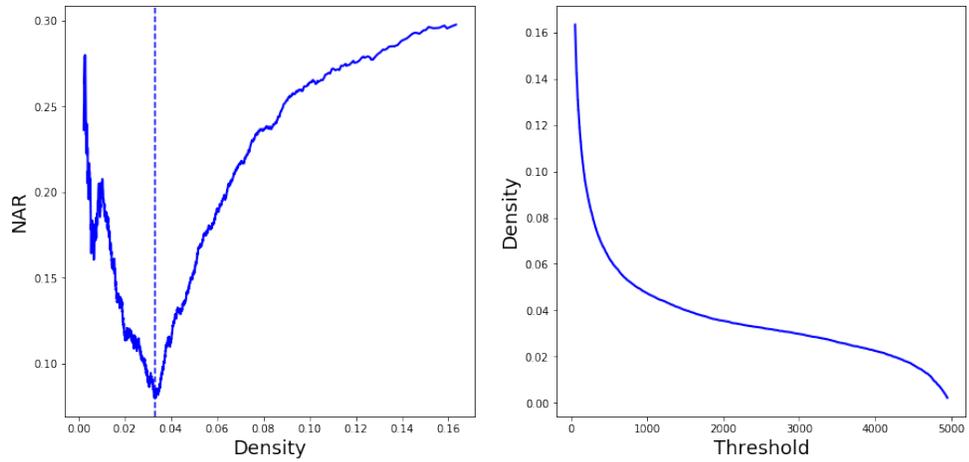
Figure 3.3: The results of applying MANIA on three sample subjects for right and left hemisphere network inference - the dashed lines show the densities chosen by MANIA for the corresponding networks. The density is calculated excluding the adjacent pairs.



(a) Subject 126426



(b) Subject 146735



(c) Subject 135124

Figure 3.4: The results of applying MANIA on three sample subjects for inter-hemispheric network inference - the dashed lines show the densities chosen by MANIA for the corresponding networks

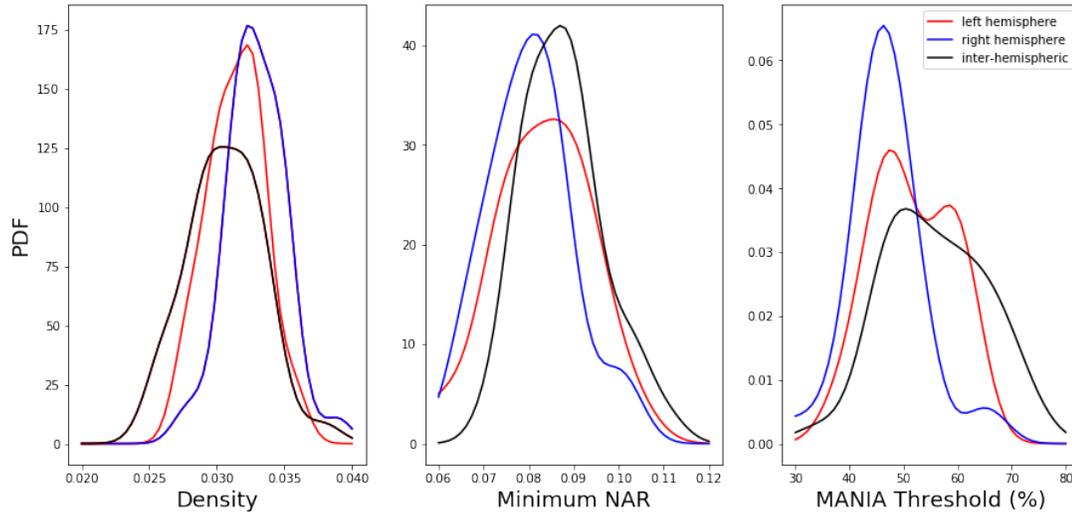


Figure 3.5: The distributions of MANIA results across 37 subjects - The distributions are calculated by kernel-density estimate using Gaussian kernels [94]

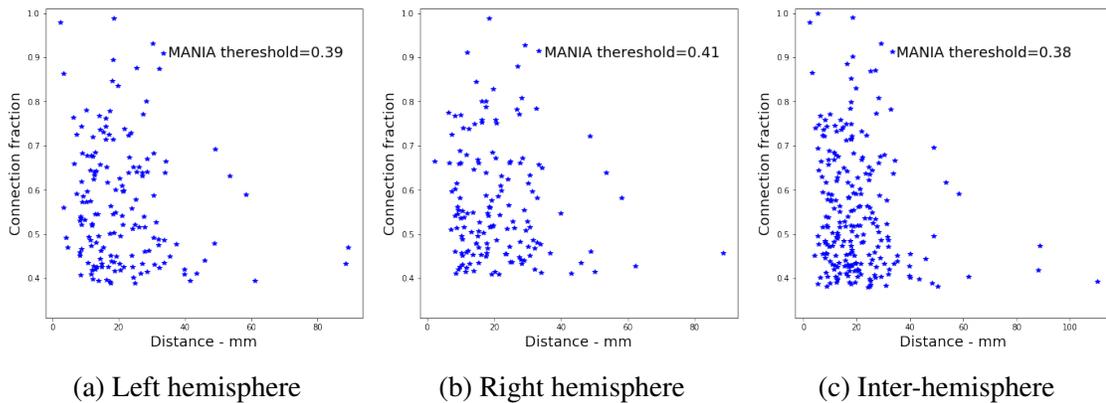


Figure 3.6: Connection probability vs axonal distance for subject 146735 in connected and none-adjacent ROIs - The graphs are presented for left, right and inter-hemispheric networks separately and MANIA thresholds are shown in the graphs. In the left hemisphere 12% of the connections are non-adjacent and 9% of connections are none-adjacent in the right hemisphere. Note that the probtrackx was run to include only streamlines extending at least 20mm (distthresh=20).

Results were obtained from three of our HCP subjects for whom both 7T and 3T MRI data were available.

The high resolution and better SNR of 7T data leads to an amplification of the connection probability for true connections and to an attenuation on the opposite case. Figure 3.7 illustrates two sample cases: one for attenuation and one for amplification in a subject. Note that the connection probability is shown in log-scale for visual purposes; thus, the amplification and attenuation in the figure span an order of magnitude.

To build a proxy for ground truth for our benchmark in each of the three subject, we applied very strict thresholds on the 7T probtrackx dataset of that subject in order to tag if an edge is or isn't present. We categorized each edge as a true positive if the connection probability in probtrackx exceeded $\delta_1 = 0.4$. Likewise, an edge was categorized as a true negative if the probability was lower than $\delta_2 = 0.01$. We then find the trivial true negative connections as those with a probability less than 0.001, which we excluded from our true negative connection set. Figure 3.8 shows the size of 7T benchmark true negatives and true positives as functions of our control parameters δ_1 and δ_2 for a sample subject 146735. Note that we have been strict in labeling connections as true positives or negatives by requiring large δ_1 and small δ_2 . This gives a sample set to measure the accuracy of MANIA that we can have a high level of confidence in. At the same time, we do not want to have easy to detect positive or negative connections, because they would artificially inflate the accuracy of any network inference method. Thus, we excluded ROI pairs in the 7T data for which the connection probability is less than 0.001 as mentioned before.

Finally, we ran MANIA on the 3T data from the same subject and quantified the accuracy against its *7T benchmark*¹. The top row of figure 3.9 shows the true positive rate and false negative rate for a threshold estimator across different density thresholds for sample subject 146735. The optimal density returned by MANIA is marked by the red vertical line. Interestingly, MANIA operates at the point where the true positive rate begins to satu-

¹Of course, this is only an indirect method of constructing a ground truth.

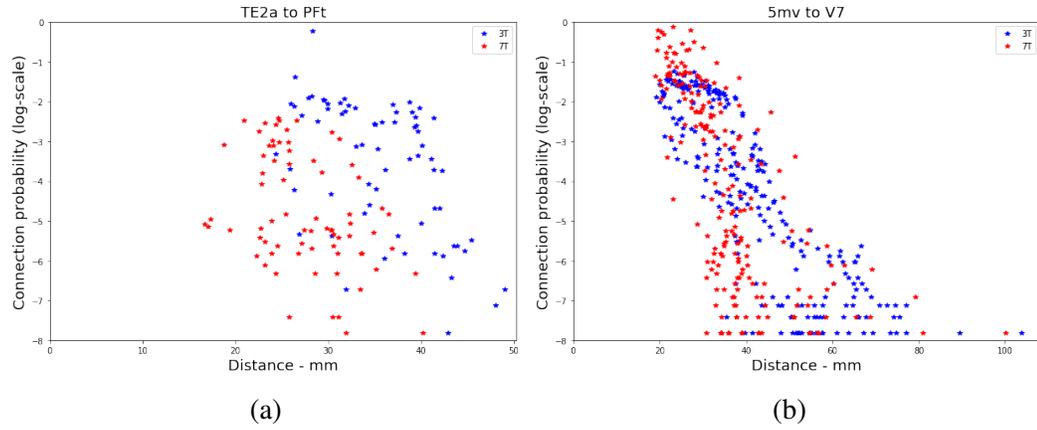


Figure 3.7: Sample source to target connection probability (*log – scale*) vs distance for subject 14673 both in 3T and 7T - a) An ROI pair in which 7T signal is attenuated compared to 3T b) An ROI pair with the opposite effect (amplification in 7T compared to 3T).

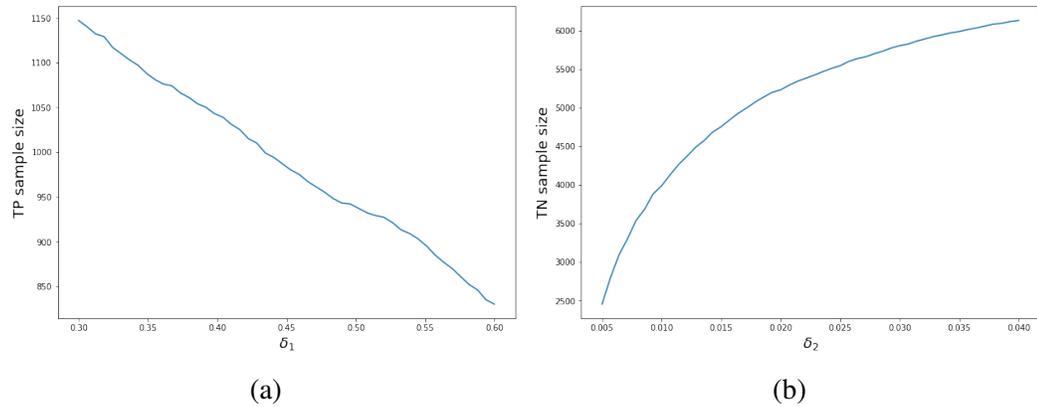
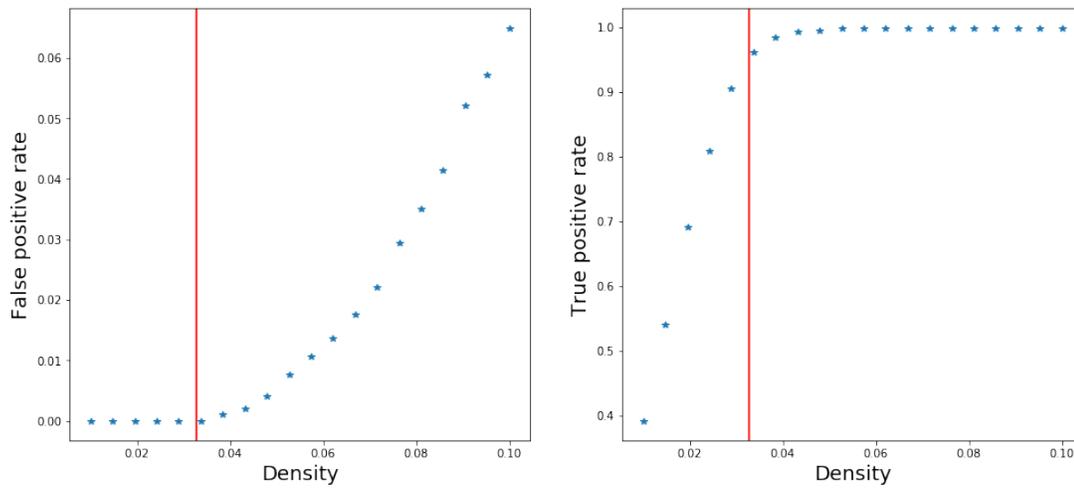


Figure 3.8: True positive and true negative sizes in subject 146735 of our benchmark dataset built using 7T data for different δ_1 and δ_2 parameters

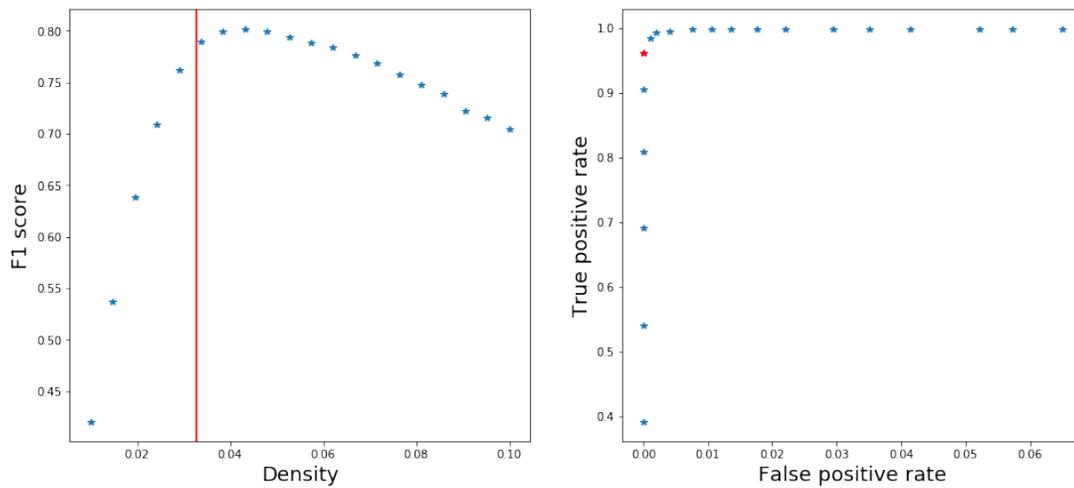
rate, and where the false positive rate begins to grow exponentially, i.e., any increase in the density would only slightly increase the true positive rate while introducing a large false positive error. There is thus a trade off between a small sensitivity increase and a large decrease in specificity.

To quantify this trade-off, we employed the F-measure [96] for the reconstructed connectivity matrices in order to assess their accuracy with a single index. The bottom left plot of figure 3.9 shows the F-measure for different network densities. MANIA manages to find a density close to the maximum F-measure. The same trade-off between sensitivity and specificity is also plotted as an ROC curve in the last plot of figure 3.9. In this plot, the operating point of MANIA is denoted by the red star. One can see that MANIA operates close to the knee of the curve, where any change in the operating point (estimated density) leads to a considerable reduction in accuracy.

All previous results were for fixed variables δ_1 and δ_2 and were shown for a sample subject 146735. As opposed to the benchmark size for true negatives, the true positive benchmark size was an order of magnitude smaller (See figure 3.8). Therefore, we ran a test to measure the robustness of our results for different values of δ_1 and for two more subjects. Moreover, instead of using maximally connected voxels between the source-target ROI-pair, we ran the same test requiring multiple number of voxels nv linking the source and the target. For example, $\delta_1 = 0.5$ and $nv = 3$ means that three voxels must be connected with a probability larger than 0.5 for a connection to exist in our benchmark true positive sample set. Figure 3.10 shows the F-measure graphs for different values of δ_1 and nv and for three subjects. As seen in the figure, MANIA is consistently operating close to the maximum F-measure, and shows a remarkable robustness to the underlying benchmark construction parameters. Thus, all our results hold for different types of benchmark dataset constructions derived from the 7T data and are valid across the subjects.



(a)



(b)

Figure 3.9: Measuring the accuracy of MANIA with $\delta_1 = 0.4$, $\delta_2 = 0.01$ Top) False positive rate and true positive rate across different density thresholds - MANIA operating points are marked by red lines Bottom) F-measure and ROC for different density thresholds - MANIA operating points are marked by red line and red star

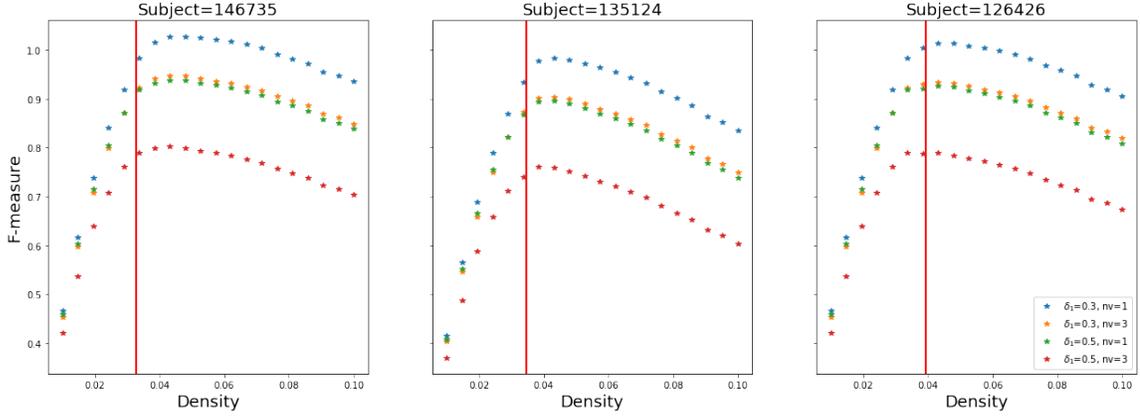


Figure 3.10: The F-measure across several schemes to build the benchmark dataset out of 7T dataset - the study is done for three subjects marked in the title of each figure.

3.2.6 Hemispheric asymmetries

As explained in section 3.2.3, MANIA was applied independently to left and right hemisphere of the 37 subjects in our HCP study. In this section, we focus on the connections consistently present in one hemisphere and not present in the other. In order to isolate these connections, we marked the connections showing 75% difference in their right and left hemisphere presence, i.e. for a connection C we marked it as hemispheric asymmetry if $|P_L(C) - P_R(C)| > 27$ where $P_R(C)$ is the number of subjects having C in their right hemisphere and $P_L(C)$ is the number of subjects having C in their left hemisphere. Table 3.3 lists all these connections. As seen in the table, there are seven such connections, two of which appear mostly in the left hemisphere while the rest of them are present mostly in the right hemisphere.

Four of the ROIs appearing in table 3.3 are located in the auditory cortex (see table 3.4). There are converging evidence from activation, connectivity, and stimulation in the literature pointing to the fact that the auditory cortex is a lateralized network [97]. Also, a47, a10p and 9-46d show network asymmetries. Interestingly, all these areas are known to have topological asymmetries reported in [23] and hinted in table 3.4.

Table 3.3: Connections with hemispheric lateralization. The first two connection are mostly present in left hemisphere and the rest of the connection are mostly present in right hemisphere while absent in the opposite hemisphere

ROI1	ROI2	# Appearances left hem.	# Appearances right hem.
MBelt	PI	35	6
9-46d	a10p	30	0
TA2	STSda	1	31
PGp	IP1	4	36
PGp	PGi	1	36
p10p	a47r	5	37
STSva	TE1m	0	28

Table 3.4: Meta information about ROIs in table 3.3

ROI	Longer Label	Location hints	Commentary
9-46d		Central portion of the dorsolateral prefrontal cortex	Relative to its supero-medial neighbor 9-46d in the left hemisphere, area 46 has less myelin (Panel B), differs modestly in functional connectivity
STSva (area 176)	Superior Temporal Sulcus ventral-anterior	Auditory Association Cortex	
STSda		Auditory Association Cortex	
MBelt	Medial Belt	Early Auditory Cortex	
PGp		Inferior Parietal Cortex	
a10p	Area anterior 10p	The orbital and polar frontal cortex	area a10p adjoins area a9-46v only in the left hemisphere.
IP1		Inferior Parietal Cortex lateral bank of intraparietal sulcus	
a47r	Area 47 r	The orbital and polar frontal cortex	The topology of area a47rs superior border differs in the left and right hem.
PGi	Area PGi	Inferior Parietal Cortex	
TA2		Auditory Association Cortex	is activated vs deactivated in the LANGAUGE MATH
PI	ParaInsular cortex	Insular and Frontal Opercular Cortex	

3.2.7 Discussion on HCP study - case study B

How dense is the human cortical connectome? The density of the human cortical connectome, the most fundamental network property, is still elusive to us at different scales. In chapter 2, we introduced MANIA as a mean to address the problem of thresholding the probabilistic tractography results in order to infer the network density for the connectome using probabilistic tractography. In the second case study of this chapter, we used MANIA to build human cortical connectome from data published by the Human Connectome Project, HCP [15]. We concluded that with the current state-of-the-art MRI acquisitions and with the control to avoid the high false positive rate of probabilistic tractography [96], the cortical connectome density is approximately 3.2%. However, note that this is not a claim on the actual density of the connectome, since there are several inherent limitations to diffusion MRI and tractography which need to be addressed prior to draw conclusions about the ultimate representative connectome.

There are multiple steps to infer brain networks from dMRI data, and each introduces parameters affecting the network density and indices, such as centrality and the clustering coefficient [98, 99]. Some of the uncertainties stem from the limited resolution of MRI scans, making it impossible to model the fibers crossing and kissing each other along the white matter pathways. However, there are uncertainties in network inference that are inherent to tractography itself [100], and more prominent in long-range connections [101]. For example, authors in [102] show that most tractography algorithms cannot recover more than one-third of the volumetric extent of well-known long-range connections. Thus the inherent limitation of tractography when dealing with longer connections, distance bias, needs conceptual development and methodological creativity.

Also, connection probabilities reported by probabilistic tractography exhibit exponential decay with inter-regional distance, even in the case of leveraging high resolution 7T MRI scans from HCP. This is in agreement with tracer studies reporting also exponential decay but for biological connection strengths [103, 104]. The extent to which the expo-

ponential decay of tractography is biologically specific or stems from the methodological artifacts (distance bias) remains to be explored [105]. A correction framework to decompose the decay exponent to the distance bias factor and the true biological factor and consequently regress out the distance bias systematically from tractography results will generate a remarkable increase in the accuracy of individual connectomes. In the presence of such a method, MANIA, or any network inference algorithm, can work on corrected data such that distance bias does not affect the final connectome.

The FDT toolbox provides a distance bias correction option by multiplying the number of streamlines that cross a voxel by the average length of those streamlines. However, there is no evidence that this simple form of distance bias correction is able to significantly improve the accuracy of the network inference. Moreover, the results do not have probability scale or meaningful span. This multiplication strategy is not meant to correct distance bias, and has only been proven to increase the accuracy of clustering analysis; hence it is recommended for studies related to brain parcellation. Authors in [106] explored distance correction using different strategies that seek to multiply the distance matrix and probabilistic tractography inferred strengths. The results lead to only a small notable increase in the overall accuracy of the identified connections, since the multiplication compromises the false positive rate in order to enhance the detection of true connections. On the other hand, the study conducted by [107] demonstrates that FDT distance correction strategies worsen the results because they compare the corrected connectome to a ground truth² connectome from tracer studies.

Distance bias, although the most important issue at hand, must be considered in conjunction with gyral bias. As authors in [89] report, “there is a bias for fiber tracking algorithms to terminate preferentially on gyral crowns, rather than the banks of sulci”. This bias will corrupt tractography algorithms to a significant extent if the seeding happens on white matter and also if volumetric analysis is used instead of a surface-based method. The

²Although tracer studies are believed to report the best in vivo connectomes, the *ground truth* is still a frequent misnomer, since tracer studies are not fully devoid of false positives and negatives.

tractography seed points in all our case studies were on gray-matter, or at its boundary to white matter. This reduces the “gyral bias” effect in our studies. Also, we moved from volumetric to surface based analysis in this chapter to account for the highly convoluted nature of the human cortex, which has an added benefit of also decreasing the gyral bias [108].

Despite the aforementioned inherent limitations of tractography, we showed that the reconstruction of highly accurate and consistent connectomes at the macro-scale is possible, albeit with sparse cortico-cortical graphs.

Finally, we should mention despite considerably denser networks for other animals reported in the literature of invasive studies, e.g., the extremely dense cortico-cortical connectome of mice (97%) [103], the human connectome must still await for such reports due to the noninvasive nature of human studies.

CHAPTER 4

MULTI-SENSORY INTEGRATION IN THE MOUSE CORTICAL CONNECTOME USING A NETWORK DIFFUSION MODEL

4.1 Introduction

Perception requires the integration of multiple sensory inputs across distributed areas throughout the brain [109]. While sensory integration at the behavioral level has been extensively studied [110], the network and system-level mechanisms underlying *Multi-Sensory Integration (MSI)* are still not well understood, especially in terms of the role that cortex plays. The traditional view of primary sensory areas as processing a single modality is rapidly shifting towards a view of the cortex as highly integrated and multi-sensory [111]. The somatosensory, visual, auditory, gustatory and other sensory streams, come together (integrate) and separate (diverge) to be processed in different parts of the cortex. The neural basis of how these sensory streams are processed and how they generate a coherent perceptual state remains elusive [112]. It is likely that this state is created and regulated by multiple structures distributed throughout the cortex working together in concert [113].

To understand the architectural principles that enable MSI, we need data and models that span the entire brain focusing not on individual neurons, regions or even circuits, but on distributed networks. The connectome is thus a potentially powerful tool for studying MSI. However, it would not be enough to just know how different brain regions are connected anatomically. Rather, we need models that combine structure (connectomics) with function [114] to address the question of which connections and paths are activated by different sensory modalities. Having both a structural network and a functional model in hand, we can begin to tackle the problem of discovering the networks that support and constrain MSI.

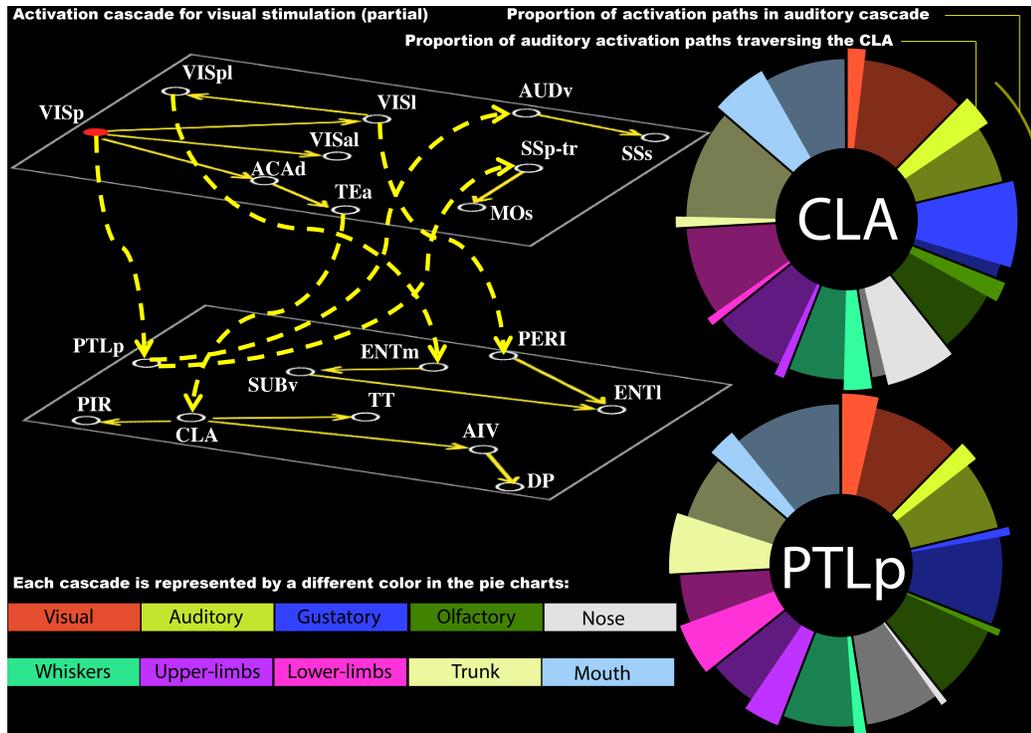


Figure 4.1: *High-level illustration of approach and main result: **Top-Left:** An illustration of a single sensory cascade, originating at the primary visual cortex (VISp). The yellow edges show anatomical connections that participate in the cascade. The ROIs at the upper layer reside on the cortical surface while the ROIs at the lower layer are deeper in the brain. The edges between the two layers are dashed. To simplify the visualization, we only include 20 ROIs at this cascade (the complete cascade includes 67 ROIs). The cascade forms a Directed Acyclic Graph (DAG), and it is produced by applying the Asynchronous Linear Threshold (ALT) model on the cerebral cortex portion of the mouse connectome – starting the cascade at VISp. **Bottom and Right:** The two most important core ROIs at the hourglass waist – Claustrum (CLA) and Posterior Parietal (PTLp) cortex – jointly cover 40% of all activation paths in the ten sensory cascades we consider. Each cascade is represented by a different color. We include two circular disks (pie charts), one for each of these two core ROIs. The proportion of activation paths in each cascade is shown by the corresponding color. The protruded portion of each circular section represents those paths that traverse the corresponding core ROI. For example, the activation paths in the auditory cascade account for about 8.5% of the total number of paths – and about 36% of those paths traverse CLA.*

Here, we adopt a variation of the *Asynchronous Linear Threshold* (ALT) network diffusion model [115] to capture the communication dynamics of networks that contribute to MSI. In particular, we focus on how information propagates throughout the brain, starting from different primary sensory regions (e.g. primary visual cortex, auditory cortex, and different somatosensory regions). The ALT model assumes that a “node” (brain region)¹ becomes active when more than a fraction of the neighboring nodes it receives afferent projections from are active. We use a variation of this model with weighted connections, where the weights are based on the connection density of the projections [115]. The ALT model is simple yet it incorporates information about distances between areas (to model connection delays) and uses local information (a thresholding nonlinearity) to potentially gate the flow of information. Figure 4.1 illustrates the notion of *activation cascade*, showing only a small part of the visual cascade according to the ALT model. With such a model, it is possible to understand how activation cascades propagate in the brain, and then combine them to study the global architecture of MSI.

We apply this model to the Mouse Connectivity Atlas (MCA) provided by the Allen Institute for Brain Science [26]. This connectome has been available since 2014, and it consists of estimates of connection density between cortical as well as subcortical regions, providing access to information about whole brain connectivity across functionally and structurally distinct regions. By coupling the ALT network diffusion model with this representation of the connectome, we can ask questions such as: what is the relative order in which different regions get activated after, say a visual or auditory stimulation? Which are the most central regions for each activation cascade? Are these cascades largely independent of each other, or are there few “bottleneck” regions through which almost all cascades go through? If so, which are these regions and what is their topological role in each cascade?

To examine the accuracy of the ALT model, we use Voltage Sensitive Dye (VSD) imag-

¹We use the terms “node”, “region” and ROI (“region of interest”) interchangeably.

ing to capture how the activity that is triggered from uni-sensory stimulation propagates throughout the cortex [27]. By comparing our modeling results to these functional datasets, we find that the ALT model predicts correctly the temporal ordering of node activations in a cascade – as long as the VSD data provides sufficient temporal resolution to permit activation time comparisons. This suggests, perhaps surprisingly, that despite the simplicity of linear threshold models, they can provide a useful approach for studying communication dynamics in brain networks.

We then aggregate the uni-sensory activation cascades predicted by the ALT model to investigate the architecture of MSI. To do so, we consider the number of activation paths that traverse each node across all of the different uni-sensory cascades inferred by the ALT model. We find that a small set of brain regions (around ten), form a “bottleneck” through which almost all such activation paths traverse. The *Claustrum (CLA)*, despite its small size, is the most central bottleneck region in the flow of sensory information from primary sources towards higher-level brain regions [116]. The *Posterior Parietal (PTLp)* cortex is the second bottleneck node, covering almost as many activation paths as CLA. Figure 4.1 visualizes the contribution of these two core regions in each of the ten sensory cascades we consider. About 40% of all activation paths, across all sensory cascades, traverse these two regions, suggesting that they play a prominent role in MSI.

Our results support the presence of a *bow-tie or hourglass architecture* in the architecture of MSI. The salient qualitative feature of an hourglass architecture is that a small number of nodes (at the waist of the hourglass) can cover almost all source-target paths [117]. In the context of MSI, this means that the multi-modal sensory input is first integrated through a small set of brain regions at the *waist* of the MSI hourglass. Then, those intermediate-level representations diverge to several higher-level association regions, providing integrated multi-sensory information for more complex cognitive tasks. Importantly, this result would not be revealed through static network analysis metrics and methods (such as betweenness centrality or rich-club), suggesting that the dynamic perspective offered by

the ALT model provides valuable insights into MSI.

4.2 Methods and Data

The ALT model focuses on how stimulation of one cortical region propagates to the rest of the brain when constrained by the underlying connectome. An appropriate metaphor could be to think of the cortex as a complex consisting of several ponds that are connected through creeks of different lengths and widths (the connectome), while a stimulus corresponds to a large rock falling in one of the ponds. ALT aims to model how that initial perturbation in the pond ripples through the connected ponds (the activation cascade). In the rest of this section, we describe the main components of our approach: the mouse connectome, the ALT diffusion model, the hourglass network analysis framework, and the VSD-based model validation.

4.2.1 Structural network

The structural (i.e., anatomical) network we analyze is a subset of the Allen Mouse Brain Connectivity Atlas [26], which is based on tracking axonal projections labeled by viral tracers. It consists of 213 ROIs that cover the entire brain – all tracer injections however were performed at the right hemisphere [26]. This means that the contralateral connections from the left hemisphere to the right and the ipsilateral connections at the left hemisphere are not mapped. For this reason we only analyze the right hemisphere connections.

The strength of the connection from a source ROI to a target ROI is quantified by a metric that Oh et al. refer to as *connection density* [26]. This metric is roughly proportional to the average number of axons projecting to a target ROI neuron from the source ROI. This metric can be thought of as the total number of axons from the source ROI to the target ROI, normalized by the size of the target ROI.

Each connection is associated with a p-value that quantifies the statistical confidence

that that connection exists [26]. We filter out connections with p-value higher than 0.05.² The reason we do not filter connections based on their weight is because there are many weak but statistically significant connections, as shown in Figure B-1 – it is well known that weak connections can play an important role in network diffusion phenomena as long as there are many of them [118].

The physical length of the connections is approximated based on the Euclidean distance between the two corresponding ROIs’ centroids. This is only an approximation but it is reasonably accurate because the mouse cortical surface is small and smooth with little folding [119, 120]. Additionally, as shown in Section 4.3.6, our results are robust to the selection of the connection weights or physical lengths.

The network we consider consists of the 67 ROIs that cover the four components of the cerebral cortex: *isocortex*, *olfactory areas*, *hippocampal formation*, and *cortical subplate*. We do not include subcortical ROIs that reside in the cerebellar cortex, cerebellar nuclei, striatum, medulla, pallidum, midbrain, pons, thalamus and hypothalamus for the following reasons. First, we are mostly interested in how sensory information that reaches the primary sensory regions of the cortex propagates throughout the rest of the cortex [121], rather than on the role of subcortical structures such as the thalamus or the superior colliculus. Most of the MSI literature focuses on the how those subcortical structures modulate and “route” sensory information in different parts of the cortex [122] – this type of processing however cannot be modeled without further sub-divisions of the thalamus and other subcortical regions and without at least some cell-type specificity. Second, modeling the interactions between cortical and sub-cortical regions, especially in the context of MSI, would require more elaborate communication models that can capture feedback mechanisms. Third, the accuracy of the inferred subcortical projections in the Allen mouse connectome is not expected to be as high as that of cortical projections.

The final network, denoted as N_c , consists of 617 directed edges between 67 nodes –

²We have repeated the analysis for other p-values in the range 0.01-0.1 – see Section B-2.

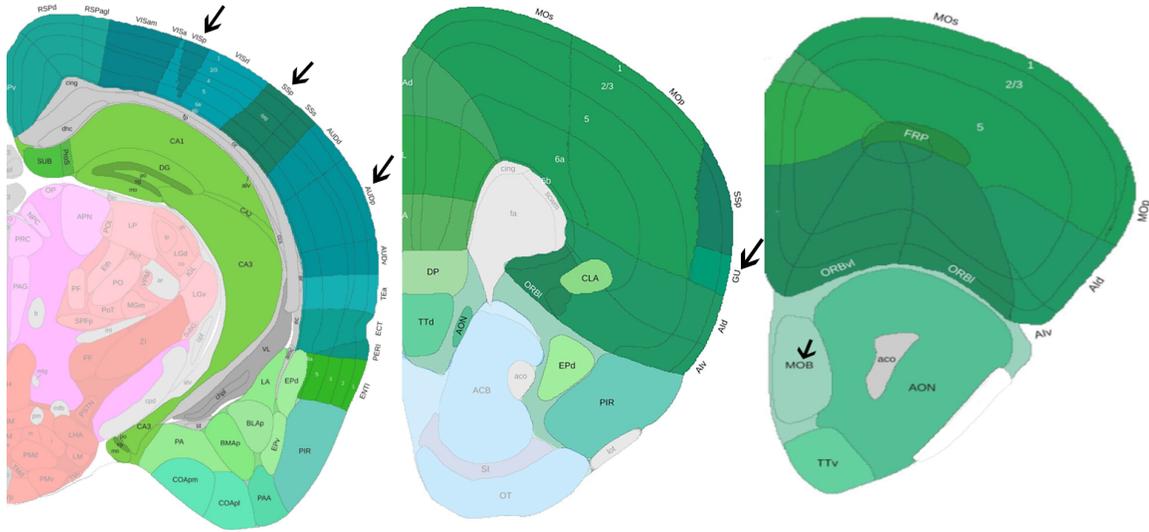


Figure 4.2: *The location of the ten primary sensory regions:* Three coronal slices from the Allen Mouse Brain Atlas with the cerebral cortex regions tinted by green and the source regions identified by an arrow. The somatosensory region includes six different sub-regions for lower limbs, upper limbs, trunk, mouth, nose and whiskers. The remaining four sensory sources are: visual (VISp), auditory (AUDp), gustatory (GU) and olfactory (MOB).

each node corresponds to one of the ROIs in our model. The density of N_c is 13.9%. The distribution of edge weights is skewed, with 80% of the edges having a weight of less than 5 and few edges having a weight of up to 40. The distribution of edge lengths is almost uniform in the 1-7mm range. The diameter of the network (maximum shortest path length between any two nodes) is 7 hops, while a node is on the average about 4 hops away from any other node. The average in-degree of each node is 9.2 connections ($\sigma=3.1$), while the out-degree distribution has the same mean but larger variability ($\sigma=8.3$). Additionally, the network N_c is strongly clustered, with an average clustering coefficient of 60% [123].

The ten primary sensory regions associated with the visual, auditory, gustatory, olfactory systems, as well as six somatosensory regions for different body parts (see Table B-1), have a special role in our analysis: they are viewed as *sources* of sensory information in the cortex [124, 125]. The location of these ROIs in the Allen Mouse Brain Atlas is shown in Figure 4.2.

4.2.2 The ALT model and activation cascades

The connectome is a structural network and so it constraints, but it does not determine by itself, the paths through which information flows in the brain. To study that flow, we also need to model how dynamic brain activity propagates on the connectome.

We choose a simple network diffusion linear threshold model [126], mostly because it involves a single parameter – more realistic neural mass models, such as Wilson-Cowan, depend on several parameters [127]. The model assumes that nodes are either *inactive* (state=0) or *active* (state=1). In the model's simplest form, a node i becomes active when more than a fraction θ of its neighbors become *active*. Here, we deploy a variation for directed and weighted networks in which each edge is associated with a communication delay and a weight, referred to as *Asynchronous Linear Threshold (ALT)* model.

Specifically, the state of node n_i is represented by $s_i(t)$, the neighbors of n_i with an incoming edge to n_i are denoted by $N_{in}(n_i)$, the communication delay from node $n_j \in N_{in}(n_i)$ is t_{ji} , while the weight of that connection is w_{ji} . Initially, the state of every node is set to 0, except the *source* node of the activation cascade, which is set to active at time $t = 0$. The state of each node n_i is then updated asynchronously based on the state of its neighbors as follows:

$$s_i(t) = 1 \quad \text{if} \quad \sum_{j \in N_{in}(i)} w_{ji} s_j(t - t_{ji}) > \theta, \quad (4.1)$$

where θ represents the activation threshold.

Figure 4.3-a illustrates the ALT model with a toy example, where we can think of each node as a cortical ROI. The (directed) edges between ROIs represent the structural connections between ROIs. In this example, the activation threshold θ is set to 1. For this value, the cascade covers the entire network node within 7 time steps. If θ was larger than 1, the cascade would not take place – the only active node would be n_1 . As we will see in the next section, this sharp transition between not having a cascade and a complete cascade

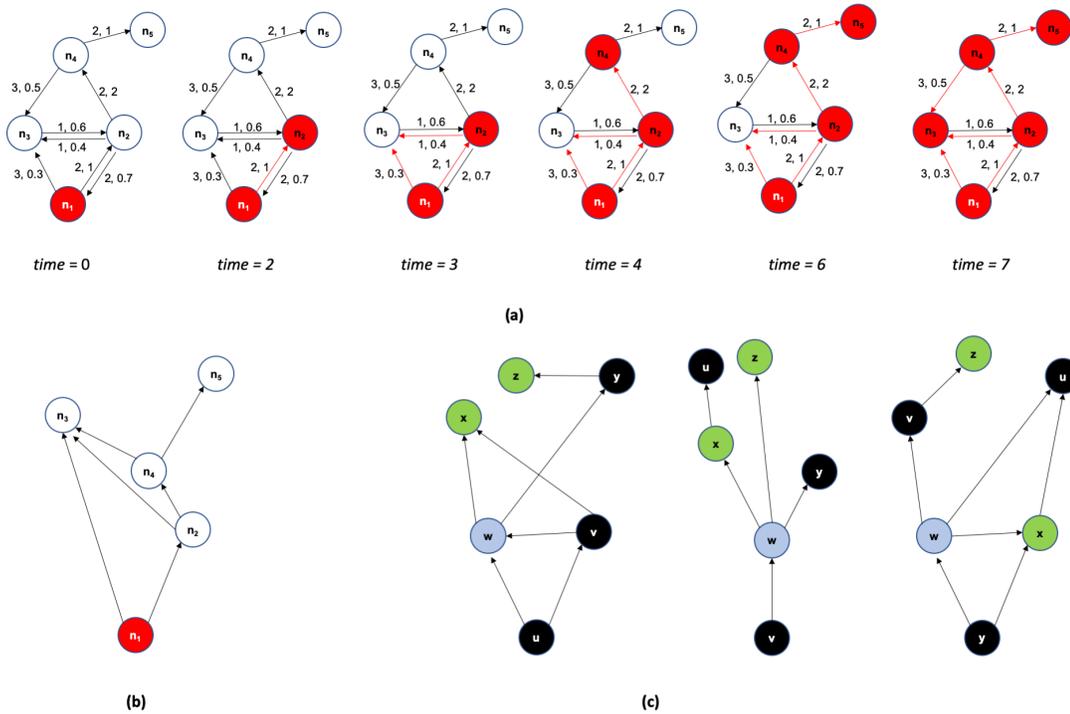


Figure 4.3: *Illustration of ALT model and τ -core analysis:* **(a)** A toy example of a 5-node network on which we run the ALT model. Each edge is marked with a communication delay, followed by a weight. The activation threshold is $\theta=1$. The black edges represent the underlying structural network while the red unidirectional edges represent the activation cascade as it unfolds over time. **(b)** The activation cascade directed acyclic graph for the previous toy example. The source of the cascade is n_1 . **(c)** A toy example with three activation cascades (the sources are nodes u , v and y). The total number of source-target paths is 11 (4 at the left, 3 at the middle, and 4 at the right). Node w has the highest path centrality ($P(w)=9/11$). If $\tau \leq 9/11$, the τ -core consists of only that node.

as θ decreases also occurs in the mouse connectome.

An activation cascade also reveals the node(s) that contribute in a causal manner in the activation of a node. For example, the activation of n_1 and of n_2 in the previous example is not sufficient to activate n_3 ; the latter is activated only when n_1 , n_2 and n_4 are all active. Suppose that t_i^a denotes the time at which n_i becomes active according to the ALT model. We say that node n_j contributes in the activation of n_i (denoted by $n_j \rightarrow n_i$) if $n_j \in N_{in}(n_i)$ and $t_i^a \geq t_j^a + t_{ji}$. In other words, the nodes j that contribute to the activation of n_i have a connection to n_i and they should be active at least t_{ji} time units before the activation of n_i . The set of relations $n_j \rightarrow n_i$ form a Directed Acyclic Graph (DAG) with a single source node that covers all nodes that participate in the cascade. The DAG edges are a subset of the connections in the underlying structural network. The cascade for the previous example is shown in Figure 4.3-b.

4.2.3 Analysis of activation cascades

An activation cascade consists of a collection of *source-target* paths, with each such path originating at the source of the cascade and terminating at a node without any outgoing edges in the cascade. A source-target path is a sequence of ROI activations that propagate in a causal manner from the source node to a target node. For instance, the cascade of Figure 4.3-b includes three source-target paths from n_1 to n_4 because the activation of the latter requires the activation of both n_2 and n_3 , and the activation of n_2 requires the activation of both n_1 and n_3 .

After constructing a cascade for each source, we use network analysis to identify the nodes that play a more central role in the collection of all cascades. The centrality metric we use has the following graph theoretic interpretation: for each node v , the *Path Centrality* $P(v)$ of node v is the fraction of all source-target paths, across all cascades, that traverse v (following the terminology of [117]). Figure 4.3-c illustrates this metric with three small cascades. Nodes with higher path centrality (PC) are expected to be more important be-

cause the activation cascades depend more heavily on them. Note that a source node is traversed by all source-target paths in its own cascade but it may have low path centrality when we consider the collection of all cascades.

The PC metric quantifies the importance of each node in isolation. We are interested, on the other hand, in the smallest set of nodes that can jointly cover almost all source-target paths in the given set of cascades. To answer this question, we adopt the τ -core definition of [117]: *the τ -core is the minimal set of nodes such that the fraction of source-target paths that traverse any node in the set is at least τ* . The τ -core problem is NP-Hard for $\tau < 1$ [117]. It can be approximated with a greedy heuristic in which the node with the highest PC joins the set in each iteration. That node is then removed from all cascades it appears in. The PC of the remaining nodes is recomputed after each iteration. If the τ -core is a small set, relative to the total number of nodes in the network, the nodes of that set can be thought of as the *bottleneck* of the activation cascades (see Figure 4.3-c for an example).

4.2.4 Comparison of modeling results with functional data

To examine whether the ALT model can accurately predict the propagation of sensory stimulation in the cortex, we need an experimental setup in which we can stimulate different sensory modalities of a living animal, and monitor at the same time and in a fine temporal resolution (in the order of a millisecond) the activity of different cortical neural populations (in a spatial resolution of few μm).

Such experiments are possible today, relying on technologies such as calcium imaging or voltage-sensitive dyes (VSD) in conjunction with fluorescence microscopy. Here, we leverage the experimental results of an earlier study to examine the accuracy of the ALT model in the context of whole-cortex imaging in mice under single sensory stimulation [27]. In brief, the experiments include five types of sensory stimulation: visual (flash), auditory (tone), whisker touch, forelimb touch, and hindlimb touch. Each stimulation experiment is repeated ten times and on several different animals (we analyze data for five animals).

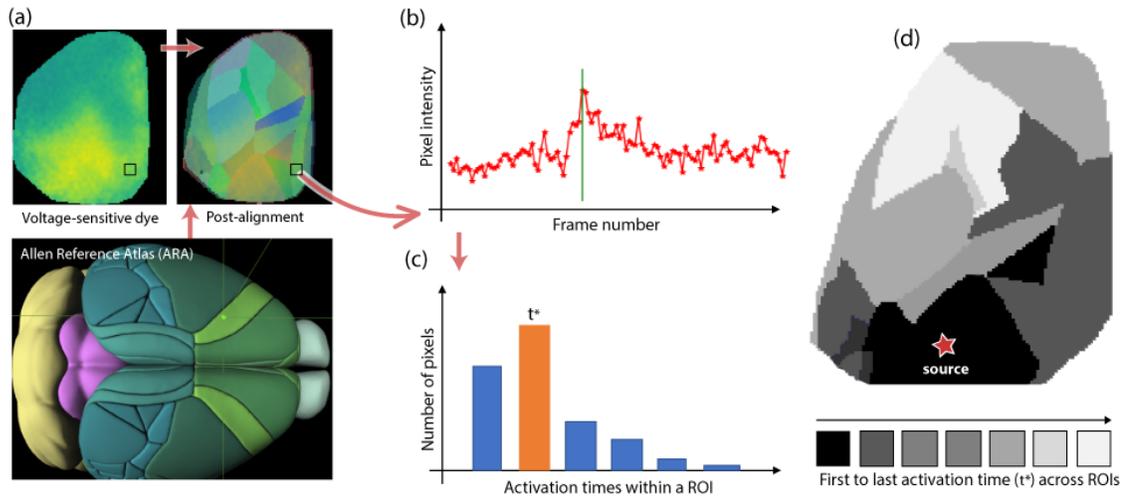


Figure 4.4: *VSD data processing pipeline*: (a) *Lower*: The Allen Reference Atlas (ARA). *Upper-left*: A sample VSD image covering most of the left cortical surface five frames after visual stimulation. *Upper-right*: The ROIs at the left ARA cortical surface mapped to the native cortical surface of an animal. (b) The activation time of a pixel is defined as the frame of maximum post-stimulus VSD signal at that pixel. (c) The activation time of an ROI is defined as the activation time of *most* pixels in that ROI. (d) The output of this pipeline is an activation time for each ROI, depicted here with a grey-scale map (black for the first ROI activation and white for the last).

The recorded images cover almost the entire cortical surface, have a temporal resolution of 6.67msec, and a size of 128×128 pixels at a spatial resolution of $50\mu\text{m}/\text{pixel}$. As it will become clear in Section 4.3.2, the experimental spatial resolution is sufficient for our purposes, given that each network node in the ALT model refers to an entire cortical ROI of the Allen mouse brain atlas. The temporal resolution, however, is marginally sufficient because for about 20-30% of all ROI pairs we cannot tell which ROI gets activated first as they appear to be activated during the same frame.

To compare the VSD-based results with our modeling results, we perform the following steps:

Step 1. Register: Register the native cortical surface of each animal to the Allen mouse brain atlas. This step is performed using an affine transformation that minimizes the least-squares error between the coordinates of the centroid of a primary sensory ROI (e.g., VISp) in the Allen atlas, and the coordinates of the pixel that first gets activated after the corre-

sponding sensory stimulation (visual in this example). We use five primary sensory ROIs to construct this transformation: VISp, SSp-bfd, SSp-ll, SSp-ul and AUD-p.

Step 2. Parcellate into ROIs: Parcellate the native cortical surface into ROIs using the Allen mouse brain atlas and the previous affine transformation. Some cortical ROIs are not visible in the VSD images: FRP, PL, ACAd, VISpl, and VISpor. Also, the MOs region is only partially captured in the VSD images.

Step 3. Estimate Activation Time: Estimate an *activation time* for each ROI in the experimental data. To perform this step, we first identify the maximum post-stimulus activity for each pixel in that ROI – this is defined as the activation time of that pixel. The image frame that corresponds to the activation time for most pixels of that ROI is defined as the activation time of the ROI. This processing pipeline is summarized in Figure 4.4.

The ALT model, on the other hand, models each cortical ROI as a single node, and it assumes that the transition of each node from inactive to active occurs instantaneously. The time unit in the ALT timeline is arbitrary – so we cannot just compare the absolute activation times between the experimental and modeling results. Instead, we examine the consistency of the temporal ordering of activations in the ALT model and in the VSD experiments. Specifically, if X and Y are two ROIs, and X is activated before Y in the modeling results, is it also the case that X is activated before Y in the VSD data? If so, we count that ROI pair as a *temporal agreement*. If X and Y are activated in the same VSD frame, we count that pair as a case of *insufficient temporal resolution*. Finally, if X is activated after Y in the VSD data, we count that ROI pair as *temporal disagreement*.

4.3 Results

4.3.1 Sensory-specific activation cascades in the mouse cortex

The ALT model requires the selection of a single parameter, the activation threshold θ . This threshold controls the *size of the cascade*, i.e., the fraction of network nodes that become active after the activation of a source node. One may expect that as θ decreases

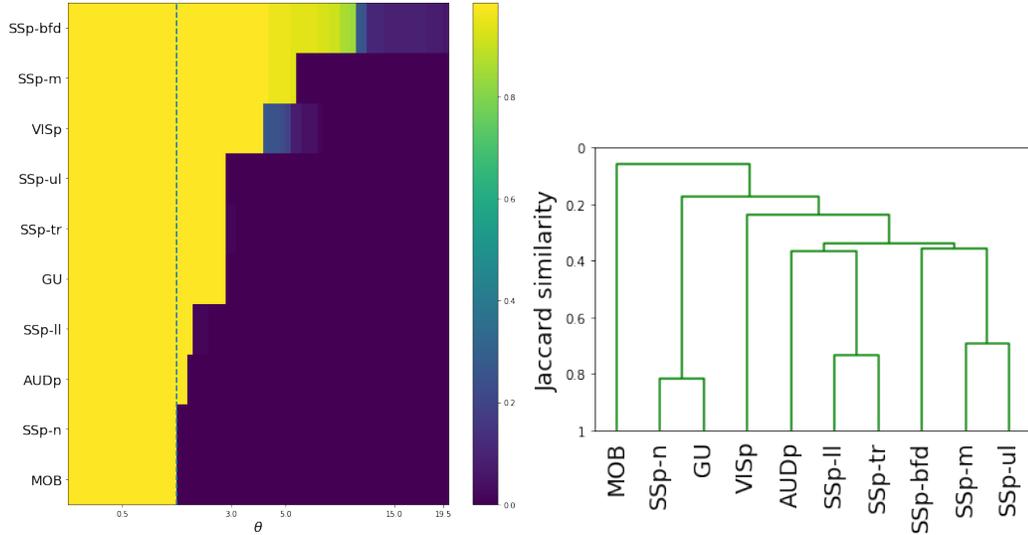


Figure 4.5: *Effect of parameter θ on cascade size, and similarity between the ten cascades:* (a) Each row of the heat map shows the fraction of activated nodes after the stimulation of a single source, for different values of the threshold θ . The selected threshold is marked with the dashed vertical line. (b) Similarity between the ten sensory cascades using the average-linkage hierarchical clustering method.

towards zero, the size of the cascade becomes *gradually* larger. This is not the case however. Figure 4.5-a shows that as θ decreases we observe a rapid transition from the absence of a cascade (where only the source node is active) to a complete cascade, in which all nodes become active. This is true for all source nodes listed in Table B-1.

The reason behind this rapid transition is the highly clustered topology of the mouse connectome. This property is often quantified by the *clustering coefficient* [128], which has an average value of 0.60 in the mouse connectome. This means that if a node v connects to two nodes u and w , there is a probability of 60% that there is an edge from u to w or in the opposite direction, “closing the triangle” formed by the three nodes u, v, w . So, if v is the source of the activation and θ is low enough so that v activates at least one of its neighbors, say u , it is highly likely that other neighbors of v receive input from u as well, increasing the chances that they will also get activated. The same argument applies to all other pairs of activated nodes – not only the source and its direct neighbors.

We choose θ so that the ALT model produces a *complete cascade*, for every source we

consider [129]. This choice is motivated by experimental results [27]: at least some activity is detected in all cortical regions after sensory stimulation (visual, auditory, touch) in anesthetized mice. It may appear counter-intuitive that visual stimulation, for example, can impact activity in regions associated with different sensory modalities (e.g., gustatory) but such interactions are possible through the many feedback connections in the connectome, and they are consistent with several prior studies which argue that there are no strictly uni-sensory regions, and that to some extent the entire cortex is a multi-sensory organ [130, 131, 113, 132].

Lower values of θ would also result in complete cascades. However as θ decreases, the dynamics of the underlying cortical networks would move away from the critical boundary between sensitivity to internal or external stimulation and stability [133]. A very low value of θ would keep the system in a constant state of activation, resembling a brain under a global epileptic seizure.

Based on the previous considerations, we select $\theta = 0.98$ as *the lowest value that results in a complete cascade across all sensory modalities*.³ We have repeated the analysis for two more values of θ (0.90 and 0.95) without any significant changes in the results (see Section B-3).

Figure 4.6 shows the complete activation cascade when the source of the stimulation is the primary visual cortex (VISp) – the corresponding cascades for the other sensory modalities are included in the Supplementary Information (see Section B-3). Note that the activation of the source triggers the activation of eleven other ROIs. Only few of them however play a major role in extending the cascade to the rest of the network: ECT (ectorhinal), PTLp (posterior parietal association), VISl (lateral visual), and POST (postsubiculum). PTLp in particular, causes the activation of seven more ROIs at the next step of the cascade. The activation of the claustrum (CLA), in this cascade, takes place through the sequence of

³The connection from AOB to COAp forms a single-edge bottleneck in the olfactory cascade. The weight of that connection in the Allen connectome is 0.46. With that value however the olfactory cascade require a different (lower) θ threshold than all other cascades. For this reason, we chose to artificially increase the weight of that connection from 0.46 to 1.

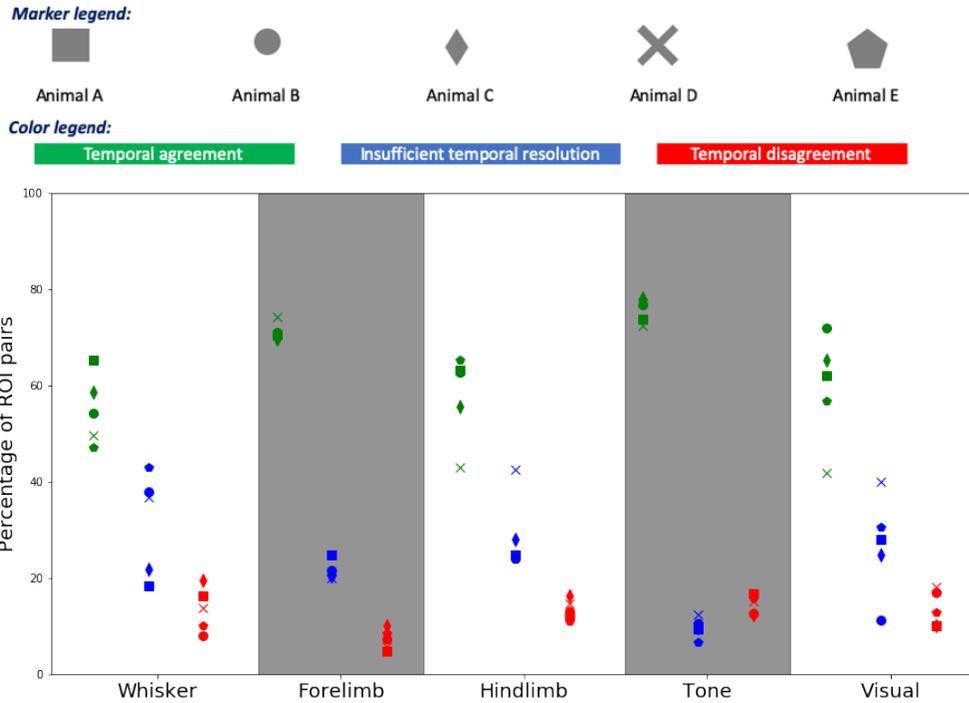
ECT, followed by TEa (temporal association).

We emphasize that the hierarchical layout shown in Figure 4.6 is specific to each sensory modality and it represents the activation cascade from the corresponding source to the rest of the cortex. This notion of hierarchy should not be confused with the hierarchical organization of the cortex [134] that results from anatomical distinctions of intracortical connections (feedforward versus feedback, based on laminar markers) [135]. The latter is an anatomical hierarchical structure, it is not specific to any particular sensory modality, and it does not convey any functional information about how one ROI may be affecting another in the presence of a specific external or internal stimulation. An activation cascade, on the other hand reveals the sequence and causal dependencies through which ROIs get activated after an initial activation at the source ROI.

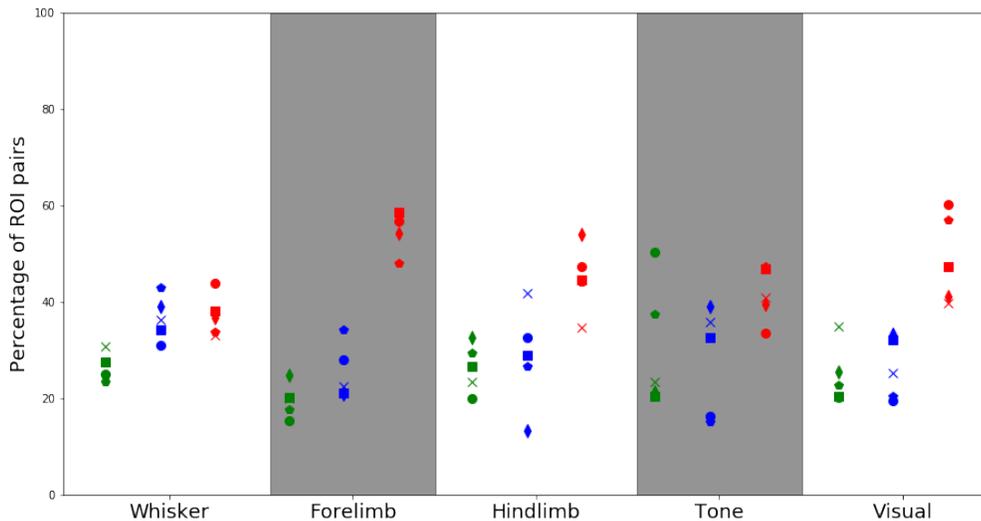
4.3.2 Model validation

We examine the validity of the ALT model predictions using functional imaging data during sensory stimulation experiments, as discussed in Section 4.2.4. The question we focus on is: *after stimulating a specific sensory modality (e.g., visual), if the ALT model predicts that an ROI X should be activated before an ROI Y , is it true that X gets activated before Y in the functional imaging data?* When this is the case, we count the pair (X, Y) as a *temporal agreement*. If X gets activated before Y in the ALT model but the opposite is true for the experimental data, we count (X, Y) as a *temporal disagreement*. Because of the finite temporal resolution in the experimental results (each frame is sampled every 7msec roughly), there are also cases where X and Y appear to be activated during the same frame, while the model always predicts a temporal difference between two activations – when that is the case, we count (X, Y) as a case of *insufficient temporal resolution*.

Figure 4.7a shows the percentage of (X, Y) ROI pairs that show temporal agreement, temporal disagreement, and insufficient temporal resolution between the activation order of X and Y in the modeling results and the mouse experiments. The plot shows results for



(a) Comparison with temporal activations based on experimental data



(b) Comparison with temporal activations based on randomized data

Figure 4.7: Comparison between model-based and experimental temporal ordering of ROI activations: a) The y-axis shows the percentage of (X, Y) ROI pairs that show temporal agreement (green), temporal disagreement (red), and insufficient temporal resolution (blue) between the activation order of X and Y in the modeling results and the mouse experiments. The plot shows results for five animals and for five sensory stimulations (a touch at the whiskers, forelimb, and hindlimb, as well as an auditory and a visual stimulation). b) The same comparison but here we have randomized the ROIs that are active during each frame, preserving the number of ROI activations in each frame.

five animals and for five sensory stimulations. Even though the variability across animals is considerable, we observe that the percentage of temporal agreement pairs, averaged across the five animals, is higher than 50% and it varies between 55% to 80% depending on the sensory modality. On the other hand, the corresponding percentage of temporal disagreement is less than 10%-20%, depending on the stimulation. In the rest of the cases, the temporal resolution is not sufficient.

We also compare the ALT modeling results with a randomized sequence of experimental activations in which we preserve the number of ROIs that are activated in each frame but assign random ROI activations during each frame. The results of that experiment are shown in Figure 4.7b. The average percentage of ROI pairs that show temporal agreement in this case varies between 20%-30%, depending on the stimulation. This comparison shows that the ALT model has significant prediction power on the temporal sequence of ROI activations, relative to a randomized baseline.

Finally, we analyzed the temporal disagreement cases between VSD experiments and modeling results to examine if certain brain regions, or pairs of regions, are over-represented in those disagreements (see Section B-4). The main result of that analysis is that the ROIs with highest disagreement cases appear at the boundary of the cortical surface at the VSD datasets and they are only partially visible. So, it is likely that the VSD data may not accurately capture the time at which those boundary regions are activated after each stimulation.

4.3.3 Similarity of sensory cascades

We next sought to answer the question of how similar the ten sensory cascades predicted by the ALT model are. The similarity between two cascades can be quantified using the *Jaccard similarity* metric. It is defined as the ratio of the common connections in two cascades over the total number of connections in those cascades.

After calculating the Jaccard similarity between every pair of cascades, we use an agglomerative hierarchical clustering algorithm to construct a dendrogram of the ten cortical

sensory cascades. This dendrogram was computed for three linkage methods: *average linkage* (the similarity between two clusters is based on the average similarity across all pairs of cascades in the two clusters), *single linkage* (based on maximum similarity), and *complete linkage* (based on minimum similarity). The resulting dendrograms are quite similar across the three linkage methods. Figure 4.5-b shows the dendrogram with average linking – the two others are shown in Figure B-13.

A first observation is that the olfactory cascade (originating at MOB) is very different than all other sensory cascades – its similarity is less than 10% with the cluster of all other cascades. This is expected given that olfaction is quite different than all other sensory processes – it bypasses the thalamus and MOB is the only primary sensory ROI in the mouse connectome that is not located in the isocortex [136, 137].

Interestingly, the two most similar cascades are the Gustatory (GU) and the somatosensory cascade of the nose (SSp-n). Further, these two cascades are quite different than all others, including the rest of the somatosensory cascades. The somatosensory cascades are quite similar to each other and they tend to cluster as follows: trunk and lower-limb (similarity of about 75%), mouth and upper-limb (about 70%), while the whiskers produce a significantly different cascade than the previous four (similarity of about 40%). This organization mirrors the anatomical layout of the somatosensory regions. The auditory and visual cascades are also quite distinct from all other cascades – but not as dramatically different as olfaction.

4.3.4 Core ROIs and hourglass architecture

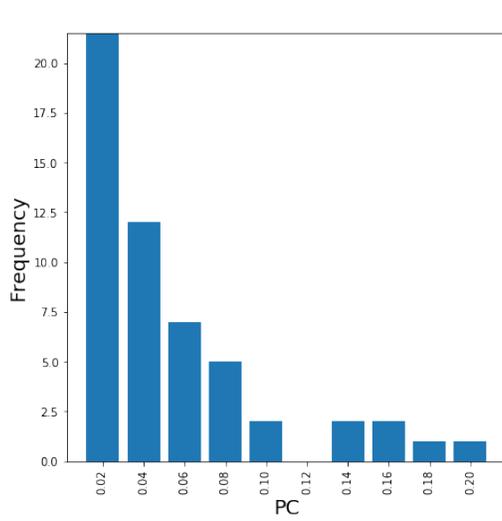
In this section, we analyze the collection of ten activation cascades (one cascade for each source) using the network analysis approach described in Section 4.2.3. The total number of source-target paths in the ten cascades is 560. The path centrality distribution, which captures how many activation paths traverse each node, is shown in Figure 4.8a. Almost half of the nodes have very low path centrality (2% or less). On the other hand, there

are four nodes with much larger path centrality – each of them covering about 20% of the source-target paths in the collection of activation cascades. These ROIs are the CLA (claustrum), SSs (supplemental somatosensory), PTLp (posterior parietal association), and AUDv (ventral auditory) areas.

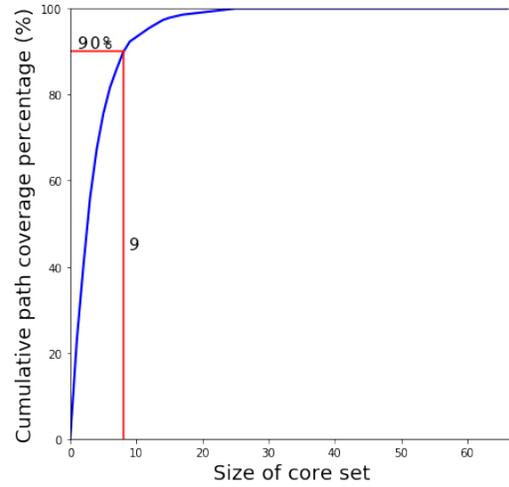
Some activation paths can traverse more than one of these highly central regions. For this reason, we compute *the minimal set of nodes that cover a given fraction τ of all source-target paths*, i.e., what is referred to as τ -core [117]. Figure 4.8b shows the fraction of covered source-target paths as we increase the size of the τ -core. The “knee-shaped” shape of this curve suggests that a small set of nodes is sufficient to cover almost all source-target paths in the activation cascades, forming a bottleneck in the MSI process. For instance, a set of nine nodes is sufficient to cover more than 90% of all source-target paths. These nine ROIs account for 14.7% of the brain volume of the ROIs we consider in the network N_c .

The small size of the τ -core, relative to the size of the network, suggests that the cortex follows an *hourglass architecture*, in which the sensory information from all different modalities is first *integrated* (other terms could be “encoded” or “compressed”) by a small set of τ -core ROIs that form the “waist” (or bottleneck) of the hourglass. Those τ -core ROIs then drive a large number of downstream ROIs that presumably operate on multi-sensory information and contribute in higher-level cognitive processes. The benefit of an hourglass architecture is that it reduces the dimensionality of the input, computing a more compact intermediate-level sensory representation at the “waist” of the hourglass. That intermediate representation is then be re-used in more than one higher-level ROIs and cognitive tasks. The hourglass architecture is visualized, at an abstract level, in Figure 4.9.

The τ -core nodes for $\tau=90\%$ are listed in Table 4.8d. Together with the percentage of *additional* source-target paths that each node contributes to the τ -core (“Path Coverage”), the table also shows the Path Centrality (PC) rank of that node. As expected, the node with the highest PC is the first node in the τ -core. After that point, the order in which nodes join the τ -core does not follow their PC ranking. The top three nodes (CLA, PTLp, AUDv) are



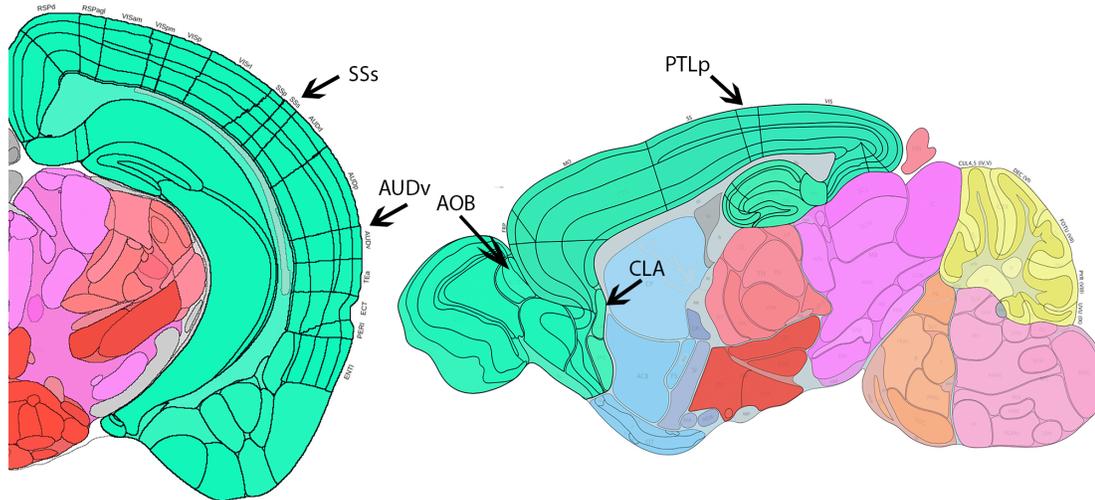
(a) Path Centrality distribution



(b) Covered source-target paths as core increases

With MOB					Without MOB	
Node	Description	Size	Path Coverage	PC Rank	Path Coverage	PC Rank
CLA	Clastrum	885	23	1	25	1
PTLp	Posterior parietal association areas	3498	17	4	19	3
AUDv	Ventral auditory area	2554	15	3	15	4
AOB	Accessory olfactory bulb	571	11	5		
SSs	Supplemental somatosensory area	5729	8	2	9	2
MOs	Secondary motor area	10098	6	8	8	5
ACAd	Anterior cingulate area, dorsal part	2789	4	17	4	13
VISl	Lateral visual area	880	3	11	4	6
ECT	Ectorhinal area	3390	2	25	2	19

(c) Core regions for $\tau=90\%$



(d) The location of the nine τ -core regions on the Allen atlas

Figure 4.8: *Path centrality and τ -core analysis*: (a) Path Centrality (PC) histogram for the 67 regions in N_c , considering all source-target paths across the ten activation cascades. (b) Cumulative path coverage by the top-X core nodes for $X=1 \dots 67$. Nine regions are sufficient to cover $\tau = 90\%$ of all paths.

sufficient to collectively cover about 60% of the activation paths. Note that none of these τ -core ROIs are a primary sensory region (i.e., a source node). Instead, they are either ROIs that are not typically associated with a single sensory modality (CLA, PTLp, ACAd, ECT) or they are ROIs that are often thought of as “secondary” or “supplemental” to a certain sensory modality (AOB, AUDv, SSs, MOs, or VISI). If we exclude the MOB cascade, the only difference is that the τ -core will not include the AOB region.

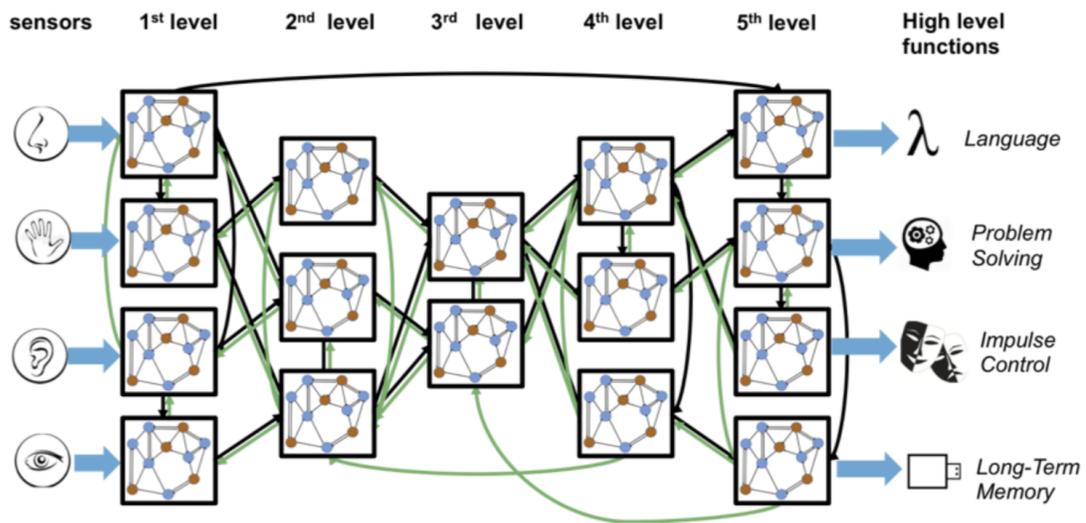


Figure 4.9: *Illustration of hourglass architecture:* A hypothetical network with feedforward, feedback and lateral connections between regions at different levels of the cortical hierarchy. Input information is provided at sensory-specific modules (left), while high-level cognitive tasks are performed by association regions at the other end of the hierarchy (right). The “hourglass feature refers to the fact that the high-dimensional input information is first integrated through through a relatively small number of highly central intermediate-level regions, before it is re-used at high-level cortical regions and tasks.

4.3.5 Location of τ -core nodes in activation cascades

Location relative to sources

In this section, we first investigate the *topological location* (rather than anatomical location) of the τ -core nodes relative to the source of each activation cascade. Are the τ -core nodes,

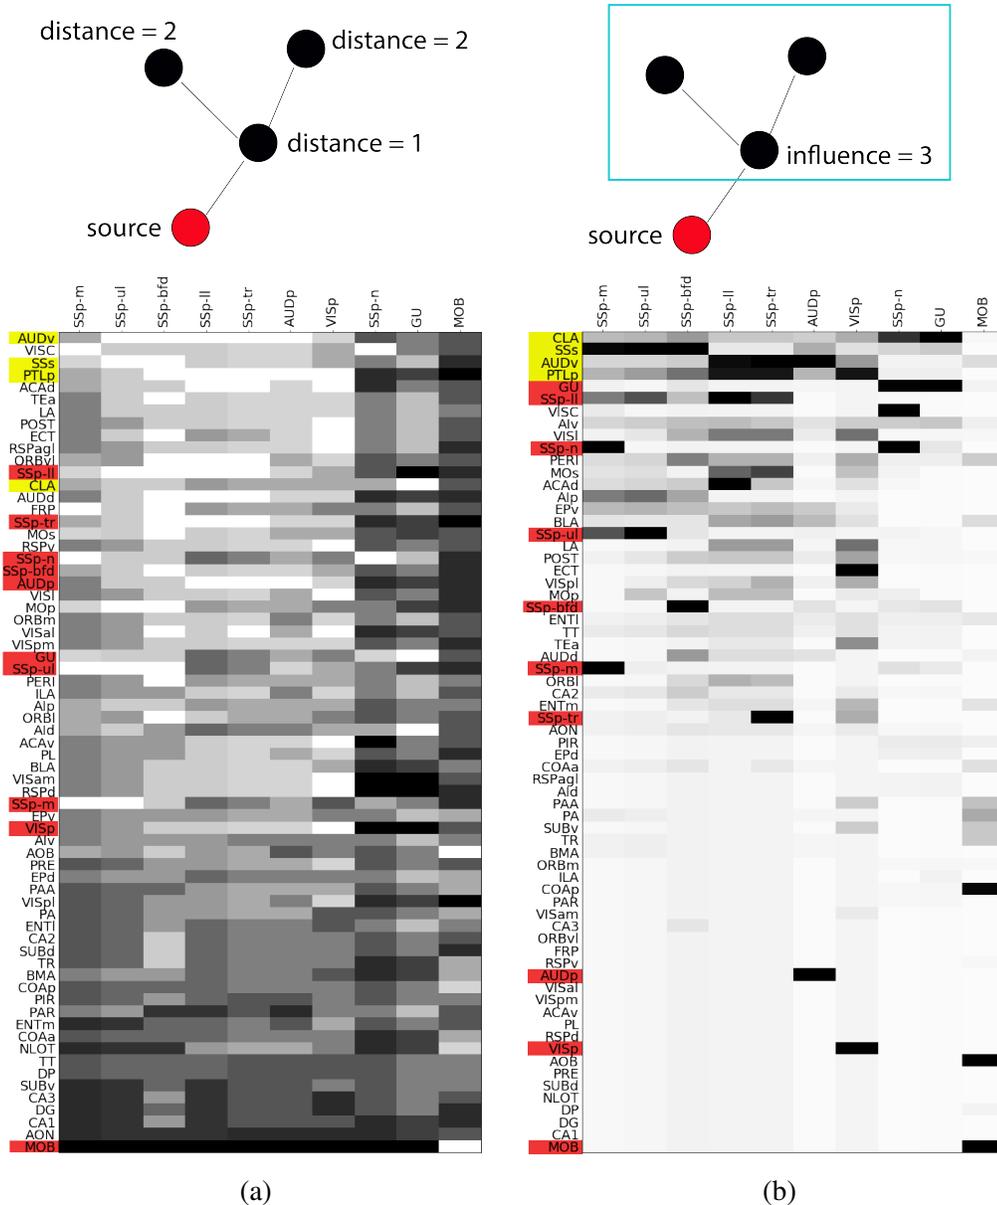


Figure 4.10: *Location-related metrics*: In both matrices, a column represents one of the ten activation cascades, originating at the node shown at the top of the column. (a) Each row represents the *source-distance* of the corresponding node from the source of that column's cascade. White denotes a distance of one hop, while black denotes the maximum distance for that cascade. Rows are ordered in terms of the average distance (in number of hops) of the corresponding node from the sources of all activation cascades (excluding the MOB cascade, which is very different). (b) Each row represents the *influence* of the corresponding node, i.e., the number of nodes that are reachable from that node in the activation cascade that the column represents. White denotes an influence of one (only that node), while black denotes an influence that covers all network nodes. Rows are ordered in terms of the average influence of the corresponding node across all activation cascades (excluding the MOB cascade).

which form the waist of the hourglass architecture, closer to the sources or targets of the activation cascades? And how does their location compare to the location of the sources and of other cortical ROIs? These questions are related to recent experimental work suggesting that cross-modal representations are constructed at early stages of the sensory information flow [138].

We focus on the top-4 τ -core nodes (CLA, PTLp, AUDv, SSs), which collectively cover about $\tau=70\%$ of all source-target paths (see Table 4.8b). Figure 4.10a visualizes in grey-scale the location of each node (matrix row) in each activation cascade (matrix column). Each source node (represented with red) is obviously at a distance of zero in its own activation cascade. Note however that source nodes can be at a much larger distance from sources of other activation cascades. For instance, the primary visual cortex (VISp) appears at the maximum distance in the nose (SSp-n) and gustatory (GU) cascades.

The four τ -core nodes we consider (represented with blue rows) *are relatively close to all source nodes*: AUDv has the lowest average distance to all of the primary source areas, while SSs and PTLp are ranked as third and fourth. The claustrum (CLA) is slightly further away from the sources, ranked 13th (out of 67) in the previous ranking. If we consider a higher value of $\tau=90\%$, the additional τ -core nodes (MOs, ACAd, VISl, and ECT) are ranked as 17th, 5th, 22nd, and 9th in terms of their average distance from sources. In summary, *all τ -core nodes appear in the top one third of the distance ranking, and so they are closer to the sources of the hourglass architecture than to its targets.*

Location relative to targets – influence

Another way to examine the location of a node v in the hourglass architecture is in terms of how many nodes appear in activation paths downstream of v – a metric that we refer to as the *influence* of v . Figure 4.10b visualizes in grey-scale the influence of each node (matrix row) in each activation cascade (matrix column). The source of a specific cascade has, by definition, maximal influence (i.e., all network nodes) in its own cascade – but it may have

a much lower influence in other cascades. Indeed, the influence of source ROIs (shown in red) does not seem to follow a coherent pattern: the gustatory (GU) and somatosensory area of the lower-limb (SSp-II) are sources with high influence but the primary visual cortex (VISp), the primary auditory cortex (AUDp) or the main olfactory bulb (MOB) are sources with low influence in other cascades.

On the other hand, the four most important τ -core nodes (CLA, PTLp, AUDv and SSs) also occupy the top-four positions in terms of influence. The next four τ -core nodes (MOs, ACAd, VISl, and ECT) have high influence as well, ranked as 12th, 13th, 9th and 20th, respectively.

Combining the previous observations about the influence of τ -core nodes as well as their distance from sources, we can summarize our findings as follows: *τ -core nodes are close to most sensory sources and they also influence the activation of many downstream nodes.* These two features place τ -core nodes at a location that allows them to both integrate sensory information from different sources as well as to use that integrated information in driving many downstream ROIs.

4.3.6 Robustness of τ -core nodes

In this section, we examine the robustness of the previous results regarding the τ -core when we randomize the edges and weights of the underlying connectome. We also examine whether the length and/or weights of these connections are responsible for the hourglass effect and for the specific regions that form the τ -core.

We create ensembles of random connectomes, derived from the mouse connectome in four different ways:

1. Randomize the weight assigned to each edge, reallocating the weights of the original connectome across randomly selected connections but maintaining the topology.
2. Randomize the physical length of each edge (and thus its communication delay in

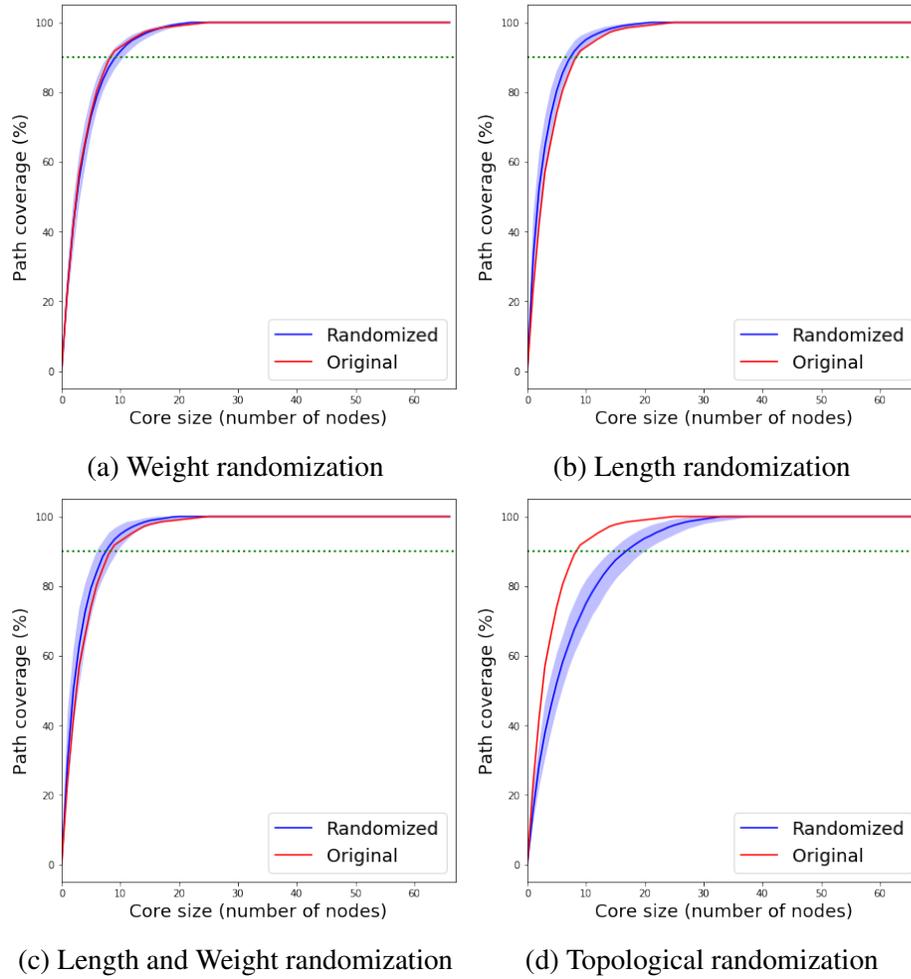
the ALT model), again reallocating randomly the lengths of the original connections but maintaining the topology.

3. Randomize both the weights and lengths assigned to each edge, as previously mentioned. We do not maintain any correlation between weights and lengths.
4. Randomize the connectome's topology by swapping connections between randomly selected pairs of nodes. This randomization method preserves the in-degree and out-degree of each node.

Figures 4.11a-4.11d focus on the first three randomization methods: weights, lengths, and their combination. In all cases, the τ -core size of the original network is contained in the 5% confidence interval of 100 randomized networks. In other words, *the weights and physical lengths of the connectome's connections do not play a significant role in the number of τ -core nodes, for any value of τ .*

On the other hand, when we randomize the topology of the connectome, the τ -core size doubles in size when $\tau=90\%$: from nine nodes in the original network to eighteen. Additionally, it takes about half of the entire network to cover all activation paths in the collection of ten A-DAGs. So, *it is the graph structure of the connectome (i.e., its topology) that leads to a small τ -core size – not the weight and/or length of the connections.*

Even though the weight and length of the connections do not have a strong effect on the τ -core size, do they affect the identity of the ROIs that participate in that τ -core? To answer this question, Table 4.11e shows the fraction of random networks that include each of the eight τ -core nodes in Table 4.8d. The Claustrum (CLA), for instance, appears in the MSI τ -core of 88% of the networks that have randomized connection lengths (89% for randomized weights) but in only 13% of the networks that have randomized topology. The results are similar for the top-6 MSI τ -core nodes: they appear in the MSI τ -core of most randomized networks when we randomize connections weights and/or lengths – but they rarely appear in the MSI τ -core when we randomize the topology. For the last two



(e) Probability that each of the eight MSI τ -core nodes appears in a randomized network

Node	Description	Probability of appearance in randomized networks			
		Length	Weight	Length & Weight	Topology
CLA	Clastrum	.88	.89	.83	.13
PTLp	Posterior parietal association areas	.97	1.00	.92	.10
AUDv	Ventral auditory area	.97	.95	.77	.09
SSs	Supplemental somatosensory area	.68	.86	.62	.14
MOs	Secondary motor area	.84	.87	.87	.11
ACAd	Anterior cingulate area dorsal part	.83	.72	.59	.14
VISI	Lateral visual area	.25	.77	.51	.21
ECT	Ectorhinal area	.12	.17	.7	.11

Figure 4.11: *Robustness results*: The effect of different connectome randomization methods on the core size. Light blue shade marks the 5% to 95% values among 100 randomization runs, while the solid blue line is the median of these runs. The red line represents the τ -core size for the original connectome. The dotted green line marks the τ -core size for $\tau=90\%$. The table at the bottom shows the fraction of random networks that include each of the eight τ -core nodes.

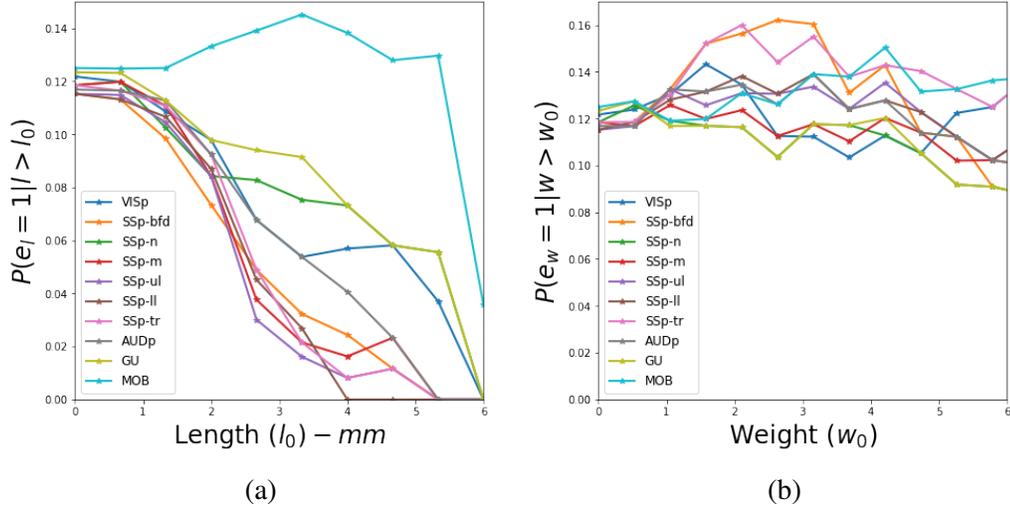


Figure 4.12: *Connections that appear in sensory cascades:* (a) Conditional probability that a connection e_l of physical length l appears in an activation cascade given that $l > l_0$. (b) Conditional probability that a connection e_w of weight w appears in an activation cascade given that $w > w_0$.

MSI τ -core nodes (VISl and ECT) their membership in the MSI τ -core is *not* as robust: randomizing connection lengths has a major effect in the appearance of VISl in the τ -core, and randomizing any aspect of the network has a major effect in the appearance of ECT in the τ -core.

4.3.7 Which anatomical connections are more important in sensory cascades?

Which anatomical connections are more important in terms of MSI? Only about half the connections of the anatomical connectome appear in sensory activation cascades, and a quarter of the former appear in only one sensory cascade.

To answer this question, we examine the conditional probability that an connection e_l of physical length l appears in an activation cascade given that $l > l_0$; we denote this probability as $P[e_l = 1 | l > l_0]$. Similarly, we define the probability $P[e_w = 1 | w > w_0]$ for an edge of weight w .

Figure 4.12 shows these two conditional probabilities separately for each of the ten

activation cascades.⁴ Even though there are significant variations across the ten cascades, all of them show that $P[e_l = 1 | l > l_0]$ decreases with l_0 , i.e., *as the length of a connection increases (and especially when l is larger than 1-2mm) it becomes less likely that it will be part of an activation cascade.*

The right part of Figure 4.12 shows the corresponding results for the connection weight conditional probability $P[e_w = 1 | w > w_0]$. With the exception of the olfactory cascade (MOB), which shows a decreasing trend, the rest of the sensory cascades show a more complex and diverse pattern. The average probability suggests that weight is *not* a significant factor in determining which connections will be part of an activation cascade (perhaps with the exception of edges with very high weights – larger than 8). One reason is that if the weight of a connection is higher than the activation threshold θ (which is about one in our modeling results), that edge is sufficient to activate a downstream node, independent of the state of other connections to the same destination node.

These results suggest that *sensory cascades spread in the cortex as a forest fire mostly through short connections connecting physically adjacent regions, rather than through the (relatively few) long connections that connect remote regions.*

4.3.8 Comparison with static network analysis methods

Finally, we asked whether we would we obtain similar results if we had relied on static network analysis metrics and approaches (such as [139, 140, 86]) that do not consider a model of dynamic network activity?

We first ask whether the path centrality metric correlates strongly with more commonly used node centrality metrics, namely: incoming or outgoing strength, betweenness centrality, closeness centrality, pagerank and eigenvector centrality [141]. The results of this comparison are included in the Supp-Info section B-7. In brief, the conclusion of that

⁴The probability for $l_0 = 0$ is roughly 0.12 because the connectome has 616 connections and the average cascade includes only 75 of these edges – so the (unconditional) probability that a connection appears in at least one activation cascade is about $75/616 \approx 0.12$.

comparison is that none of these centrality metrics correlates well with path centrality, computed over all source-target sensory activation paths. The identification of nodes that play an important role in multisensory integration requires an approach that models communication dynamics over the anatomical network – not just the topological properties of the latter [142, 97].

We also examined the similarity between the nine τ -core nodes, the five rich-club nodes [143, 144, 145], as well as the thirteen core-periphery nodes computed using Rombach’s method [146]. The complete results appear in B-7. The overlap between the rich-club and the τ -core nodes consists of only the claustrum (CLA) and the supplementary motor region (MOs). The overlap between Rombach’s core nodes and the τ -core nodes consists of only MOS and AOB.

In summary, *this analysis highlights the marked differences between ALT-based network diffusion modeling and static network analysis methods based on centrality metrics or core-periphery concepts.*

4.4 Discussion

This study built a suite of tools to quantify and characterize multisensory integration from the perspective of communication dynamics. We chose to analyze the mouse connectome due to the availability of a detailed meso-scale anatomical map and because the problem of multisensory integration is relatively well-studied in rodents [122]. In particular, we analyzed the cortical brain network from the Allen Mouse Connectivity Atlas [26] and focused on “early” dynamics of sensory integration. By *early* we specifically mean the first wave of cortical activity, starting from primary sensory areas and propagating to the whole-hemisphere. This point differentiates our work from earlier models of multisensory integration [147, 148].

The underlying anatomical pathways recruited by diverse sensory modalities (e.g., visual, auditory, somatosensory) branch out rapidly and become increasingly complex as they

reach the higher associative cortical areas [27]. To capture the nonlinear and asynchronous nature of these dynamics, we used a diffusion model that can capture both non-linearity and communication delays, called Asynchronous Linear Thresholding (ALT). We found that ALT can closely recapitulate sensory cascades when compared with VSD datasets from sensory stimulation experiments on mice [27].

ALT is a phenomenological model that aims to describe (but not explain mechanistically) the diffusion of information at the ROI level. As such, it does not model the underlying mechanisms of neuronal communication through chemical or electrical synapses, and it is also quite different than neural mass models that capture coarse grained activity of large populations of neurons and synapses [149]. Similar to neuroimaging modalities of similar spatial resolution, such as task-based fMRI, ALT aims to capture how a certain activity, namely the stimulation of a primary sensory region, causes the activation of other brain regions. A difference with fMRI or MEG however, is that the resulting activation cascades can be analyzed to infer interactions between ROIs that participate in the cascade, as described above. Additionally, ALT depends on the communication delays and weights of the connections between brain regions. Thus, it produces a timeline of activation events, one for each ROI that participates in the cascade. The length of the time interval between activation events, at least in relative terms, can be compared with experimental results from neuroimaging modalities that have fine temporal resolution, such as VSD (as in this study) or calcium imaging.

On the other hand, the ALT model cannot capture more complex dynamics, such as sustained oscillations at certain frequencies or feedback from a newly activated region back to regions that were activated earlier. Modeling such dynamics would require more elaborate neural mass models – their validation however would require whole-brain neuroimaging data of fine spatial and temporal resolution (higher temporal resolution than fMRI and higher spatial resolution than EEG). Additionally, those more complex models may not be necessary when the goal is to map the feedforward initial propagation of brain activity after

a sensory stimulation.

Our primary result is that relatively few cortical regions are responsible for integrating almost all sensory information. This finding supports the idea that multisensory integration is performed through an “hourglass architecture”. The benefit of an hourglass architecture is that it first reduces the input dimensionality of the sensory stimulus at few intermediate-level modules that reside at the “hourglass waist”. Second, it re-uses those compressed intermediate-level representations across higher-level tasks, reducing redundancy between the latter. The hourglass analysis framework was first developed in [117] and it has been recently applied in the connectome of *C.elegans* [150]. There are two fundamental differences in our study: i) we rely on communication dynamics while [150] uses a pre-defined set of “routing mechanisms” to construct ensembles of sensory pathways, and ii) the sources and targets of the *C.elegans* cascade are the sensory and motor neurons, respectively. Nonetheless, it is interesting that the “hourglass architecture” emerges in both studies. One possibility is that this architecture is selected by evolution because it drives a network towards reusing a small set of intermediate functions in constructing a range of redundant output functions.

Rather than studying each sensory cascade in isolation, our analysis framework is based on the combination of all sensory cascades. Comparative and competitive cascades have been studied in [151] to quantify the combined effect of multiple cascades, using simultaneously activated source nodes. Sensory stimuli of different modalities, however, do not need to arrive at the cortex simultaneously in order to be integrated [152]. Different sensory stimuli travel at different speeds through body receptors [153]. The analysis framework that we followed constructs uni-sensory cascades and merges their activation paths. We have also experimented with cascades originating from two simultaneously activated sources (see Section B-6). The hourglass architecture is still observed in that case, and the core nodes are mostly the same.

There is mounting evidence that points to the fact that many computations in the brain

are multi-sensory. Even primary sensory cortices respond to sensory stimuli from different modalities. For instance, there is significant activity in olfactory cortex after delivering taste stimuli to the tongue [154]. Or, many units in the auditory cortex respond to visual stimuli [132]). A number of studies have shown primary-like responses from multiple sensory modalities [130, 131], advancing the notion that much of cortex is multi-sensory [113]. Our results shed new light on this debate, and support the findings that many regions in early sensory cascades are multi-sensory, the most important of which are the claustrum and the posterior parietal cortex.

The claustrum is known for its anatomical uniqueness [155] and its precise function has been enigmatic [156, 116]. The late F. Crick has hypothesized that the claustrum may be a potential gateway to consciousness [157]. More recent results have demonstrated that the claustrum is important in gating sensory information, and in attention mechanisms in visual perception [156]. These and other studies provide strong evidence that the claustrum is a crucial node for multisensory integration. Our findings are in agreement with this growing body of work.

The posterior parietal cortex is the second most important region for MSI, according to our analysis. This stems from its strong and immediate connectivity to primary sensory regions and its projections to motor areas. PTLps connectivity points to its role in integrating sensory information to direct immediate motor commands [158]. The experiments of Nikbakht et al. [159] provide direct evidence that PTLp is multisensory both at the behavioural and neurophysiological levels and it provides sensory independent information about the orientation and categorization of objects in the environment. We refer the reader to [160] for further discussion.

Three additional core nodes have been associated with specific sensory modalities in the past: Ventral auditory area (AUDv), Supplemental somatosensory area (SSs), and Lateral visual area (VISl). However, all of them have also been implicated to some extent with multisensory processing. For instance, Hishida et al. [161] found that activity propagating

to the parietal association area passes through the ventral auditory region, irrespective of the sensory stream source (visual, auditory or somatosensory). Similarly, [162] suggests that SSs has a major role in bringing context to the sensory pathways [163]. VISl is on the dorsal visual stream and is associated mostly with spatial location and action guidance [164], while the other secondary visual areas participate in the ventral stream [165].

An important next step is the inclusion of subcortical regions in the diffusion model. Various nuclei in the thalamus or the superior colliculus are known to be crucial in MSI [166]. An integrated model of both cortical and subcortical activity, in the context of sensory integration, may help to explain the complex cortico-thalamic feedback mechanisms [167, 168]. Disentangling such communication dynamics requires working at finer spatial resolutions, and including distinct cortical layers (laminar divisions) [169]. Incorporating high-resolution [170] and cell type specific [135] information about connectivity in our pipeline is an interesting direction, that can allow segregation of different communication channels. By increasing the spatial resolution and the complexity of the model dynamics, we could capture more complex population-level activity patterns (e.g., Wilson-Cowen [149]) and better understand the role of mesoscale network topology in constructing a coherent perceptual state from raw sensory streams.

CHAPTER 5

CLOSING REMARKS

5.1 Conclusions

We have now come to the conclusions of this thesis. The ultimate aim of which was a topological analysis of the macroscale brain connectome, both structurally (i.e. only anatomical) and also by the means of network dynamic models. The path to this goal started at the human connectome. Although recent advances in MRI have provided neuroscientists with an abundance of human brain scans, unfortunately these scans are only indirect measurements of the structural connections in the brain and often susceptible to subjective interpretation, to an extent that the most fundamental property of the human connectome, its density, is still elusive and debated.

Many sophisticated steps from MRI acquisition to connection probability inference for the human connectome often end with proportional or weight-based thresholding that is arbitrary and study dependent. In the pursuit of a better, more reproducible, and objective connectome, we started our work by introducing a network inference framework dubbed **MANIA**. MANIA relaxes the need to subjectively fix a threshold in probabilistic tractography studies to infer connectome density or filter out spurious connections. We rigorously validated our method using synthetic data from phantom MRI acquisitions, and also using high-resolution MRI data. After showing the accuracy of our method, consequent topological analyses for human brains were based on connectomes inferred by MANIA. In particular, we ran a subsequent case study of 27 healthy subjects to model the major the depression disorder circuitry. MANIA circuitry was compatible with clinical practices that seek to cure the disorder using deep brain stimulation. Further, we presented one of the first cortical network on the renowned data repository used by neuroscientists from around

the world: the human connectome database (<https://db.humanconnectome.org>).

Although the human connectome is undoubtedly the most interesting to study dynamic processes on, it is unfortunately undirectional for in-vivo dMRI connectomes. Furthermore, the accuracy and resolution of the connectomes for other mammals, notably macaque and mice, are still significantly ahead of those of humans. The main reason behind this is the possibility of using intrusive surgeries and introducing bio-markers in the brain of these mammals. For example, The Allen Institute published a resolute atlas of the mouse brain which modelled connectivity between approximately five hundred regions of the mouse brain. Apart from the resolution and greater confidence in the connectome, the connectome is directed network as opposed to that inferred from diffusion data, as in the case of human MRIs.

More suitable dataset plus the fact the neural networks are believed to have same dynamics at different scales across the species convinced us to use mouse connectome for the study of a fundamental dynamic: Multi-sensory integration.

We proposed a framework to model sensory stream flows and integration on the anatomical brain connectomes. The framework is accompanied by rigorous experimental validation using the state-of-the-art microscopy with both high temporal resolution (< 7 millisecond) and high spatial resolution (< 8 micrometer). After validating our computational model, we apply it on ten sensory regions of the mouse cortical connectome and found that the multi-sensory integration and unified perception view of the world is facilitated through an hourglass architecture in which few core nodes integrate the sensory information and disseminate a digested view to many consumer nodes.

Therefore, the contributions of this thesis can be categorized into the following two groups of table 5.1

Table 5.1: Summary of the contributions in this thesis by chapter and type - group 1 is related to the anatomical analysis while group 2 is related to the MSI dynamic

Contribution group	Chapters	Type	List
Group 1	Chapter 2	Methodological	Introducing of network inference framework (MANIA) for converting probabilistic tractography data to actual anatomical connectome and extending the method to support both individual and group connectome inferences.
	Chapter 3	Finding	Constructing networks to model the brain circuitry implicated in mental disorders and identifying the BA25 to Neucleus Accumbens as the most important connection in the circuit.
	Chapter 3	Finding	Presenting cortical anatomical connectome for the human brain using state-of-the-art diffusion data and parcellation published by HCP.
Group 2	Chapter 3	Finding	Comparison between 7T diffusion data and 3T counterparts: We validated MANIA in 3T by benchmarking better data at 7T.
	Chapter 4	Methodological	Introducing a computational framework to model the propagation of sensory information in the brain at the early cascade phase.
	Chapter 4	Finding	Presenting evidence of an hourglass architecture in the process of MSI in the mouse cortex.
	Chapter 4	Finding	Showing that nodes in the hourglass waist that we discovered in MSI have negligible intersection with nodes ranked high by static network centrality metrics.
	Chapter 4	Finding	Presenting sensory cascades from 10 primary sensory regions in the mouse cortex - we compared these cascades together and showed that the olfactory cascade is the most different of all.

5.2 Limitations

5.2.1 Lack of ground truth to quantify the accuracy of tractography methods

In the course of this thesis, we have discussed the impact of many algorithmic and methodological choices and their shortcomings in inferring human brain anatomical connectome, e.g., definition of nodes in networks, distance bias, gyral bias and seeding strategy. Unfortunately measuring the ultimate accuracy of different connectome inference methodologies in the space of all these choices is only possible indirectly for the lack of ground truth connectome. Thus, we resorted to case studies on known circuitry (major depression disorder), diffusion MR phantom study, and high resolution MR images to quantify the accuracy of our method in a diverse way. Although, there have been efforts to directly validate tractography results using chemical tracers, but they have their own limitations and biases. All in all, the ultimate benchmark to measure the accuracy of inferred anatomical connectomes is yet to come.

5.2.2 Undirected nature of diffusion data

A major limitation in mapping brain connections from tractography data is that diffusion MRI does not provide any information about the direction of connections. Although, we exploited this limitation to our advantage in chapter 2 and introduced MANIA, the same limitation prohibited us from using MANIA connectomes in studying MSI dynamics of chapter 4. Studying communication dynamics requires direction of connections in the underlying network, and the diffusion MRI cannot provide it. Thus, the study of brain-wide causal communication dynamics (such as MSI) on anatomical connectomes is still impossible in-vivo for humans.

5.2.3 Partial validation of the ALT model

We ran ALT model from ten primary sensory regions in chapter 4. However, we examined the validity of the ALT model predictions for 5 primary sensory regions since VSD experimental data for other sensory regions was not available at the time of this study. Also, the temporal resolution of VSD did not allow us to quantify the accuracy of ALT predictions thoroughly; thus, there were cases in which the results were undecided in terms of correctness.

5.3 Extensions

Throughout the thesis we have suggested several directions in which this research can be extended in order to address its limitations or expand its reach. Here, we elaborate upon the most immediate or interesting ones:

- Thresholding is not the only problem standing in the way of an accurate connectome. The results of probabilistic tractography are severely compromised by distance bias, i.e., it is harder to find longer connections compared to short ones. It is now established that the human brain has topological features that are only possible because of the presence of the direct communication of distant regions, hence the long connections. Addressing the distance bias problem by systematic modelling of the probabilistic tractography will lead to a reduction of the bias of distance on probabilistic streamline counts. This will be a considerable contribution to the field of connectomics.
- The weights of anatomical connectomes are still very controversial; how would one convert the tractography results to a weighted connectome which is biologically relevant - pruning out the spurious connections and yet constructing a biologically weighted connectivity matrix? In MANIA, we take a different stand on the weight issue compared to [91]: instead of interpreting the probabilistic spatial maps that are

generated using probabilistic tractography as strength-related weights, we generated a binary network, but with a separate confidence value for each edge. From this point of view, MANIA generates a binary network, but with the relative confidence we have in the presence or absence of each edge.

- *Comparative connectomes* and *connectome across the age* are two important and active trends in neuroscience [171, 172]. Studying variations in the brain across species and ages using the toolkits developed in this thesis, or studying the topological features (such as the hourglass structure) across species and ages, are important untapped directions for future research.
- The studies in chapters 3 (case study B) and 4 are based on cortical brain regions. The cortex is believed to house higher associative/cognitive functions, but the link to subcortical regions cannot be ignored. Given that the cortex is the outer layer of the brain, the quality and availability of data from the cortex is higher than that of the subcortical regions. However, subcortical data is also sporadically available, providing a great opportunity to extend the work in this thesis. For example, is the hourglass structure also seen in the subcortical regions? Is the distance bias model of chapter 3 valid for data from subcortical regions?
- The results here apply to the macroscale brain. Given the availability of data and pipeline for mesoscale analysis, it will be worth exploring how the topological features vary at that scale.
- We studied the multi sensory brain integration using the primary sensory regions. One can run the same analysis in other cortical or subcortical regions. Since photostimulation in mice has made it possible to study the cascade from non-sensory regions, the study of these cascades and their differences to sensory cascades remain open and interesting directions for future research.

Appendices

APPENDIX A: LIST OF HCP SUBJECTS USED IN CHAPTER 3

Table A.1: List of the subjects used in our study

Number	HCP ID	Release	Quarter	Gender	Age
1	126426	S1200	Q12	F	31-35
2	135124	S1200	Q13	F	31-35
3	137431	S1200	Q13	F	31-35
4	144125	S900	Q12	F	31-35
5	146735	S1200	Q12	F	31-35
6	152427	S1200	Q13	F	31-35
7	153227	S900	Q12	F	31-35
8	177140	S1200	Q13	F	31-35
9	180533	S1200	Q12	F	31-35
10	186545	S1200	Q13	F	31-35
11	188145	S1200	Q12	F	31-35
12	192237	S1200	Q11	F	31-35
13	206323	S1200	Q12	F	31-35
14	227533	S1200	Q08	F	31-35
15	248238	S1200	Q13	F	31-35
16	360030	S1200	Q13	F	31-35
17	361234	S900	Q12	F	31-35
18	362034	S1200	Q13	F	31-35
19	368753	S1200	Q12	F	31-35
20	401422	S1200	Q08	F	31-35
21	413934	S1200	Q09	F	31-35
22	453542	S1200	Q12	F	31-35
23	463040	S1200	Q13	F	31-35
24	468050	S1200	Q08	F	31-35
25	481042	S1200	Q12	F	31-35
26	825654	S1200	Q12	F	31-35
27	911849	S1200	Q13	F	31-35
28	917558	S900	Q12	F	31-35
29	992673	S900	Q12	F	31-35
30	558960	S1200	Q12	F	31-35
31	569965	S1200	Q13	F	31-35
32	644246	S1200	Q12	F	31-35
33	654552	S1200	Q12	F	31-35
34	680452	S1200	Q13	F	31-35
35	701535	S1200	Q13	F	31-35
36	804646	S1200	Q13	F	31-35
37	814548	S1200	Q13	F	31-35

APPENDIX B: SUPPLEMENTARY INFORMATION FOR CHAPTER 4

B-1 Sources of sensory cascades

The ten cortical regions we consider as sources of sensory activation cascades are shown in Table B-1.

Table B-1: *Sources of sensory cascades*

Acronym	Description	Major region	Size (number of voxels)
AUDp	Primary auditory area	Isocortex	2689
VISp	Primary visual area	Isocortex	6227
GU	Gustatory areas	Isocortex	2104
SSp-n	Primary somatosensory area, nose	Isocortex	1358
SSp-bfd	Primary somatosensory area, barrel field	Isocortex	10306
SSp-ll	Primary somatosensory area, lower limb	Isocortex	3254
SSp-m	Primary somatosensory area, mouth	Isocortex	2924
SSp-tr	Primary somatosensory area, trunk	Isocortex	4799
SSp-ul	Primary somatosensory area, upper limb	Isocortex	5406
MOB	Main olfactory bulb	Olfactory Areas	16978

B-2 Connectome filtering

We have repeated the analysis for the following p-values: 0.01, 0.02, \dots 0.1. When the p-value is smaller than 0.03, the network becomes disconnected. When the p-value is higher than 0.04, the core nodes remain the same as in Table 4.8d.

Figure B-1 shows the network density for different p-values (left) as well as the connection density (edge weight) versus p-value for each structural connection among the 67 ROIs we consider (right).

B-3 The activation cascade of each sensory source

The ten activation cascades, one for each sensory source, are shown in Figures B-2 through B-11.

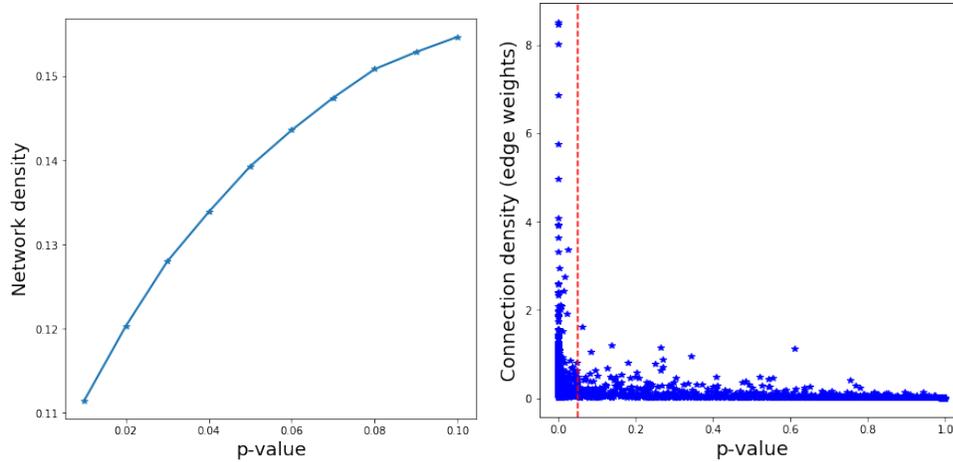


Figure B-1: *Connectome filtering*. **Left:** Network density versus edge p-value. **Right:** Connection density (edge weight) versus edge p-value

Sensitivity to activation threshold θ

Recall that if θ is higher than 0.98, the cascade of some sources will not be complete. We have also computed the core nodes for two lower values of θ : 0.9 and 0.95. The 70%-core remains exactly the same. The 90%-core for $\theta = 0.9$ includes PERI instead of ECT.

B-4 Analysis of “disagreement cases” between VSD data and ALT modeling results

In this section we analyze the cases in which the VSD experimental results predict a different temporal ordering than the ALT modeling results. Recall that a disagreement refers to a pair of ROIs (e.g., X and Y) for which the activation order in ALT is different than that in VSD. Thus, every disagreement case involves two different ROIs.

We first split the ROI pairs in two sets, the disagreement cases and the agreement cases. We measured the Euclidean distance between the center of the two ROIs in each pair of the two sets. The hypothesis that the sample mean of these distances is the same could not be rejected with a p-value of 5%.

We did the same for the connection weight between connected ROI pairs, comparing the average weight of pairs in the set of agreements and the set of disagreements. Again,

the hypothesis that the sample mean of these distances is the same could not be rejected with a p-value of 5%.

So, the distance or connection weight between two ROIs does not predict whether they will be a disagreement case.

Are there certain ROIs that appear in surprisingly many disagreement cases? To answer this question, we measure the number of times X that each ROI appears in the N disagreement cases that are observed in the cascade of a certain sensory modality (considering the datasets from all five animals). If an ROI appears in $k > \mu + 3\sigma$ disagreement cases, we conclude that it is significantly over-represented in the group of disagreement cases. The mean and standard deviation in the previous inequality are calculated based on the null model that the disagreements involve randomly chosen pairs of distinct ROIs. So, if K is the number of ROIs and we have N disagreements (for a specific sensory cascade), then the probability of selecting a specific ROI in each pair is:

$$p = \frac{1}{K} + \left(1 - \frac{1}{K}\right) \times \frac{1}{K-1} \quad (\text{B.1})$$

Given that we sample N pairs, the null model is that we will sample each ROI a number of times X , where X follows the *Binomial*(N, p). So, $\mu = Np$ and $\sigma^2 = Np(1-p)$.

Figure B-12 shows the results for this study. Note that in each of the five cascades, it is only one or two ROIs that are significantly over-represented in disagreement cases. In the visual cascade for instance, it is the SSp-bfd ROI that is present in surprisingly many disagreements.

We further analyzed these disagreements in which one of the two ROIs is an outlier, asking whether the second ROI is also over-represented. Figure B-12f identifies such ROI pairs – four of the cascades have only one pair while the Forelimb stimulation cascade has none. ACAd is an ROI that appears in many disagreements in the visual, whisker and tone stimulation cascades – note that this ROI appears at the boundary of the cortical

surface at the VSD datasets and it is only partially visible. So, it is likely that the VSD data do not reflect accurately the time at which that region is activated after each stimulation. To a smaller degree, the same may be true for TEa, which is the outlier in the hindlimb stimulation cascade.

B-5 Similarity between activation cascades with single and complete linkage

Figure B-13 shows hierarchical clustering dendrograms quantifying the similarity between the ten activation cascades based on single (left) and complete (right) linkage. The corresponding average linkage plot is shown in Figure 4.5-b.

B-6 Activation cascades when two sensory sources are activated simultaneously

We have also considered the case of two simultaneously active sources, considering all possible pairs of sources ($10 \times 9 / 2 = 45$ cascades). The hourglass analysis is summarized in Figure B-14. The core nodes are the same with the single-source case, except that the anterior cingulate area – dorsal part (ACAd) and the Ectorhinal area (ECT) are replaced by the perirhinal area (PERI).

B-7 Comparison with other network analysis metrics

We first ask whether the path centrality metric correlates strongly with more commonly used node centrality metrics, namely: incoming or outgoing strength (the equivalent of “degree” for directed and weighted networks), betweenness centrality (fraction of all shortest paths traversing a node), closeness centrality (inversely related to average shortest path distance from that node to any other node), pagerank and eigenvector centrality (two related “influence” metrics that assign a higher score to a node that is connected to other highly-scored nodes compared to a node that has the same number of connections to low-scored nodes) [141]. Given that we are mostly interested in the dissemination of sensory

information over the network, one may expect that the pagerank and eigenvector centrality metrics would be more highly correlated with path centrality because such cascades do not necessarily follow shortest paths [173].

Table B-2: Top-five nodes according to different centrality metrics

Centrality metric	Top-five nodes
Out-strength	ENT1 - CA2 - PIR - FRP - EPv
In-strength	CA1 - MOs - ENT1 - DP - FRP
Betweenness	ENT1 - CLA - PERI - VIS1 - MOs
Closeness	ENT1 - PERI - CLA - PIR - VIS1
Pagerank	VIS1 - RSPagl - VISpm - FRP - RSPv
Eigenvector	MOs - FRP - AOB - CLA - MOp
Path Centrality	CLA - SSs - PTLp - AUDv - MOs

Table B-2 shows the top-five nodes in the network for each centrality metric. Figure B-15 shows scatter plots for the previous centrality metrics, comparing each of them with path centrality. The plots also show Kendall’s τ rank correlation coefficient. All centrality metrics are computed using *networkx* [174].

Strength (either for outgoing or incoming edges) only considers the local connections of each node; nodes in the hippocampal formation (CA1, CA2, ENT1) are among the most strongly connected, while none of the τ -core nodes ranks highly in terms of strength. The betweenness and closeness metrics are both based on shortest paths; the lateral entorhinal (ENT1) ranks highest in terms of that metric, while the claustrum (CLA) is the only τ -core node in the top-5 according to these two metrics. The highest ranked node based on pagerank is the lateral visual area (VIS1) while the highest ranked node based on eigenvector centrality is (by far) the secondary motor area (MOs).

Table B-3 shows the five rich-club nodes. The rich-club analysis is performed on unweighted and undirected networks and the rich-club coefficient is computed based on 1000 random surrogate networks, as described in [143]. The rich-club coefficient peaks at 1.15 for nodes with total degree over 23. Only two of the nodes in the rich-club overlap with the hourglass τ -core (CLA and MOs).

Table B-3: Rich-club nodes, Rombach core nodes nodes, and hourglass core nodes ($\tau=90\%$): the overlap between the first two sets with the hourglass core nodes is highlighted in red.

Mode	Nodes
Rich club (in order of total degree)	BLA, CLA , ENTI, MOs , PERI
Rombach (in order of Rombach score [146])	CA2, CA1, FRP, ENTI, MOs , EPv, PIR, MOp DP, COAa, NLOT, PAA, AOB
Hourglass Core for $\tau=90\%$ (in order of path coverage)	CLA, PTLp, AUDv, AOB SSs, MOs, ACAd, VISI, ECT

We also performed a core-periphery analysis on weighted but undirected networks using Rombach's method [146]. The core nodes according to this method are reported in Table B-3. Again, only two of the nodes in that core set overlap with the hourglass τ -core (MOs and AOB).

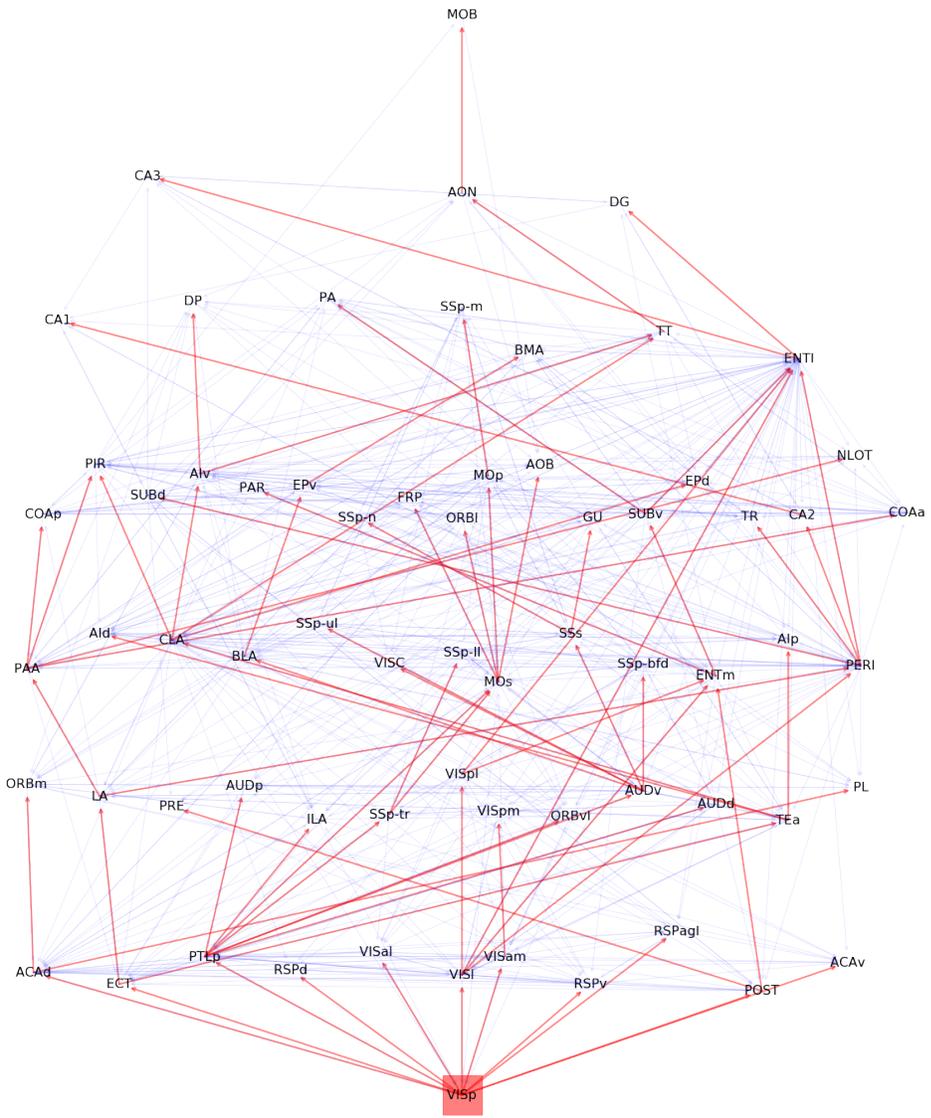


Figure B-2: *Visual cascade* (source: VISp)

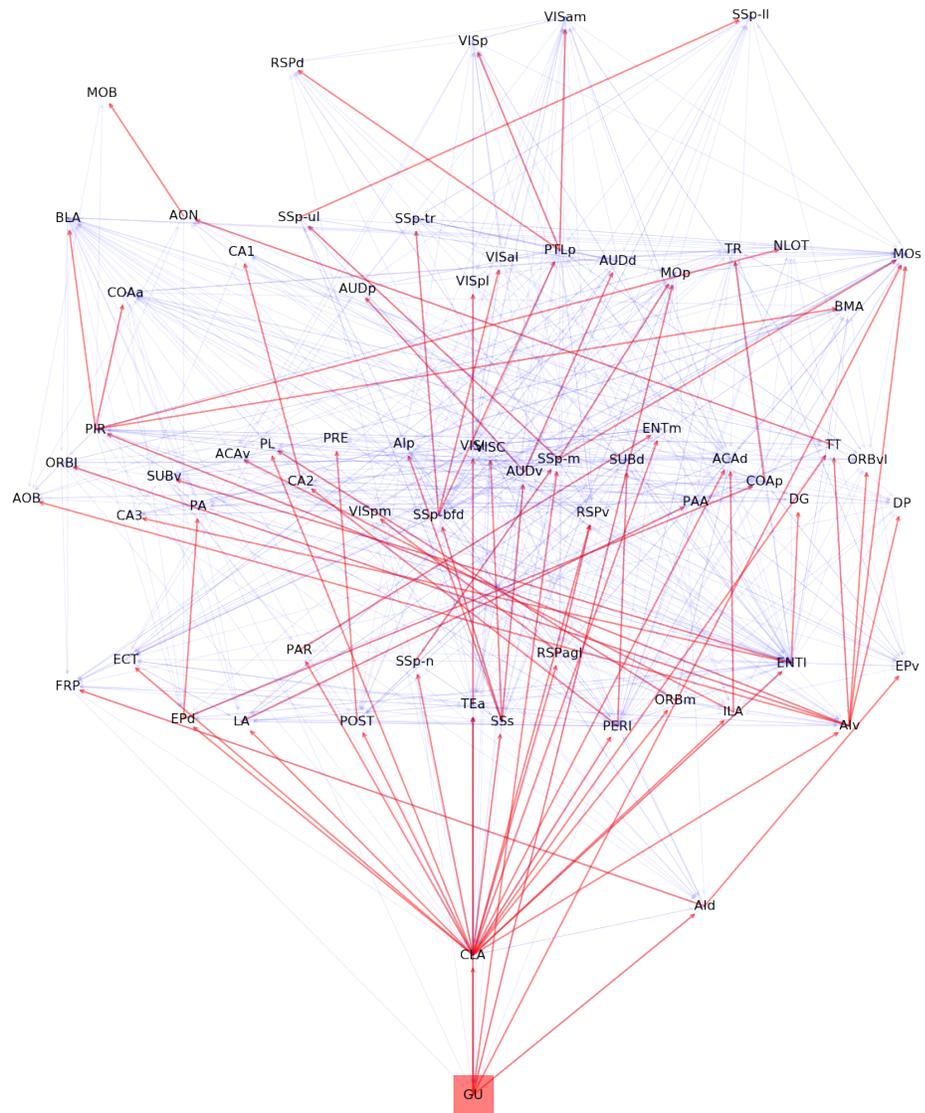


Figure B-4: *Gustatory cascade* (source: GU)

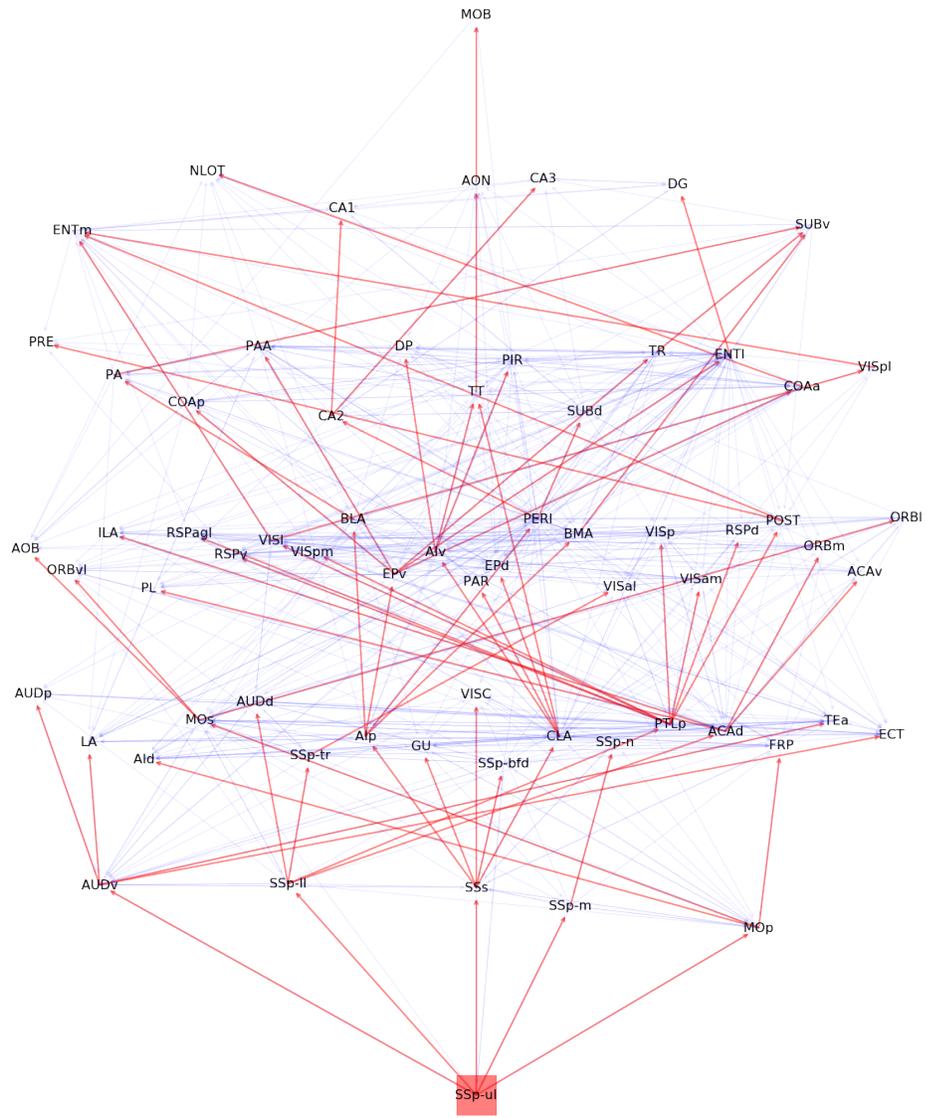


Figure B-5: *Upper-limb somatosensory cascade* (source: SSp-ul)

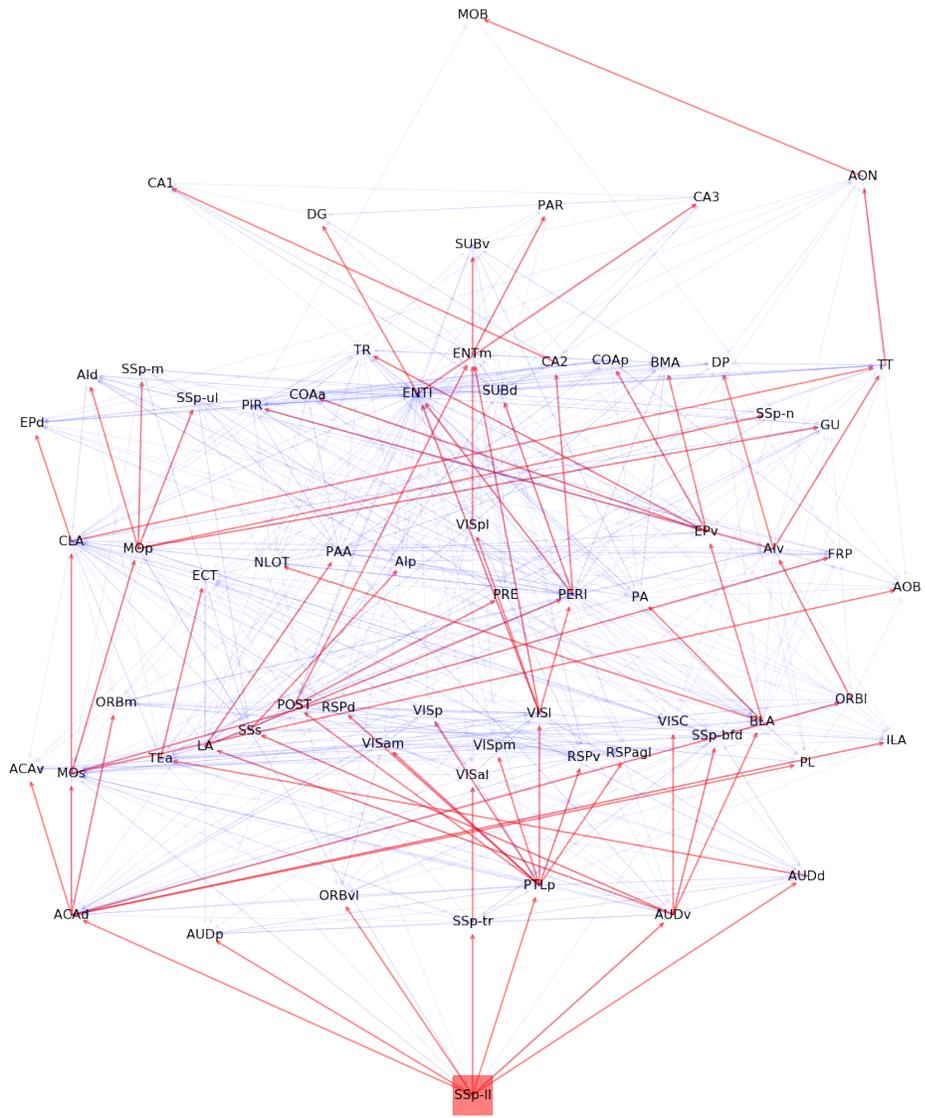


Figure B-6: Lower-limb somatosensory cascade (source: SSp-II)

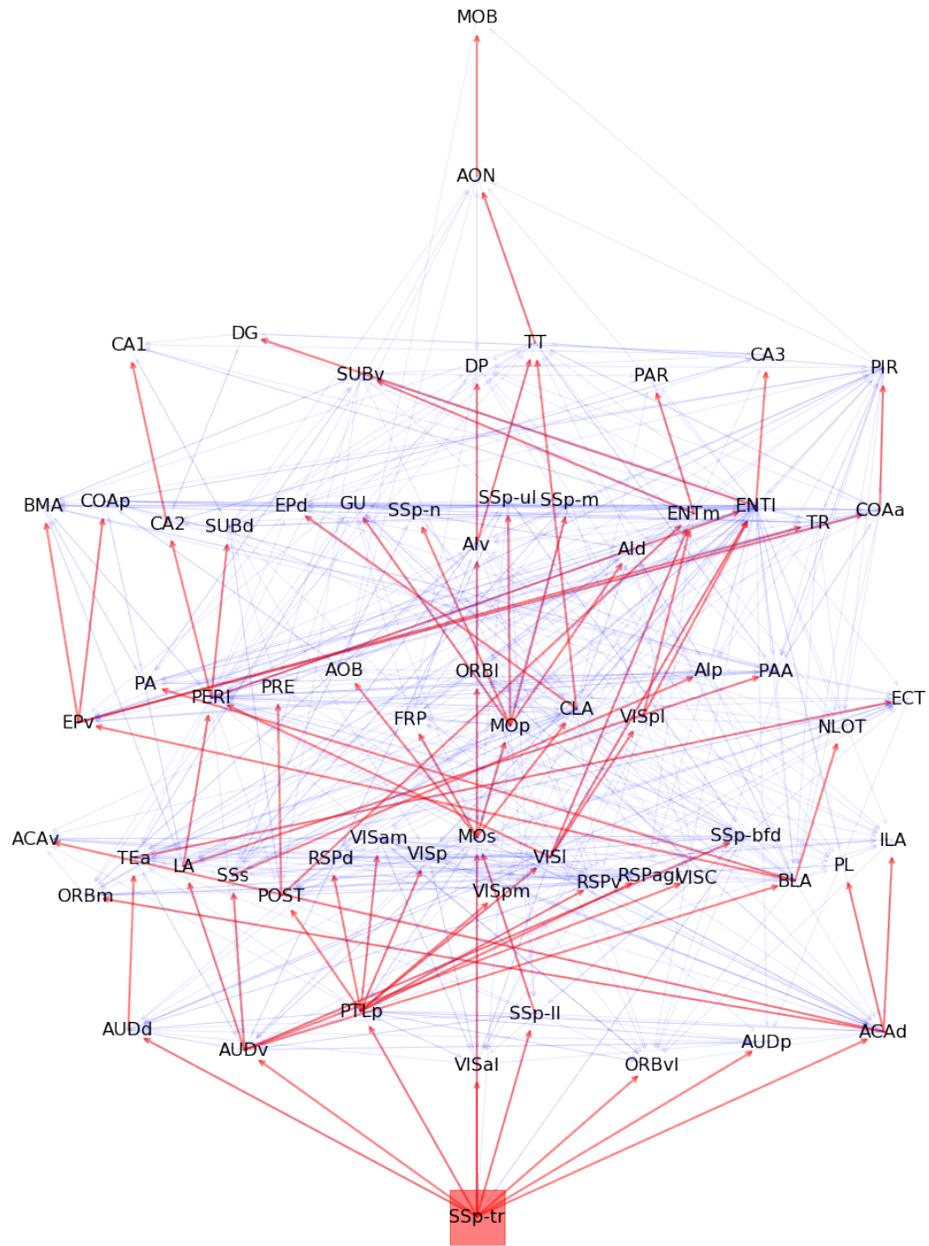


Figure B-8: *Trunk somatosensory cascade* (source: SSp-tr)

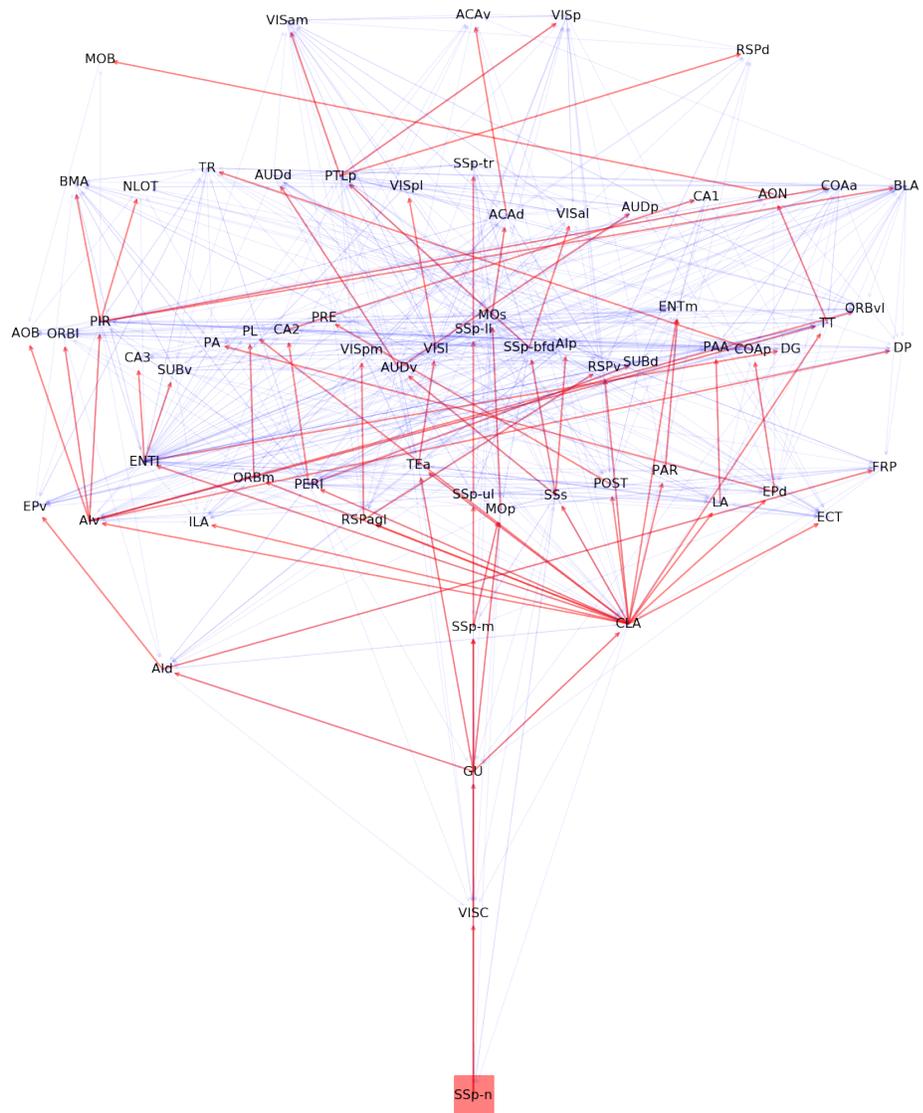
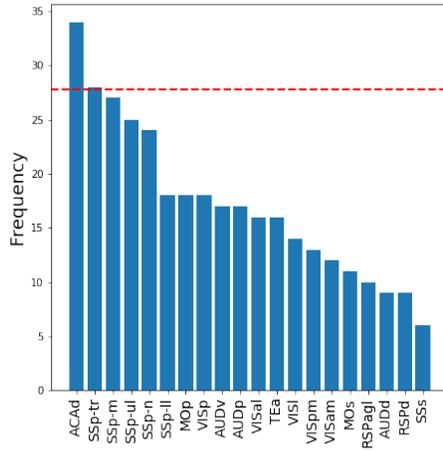
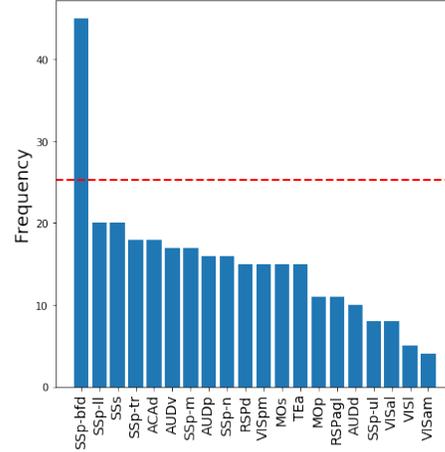


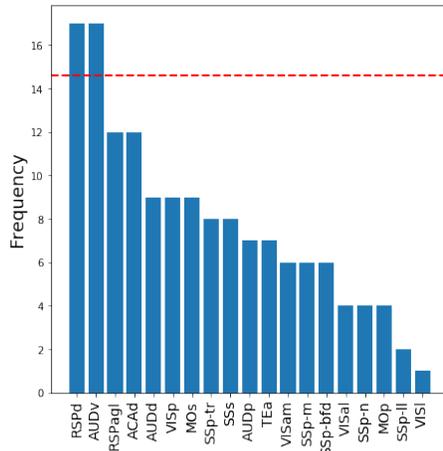
Figure B-10: *Nose somatosensory cascade* (source:SSp-n)



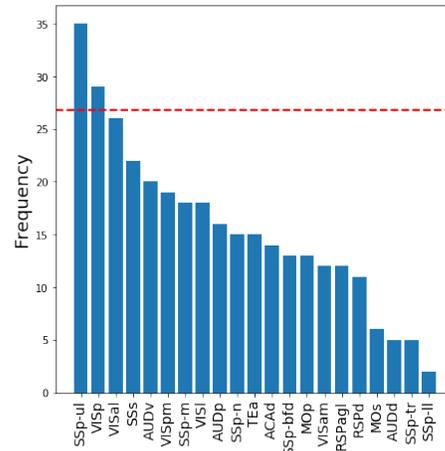
(a) Whisker cascade



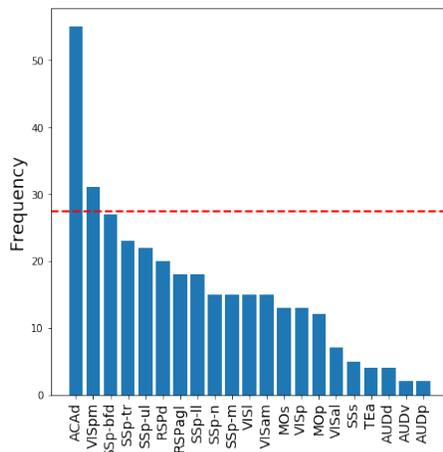
(b) Visual cascade



(c) Forelimb cascade



(d) Hindlimb cascade



(e) Tone cascade

Stimulus	First ROI	Second ROI
Visual	ACAd*	SSp-bfd
Whisker	ACAd*	AUDp
Tone	ACAd*	SSp-ul
Hindlimb	TEa*	ViSp

(f) Disagreement ROI pairs

Figure B-12: (a)-(e): The number of disagreements between VSD and ALT that involve each ROI, over the 21 ROIs that appear in the VSD data. The outlier ROIs have frequencies that exceed the red dashed line. (f) Disagreement ROI pairs that appear in all five animal datasets - each row corresponds to a different stimulus. ROIs at the boundary of the VSD visible cortical surface are marked by a star.

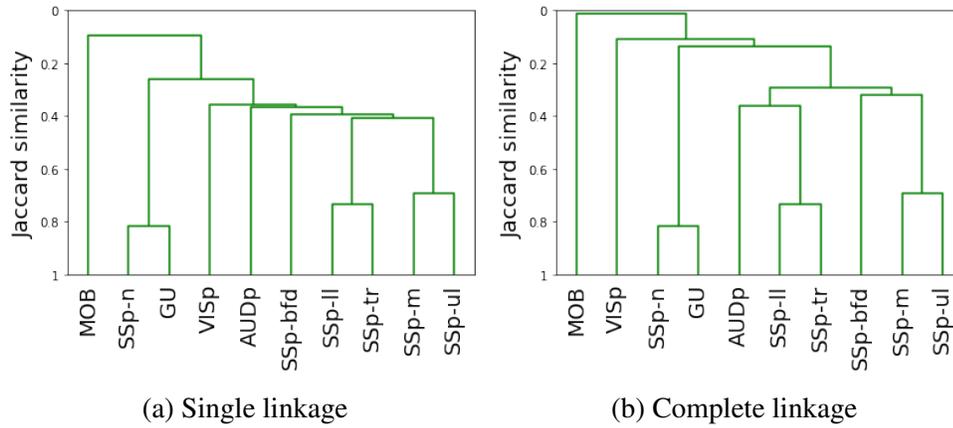
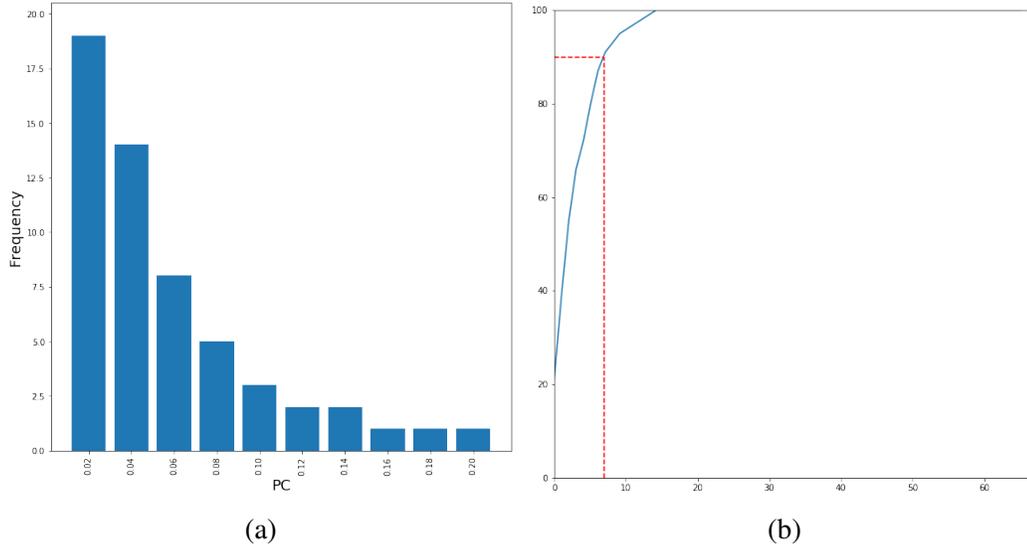


Figure B-13: *Similarity between activation cascades.*



Node	Description	Path Coverage	PC Rank
CLA	Clastrum	21	1
PTLp	Posterior parietal association areas	18	3
SSs	Supplemental somatosensory area	16	2
AOB	Accessory olfactory bulb	11	6
AUDv	Ventral auditory area	8	5
PERI	Perirhinal area	8	4
MOs	Secondary motor area	5	9
VISl	Lateral visual area	4	11

(c)

Figure B-14: *Simultaneous activation of two source nodes:* (a) Path Centrality (PC) histogram for the 67 cortical nodes, considering all source-target paths across the $10 \times 9/2 = 45$ activation cascades. (b) Cumulative path coverage by the top- X core nodes for $X=1 \dots 67$. (c) Eight nodes are enough to cover $\tau = 90\%$ of all paths.

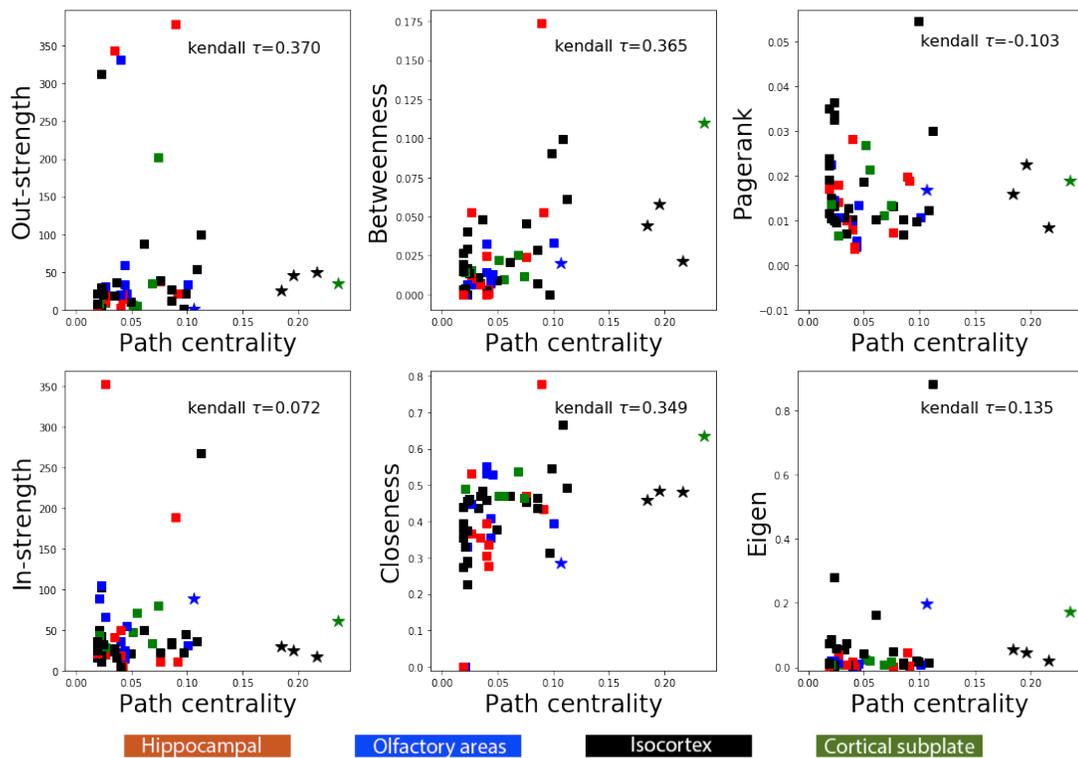


Figure B-15: *Path centrality compared to other centrality metrics.* Each point corresponds to one of the 67 nodes in the network. The four star nodes constitute the hourglass core for $\tau = 70\%$. All nodes are color-coded based on the broader brain region they belong to (isocortex, hippocampal formation, cortical subplate, olfactory areas).

REFERENCES

- [1] M. E. Newman, “The structure and function of complex networks,” *SIAM review*, vol. 45, no. 2, pp. 167–256, 2003.
- [2] A.-L. Barabási, *Linked: The new science of networks*, 2003.
- [3] A.-L. Barabasi and Z. N. Oltvai, “Network biology: Understanding the cell’s functional organization,” *Nature reviews genetics*, vol. 5, no. 2, p. 101, 2004.
- [4] N. Eagle, A. S. Pentland, and D. Lazer, “Inferring friendship network structure by using mobile phone data,” *Proceedings of the national academy of sciences*, vol. 106, no. 36, pp. 15 274–15 278, 2009.
- [5] I. Fountalis, C. Dovrolis, A. Bracco, B. Dilkina, and S. Keilholz, “ δ -maps: From spatio-temporal data to a weighted and lagged network between functional domains,” *Applied Network Science*, vol. 3, no. 1, p. 21, 2018.
- [6] O. Sporns, D. R. Chialvo, M. Kaiser, and C. C. Hilgetag, “Organization, development and function of complex brain networks,” *Trends in cognitive sciences*, vol. 8, no. 9, pp. 418–425, 2004.
- [7] O. Sporns, G. Tononi, and R. Kötter, “The human connectome: A structural description of the human brain,” *PLoS computational biology*, vol. 1, no. 4, e42, 2005.
- [8] H. Markram, “The human brain project,” *Scientific American*, vol. 306, no. 6, pp. 50–55, 2012.
- [9] S. E. Petersen and O. Sporns, “Brain networks and cognitive architectures,” *Neuron*, vol. 88, no. 1, pp. 207–219, 2015.
- [10] A. Fornito and E. T. Bullmore, “Connectomics: A new paradigm for understanding brain disease,” *European Neuropsychopharmacology*, vol. 25, no. 5, pp. 733–748, 2015.
- [11] A. K. Barbey, “Network neuroscience theory of human intelligence,” *Trends in cognitive sciences*, vol. 22, no. 1, pp. 8–20, 2018.
- [12] M. P. van den Heuvel, L. H. Scholtens, L. F. Barrett, C. C. Hilgetag, and M. A. de Reus, “Bridging cytoarchitectonics and connectomics in human cerebral cortex,” *Journal of Neuroscience*, vol. 35, no. 41, pp. 13 943–13 948, 2015.

- [13] J. G. White, E. Southgate, J. N. Thomson, and S. Brenner, “The structure of the nervous system of the nematode *caenorhabditis elegans*,” *Philos Trans R Soc Lond B Biol Sci*, vol. 314, no. 1165, pp. 1–340, 1986.
- [14] M. P. Van den Heuvel, E. T. Bullmore, and O. Sporns, “Comparative connectomics,” *Trends in cognitive sciences*, vol. 20, no. 5, pp. 345–361, 2016.
- [15] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H. Consortium, *et al.*, “The wu-minn human connectome project: An overview,” *Neuroimage*, vol. 80, pp. 62–79, 2013.
- [16] H. Markram, K. Meier, T. Lippert, S. Grillner, R. Frackowiak, S. Dehaene, A. Knoll, H. Sompolinsky, K. Verstreken, J. DeFelipe, *et al.*, “Introducing the human brain project,” *Procedia Computer Science*, vol. 7, pp. 39–42, 2011.
- [17] X. Chen, B. Errangi, L. Li, M. F. Glasser, L. T. Westlye, A. M. Fjell, K. B. Walhovd, X. Hu, J. G. Herndon, T. M. Preuss, and J. K. Rilling, “Brain aging in humans, chimpanzees (*pan troglodytes*), and rhesus macaques (*macaca mulatta*): Magnetic resonance imaging studies of macro- and microstructural changes,” *Neurobiology of Aging*, vol. 34, no. 10, pp. 2248–2260, Oct. 2013.
- [18] J. C. Beucke, J. Sepulcre, M. C. Eldaief, M. Sebold, N. Kathmann, and C. Kaufmann, “Default mode network subsystem alterations in obsessivecompulsive disorder,” *British Journal of Psychiatry*, vol. 205, no. 05, pp. 376–382, Nov. 2014.
- [19] I. Elman, D. Borsook, and N. D. Volkow, “Pain and suicidality: Insights from reward and addiction neuroscience,” *Progress in Neurobiology*, vol. 109, pp. 1–27, Oct. 2013.
- [20] A. Peterson, J. Thome, P. Frewen, and R. A. Lanius, “Resting-state neuroimaging studies: A new way of identifying differences and similarities among the anxiety disorders?” *The Canadian Journal of Psychiatry*, vol. 59, no. 6, pp. 294–300, Jun. 2014.
- [21] D. Seminowicz, H. Mayberg, A. McIntosh, K. Goldapple, S. Kennedy, Z. Segal, and S. Rafi-Tari, “Limbicfrontal circuitry in major depression: A path modeling metaanalysis,” *NeuroImage*, vol. 22, no. 1, pp. 409–418, May 2004.
- [22] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, *et al.*, “An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest,” *Neuroimage*, vol. 31, no. 3, pp. 968–980, 2006.
- [23] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, *et al.*, “A multi-modal

- parcellation of human cerebral cortex,” *Nature*, vol. 536, no. 7615, pp. 171–178, 2016.
- [24] A. Avena-Koenigsberger, B. Misic, and O. Sporns, “Communication dynamics in complex brain networks,” *Nature Reviews Neuroscience*, vol. 19, no. 1, p. 17, 2018.
- [25] M. A. de Reus and M. P. van den Heuvel, “Simulated rich club lesioning in brain networks: A scaffold for communication and integration?” *Frontiers in human neuroscience*, vol. 8, p. 647, 2014.
- [26] S. W. Oh, J. A. Harris, L. Ng, B. Winslow, N. Cain, S. Mihalas, Q. Wang, C. Lau, L. Kuan, A. M. Henry, M. T. Mortrud, B. Ouellette, T. N. Nguyen, S. A. Sorensen, C. R. Slaughterbeck, W. Wakeman, Y. Li, D. Feng, A. Ho, E. Nicholas, K. E. Hirokawa, P. Bohn, K. M. Joines, H. Peng, M. J. Hawrylycz, J. W. Phillips, J. G. Hohmann, P. Wohnoutka, C. R. Gerfen, C. Koch, A. Bernard, C. Dang, A. R. Jones, and H. Zeng, “A mesoscale connectome of the mouse brain,” *Nature*, vol. 508, no. 7495, pp. 207–214, 2014. arXiv: NIHMS150003.
- [27] M. H. Mohajerani, A. W. Chan, M. Mohsenvand, J. LeDue, R. Liu, D. A. McVea, J. D. Boyd, Y. T. Wang, M. Reimers, and T. H. Murphy, “Spontaneous cortical activity alternates between motifs defined by regional axonal projections,” *Nature neuroscience*, vol. 16, no. 10, p. 1426, 2013.
- [28] K. Shadi, S. Bakhshi, D. A. Gutman, H. S. Mayberg, and C. Dvornik, “A symmetry-based method to infer structural brain networks from probabilistic tractography data,” *Frontiers in neuroinformatics*, vol. 10, p. 46, 2016.
- [29] O. Sporns, G. Tononi, and R. Ktter, “The human connectome: A structural description of the human brain,” *PLoS Computational Biology*, vol. 1, no. 4, e42, 2005.
- [30] D. C. V. Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, and K. Ugurbil, “The wu-minn human connectome project: An overview,” *NeuroImage*, vol. 80, pp. 62–79, Oct. 2013.
- [31] S. Olaf, *Discovering the human connectome*. MIT press, 2012.
- [32] O. Sporns, “Making sense of brain network data,” *Nature Methods*, vol. 10, no. 6, pp. 491–493, Jun. 2013.
- [33] O. Ciccarelli, M. Catani, H. Johansen-Berg, C. Clark, and A. Thompson, “Diffusion-based tractography in neurological disorders: Concepts, applications, and future developments,” *The Lancet Neurology*, vol. 7, no. 8, pp. 715–727, Aug. 2008.

- [34] A. Fornito and E. T. Bullmore, “Connectomics: A new paradigm for understanding brain disease,” *European Neuropsychopharmacology*, vol. 25, no. 5, pp. 733–748, May 2015.
- [35] H. S. Mayberg, A. M. Lozano, V. Voon, H. E. McNeely, D. Seminowicz, C. Hamani, J. M. Schwalb, and S. H. Kennedy, “Deep brain stimulation for treatment-resistant depression,” *Neuron*, vol. 45, no. 5, pp. 651–660, Mar. 2005.
- [36] D. S. Bassett, E. Bullmore, B. A. Verchinski, V. S. Mattay, D. R. Weinberger, and A. Meyer-Lindenberg, “Hierarchical organization of human cortical networks in health and schizophrenia,” *Journal of Neuroscience*, vol. 28, no. 37, pp. 9239–9248, Sep. 10.
- [37] R. L. Buckner, A. Z. Snyder, B. J. Shannon, G. LaRossa, R. Sachs, A. F. Fotenos, Y. I. Sheline, W. E. Klunk, C. A. Mathis, J. C. Morris, *et al.*, “Molecular, structural, and functional characterization of alzheimer’s disease: Evidence for a relationship between default activity, amyloid, and memory,” *Journal of Neuroscience*, vol. 25, pp. 7709–7717, 34 2005.
- [38] S. Mori and P. C. M. van Zijl, “Fiber tracking: Principles and strategies - a technical review,” *NMR in Biomedicine*, vol. 15, no. 7-8, pp. 468–480, 2002.
- [39] T. Behrens, H. J. Berg, S. Jbabdi, M. Rushworth, and M. Woolrich, “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?” *NeuroImage*, vol. 34, no. 1, pp. 144–155, Jan. 2007.
- [40] S. Mori, B. J. Crain, V. P. Chacko, and P. C.M. V. Zijl, “Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging,” *Annals of Neurology*, vol. 45, no. 2, pp. 265–269, Feb. 1999.
- [41] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, V. J. Wedeen, R. Meuli, and J.-P. Thiran, “Mapping human whole-brain structural networks with diffusion MRI,” *PLoS ONE*, vol. 2, no. 7, O. Sporns, Ed., e597, Jul. 4.
- [42] M. Daianu, N. Jahanshad, T. M. Nir, A. W. Toga, C. R. Jack Jr, M. W. Weiner, and P. M. Thompson for the Alzheimer’s Disease Neuroimaging Initiative, “Breakdown of brain connectivity between normal aging and alzheimer’s disease: A structural k-core network analysis,” *Brain connectivity*, pp. 407–422, 2013.
- [43] A. M. McIntosh, S. M. Maniega, G. K. S. Lymer, J. McKirdy, J. Hall, J. E. Sussmann, M. E. Bastin, J. D. Clayden, E. C. Johnstone, and S. M. Lawrie, “White matter tractography in bipolar disorder and schizophrenia,” *Biological Psychiatry*, vol. 64, no. 12, pp. 1088–1092, Dec. 2008.

- [44] J. S. Damoiseaux and M. D. Greicius, “Greater than the sum of its parts: A review of studies combining structural connectivity and resting-state functional connectivity,” *Brain Structure and Function*, vol. 213, no. 6, pp. 525–533, Jun. 30.
- [45] J. T. Duda, P. A. Cook, and J. C. Gee, “Reproducibility of graph metrics of human brain structural networks,” *Frontiers in Neuroinformatics*, vol. 8, May 7.
- [46] L. L., R. JK, P. TM, G. MF, and H. X, “The effects of connection reconstruction method on the interregional connectivity of brain networks via diffusion tractography,” *Human brain mapping*, vol. 33, 21928316, pp. 1894–913, 8 2012.
- [47] B. C. M. van Wijk, C. J. Stam, and A. Daffertshofer, “Comparing brain networks of different size and connectivity density using graph theory,” *PLoS ONE*, vol. 5, no. 10, O. Sporns, Ed., e13701, Oct. 28.
- [48] G. Gong, Y. He, L. Concha, C. Lebel, D. W. Gross, A. C. Evans, and C. Beaulieu, “Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography,” *Cerebral Cortex*, vol. 19, no. 3, pp. 524–536, Jun. 20.
- [49] D. K. Jones, T. R. Knösche, and R. Turner, “White matter integrity, fiber count, and other fallacies: The do’s and don’ts of diffusion mri,” *Neuroimage*, vol. 73, Elsevier 2013.
- [50] Y. Li, Y. Liu, J. Li, W. Qin, K. Li, C. Yu, and T. Jiang, “Brain anatomical network and intelligence,” *PLoS Computational Biology*, vol. 5, no. 5, O. Sporns, Ed., e1000395, May 29.
- [51] S. Jbabdi and H. Johansen-Berg, “Tractography: Where do we go from here?” *Brain Connectivity*, vol. 1, no. 3, pp. 169–183, Sep. 2011.
- [52] E. C. Robinson, A. Hammers, A. Ericsson, A. D. Edwards, and D. Rueckert, “Identifying population differences in whole-brain structural networks: A machine learning approach,” *NeuroImage*, vol. 50, no. 3, pp. 910–919, Apr. 2010.
- [53] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns, “Mapping the structural core of human cerebral cortex,” *PLoS Biology*, vol. 6, no. 7, K. J. Friston, Ed., e159, Jul. 1.
- [54] G. J. Parker, H. A. Haroon, and C. A. Wheeler-Kingshott, “A framework for a streamline-based probabilistic index of connectivity (pico) using a structural interpretation of MRI diffusion measurements,” *Journal of Magnetic Resonance Imaging*, vol. 18, no. 2, pp. 242–254, Jul. 21.

- [55] M. Descoteaux, R. Deriche, T. Knosche, and A. Anwander, “Deterministic and probabilistic tractography based on complex fibre orientation distributions,” *IEEE Transactions on Medical Imaging*, vol. 28, no. 2, pp. 269–286, Feb. 2009.
- [56] C. Reveley, A. K. Seth, C. Pierpaoli, A. C. Silva, D. Yu, R. C. Saunders, D. A. Leopold, and F. Q. Ye, “Superficial white matter fiber systems impede detection of long-range cortical connections in diffusion MR tractography,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 21, E2820–E2828, May 11.
- [57] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, “Characterization and propagation of uncertainty in diffusion-weighted MR imaging,” *Magnetic Resonance in Medicine*, vol. 50, no. 5, pp. 1077–1088, Oct. 24.
- [58] M. Jenkinson and S. Smith, “A global optimisation method for robust affine registration of brain images,” *Medical Image Analysis*, vol. 5, no. 2, pp. 143–156, Jun. 2001.
- [59] J. L. Lancaster, M. G. Woldorff, L. M. Parsons, M. Liotti, C. S. Freitas, L. Rainey, P. V. Kochunov, D. Nickerson, S. A. Mikiten, and P. T. Fox, “Automated talairach atlas labels for functional brain mapping,” *Human Brain Mapping*, vol. 10, no. 3, pp. 120–131, 2000.
- [60] R. C. Craddock, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg, “Disease state prediction from resting state functional connectivity,” *Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1619–1628, Dec. 2009.
- [61] G. A. James, M. E. Kelley, R. C. Craddock, P. E. Holtzheimer, B. W. Dunlop, C. B. Nemeroff, H. S. Mayberg, and X. P. Hu, “Exploratory structural equation modeling of resting-state fmri: Applicability of group models to individual subjects,” *NeuroImage*, vol. 45, no. 3, pp. 778–787, Apr. 2009.
- [62] M. A. de Reus and M. P. van den Heuvel, “The parcellation-based connectome: Limitations and extensions,” *NeuroImage*, vol. 80, pp. 397–404, Oct. 2013.
- [63] A. Zalesky, A. Fornito, I. H. Harding, L. Cocchi, M. Ycel, C. Pantelis, and E. T. Bullmore, “Whole-brain anatomical networks: Does the choice of nodes matter?” *NeuroImage*, vol. 50, no. 3, pp. 970–983, Apr. 2010.
- [64] M. Petrides, “Lateral prefrontal cortex: Architectonic and functional organization,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 360, no. 1456, pp. 781–795, Apr. 29.
- [65] N. Tzourio, L. Petit, E. Mellet, C. Orssaud, F. Crivello, K. Benali, G. Salamon, and B. Mazoyer, “Use of anatomical parcellation to catalog and study structure-

function relationships in the human brain,” *Human Brain Mapping*, vol. 5, no. 4, pp. 228–232, 1997.

- [66] T. Blumensath, S. Jbabdi, M. F. Glasser, D. C. V. Essen, K. Ugurbil, T. E. Behrens, and S. M. Smith, “Spatially constrained hierarchical parcellation of the brain with resting-state fmri,” *NeuroImage*, vol. 76, pp. 313–324, Aug. 2013.
- [67] R. C. Craddock, G. James, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg, “A whole brain fmri atlas generated via spatially constrained spectral clustering,” *Human Brain Mapping*, vol. 33, no. 8, pp. 1914–1928, Jul. 18.
- [68] S. Makeig, G. Brown, S. Kindermann, T.-P. Jung, A. Bell, T. Sejnowski, and M. McKeown, “Response from martin mckeown, makeig, brown, jung, kindermann, bell and sejnowski,” *Trends in Cognitive Sciences*, vol. 2, no. 10, p. 375, Oct. 1998.
- [69] B. Thirion, G. Varoquaux, E. Dohmatob, and J.-B. Poline, “Which fmri clustering gives good brain parcellations?” *Frontiers in Neuroscience*, vol. 8, Jul. 1.
- [70] D. S. Bassett, J. A. Brown, V. Deshpande, J. M. Carlson, and S. T. Grafton, “Conserved and variable architecture of human white matter connectivity,” *NeuroImage*, vol. 54, no. 2, pp. 1262–1279, Jan. 2011.
- [71] F. Schalekamp and A. v. Zuylen, “Rank aggregation: Together we’re strong,” in *2009 Proceedings of the Eleventh Workshop on Algorithm Engineering and Experiments (ALENEX)*, 2009, pp. 38–51.
- [72] N. Ailon, M. Charikar, and A. Newman, “Aggregating inconsistent information, Ranking and clustering,” *Journal of the ACM*, vol. 55, no. 5, pp. 1–27, Oct. 1.
- [73] C. Thomas, F. Q. Ye, M. O. Irfanoglu, P. Modi, K. S. Saleem, D. A. Leopold, and C. Pierpaoli, “Anatomical accuracy of brain connections derived from diffusion MRI tractography is inherently limited,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 46, pp. 16 574–16 579, Nov. 3.
- [74] L. Li, J. K. Rilling, T. M. Preuss, M. F. Glasser, F. W. Damen, and X. Hu, “Quantitative assessment of a framework for creating anatomical brain networks via global tractography,” *NeuroImage*, vol. 61, no. 4, pp. 1017–1030, Jul. 2012.
- [75] D. M. Morris, K. V. Embleton, and G. J. Parker, “Probabilistic fibre tracking: Differentiation of connections from chance events,” *NeuroImage*, vol. 42, no. 4, pp. 1329–1339, Oct. 2008.
- [76] K. Taljan, C. McIntyre, and K. Sakaie, “Anatomical connectivity between subcortical structures,” *Brain Connectivity*, vol. 1, no. 2, pp. 111–118, Aug. 2011.

- [77] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot, “Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain,” *NeuroImage*, vol. 15, no. 1, pp. 273–289, Jan. 2002.
- [78] M. van den Heuvel, C. Stam, M. Boersma, and H. H. Pol, “Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain,” *NeuroImage*, vol. 43, no. 3, pp. 528–539, Nov. 2008.
- [79] J. Power, A. Cohen, S. Nelson, G. Wig, K. Barnes, J. Church, A. Vogel, T. Laumann, F. Miezin, B. Schlaggar, and S. Petersen, “Functional network organization of the human brain,” *Neuron*, vol. 72, no. 4, pp. 665–678, Nov. 2011.
- [80] B. T. T. Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zillei, J. R. Polimeni, B. Fischl, H. Liu, and R. L. Buckner, “The organization of the human cerebral cortex estimated by intrinsic functional connectivity,” *Journal of Neurophysiology*, vol. 106, no. 3, pp. 1125–1165, Sep. 2011.
- [81] R. L. Buckner, J. R. Andrews-Hanna, and D. L. Schacter, “The brain’s default network,” *Annals of the New York Academy of Sciences*, 2008.
- [82] M. Kleinnijenhuis, T. van Mourik, D. G. Norris, D. J. Ruiteer, A.-M.v. C. van Walsum, and M. Barth, “Diffusion tensor characteristics of gyrencephaly using high resolution diffusion mri in vivo at 7t,” *Neuroimage*, vol. 109, pp. 378–387, 2015.
- [83] M. Newman, *Networks*. Oxford University Press, Mar. 25, ISBN: 9780199206650.
- [84] M. D. Greicius, K. Supekar, V. Menon, and R. F. Dougherty, “Resting-state functional connectivity reflects structural connectivity in the default mode network,” *Cerebral Cortex*, vol. 19, no. 1, pp. 72–78, Apr. 9.
- [85] M. Petrides and D. N. Pandya, “Efferent association pathways from the rostral prefrontal cortex in the macaque monkey,” *Journal of Neuroscience*, vol. 27, no. 43, pp. 11 573–11 586, Oct. 24.
- [86] M. P. Van Den Heuvel and O. Sporns, “Rich-club organization of the human connectome,” *Journal of Neuroscience*, vol. 31, no. 44, pp. 15 775–15 786, 2011.
- [87] C. Destrieux, B. Fischl, A. Dale, and E. Halgren, “Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature,” *Neuroimage*, vol. 53, no. 1, pp. 1–15, 2010.

- [88] D. C. Van Essen, C. J. Donahue, and M. F. Glasser, “Development and evolution of cerebral and cerebellar cortex,” *Brain, behavior and evolution*, vol. 91, pp. 158–169, 2018.
- [89] K. Schilling, Y. Gao, V. Janve, I. Stepniewska, B. A. Landman, and A. W. Anderson, “Confirmation of a gyral bias in diffusion mri fiber tractography,” *Human brain mapping*, vol. 39, no. 3, pp. 1449–1466, 2018.
- [90] J. A. Roberts, A. Perry, A. R. Lord, G. Roberts, P. B. Mitchell, R. E. Smith, F. Calamante, and M. Breakspear, “The contribution of geometry to the human connectome,” *Neuroimage*, vol. 124, pp. 379–393, 2016.
- [91] C. J. Donahue, S. N. Sotiropoulos, S. Jbabdi, M. Hernandez-Fernandez, T. E. Behrens, T. B. Dyrby, T. Coalson, H. Kennedy, K. Knoblauch, D. C. Van Essen, *et al.*, “Using diffusion tractography to predict cortical connection strength and distance: A quantitative comparison with tracers in the monkey,” *Journal of Neuroscience*, vol. 36, no. 25, pp. 6758–6770, 2016.
- [92] M. P. Van den Heuvel and O. Sporns, “An anatomical substrate for integration among functional networks in human cortex,” *Journal of Neuroscience*, vol. 33, no. 36, pp. 14 489–14 500, 2013.
- [93] R. F. Betzel and D. S. Bassett, “Specificity and robustness of long-distance connections in weighted, interareal connectomes,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 21, E4880–E4889, 2018.
- [94] D. W. Scott, *Multivariate density estimation: Theory, practice, and visualization*. John Wiley & Sons, 2015.
- [95] K. Uğurbil, J. Xu, E. J. Auerbach, S. Moeller, A. T. Vu, J. M. Duarte-Carvajalino, C. Lenglet, X. Wu, S. Schmitter, P. F. Van de Moortele, *et al.*, “Pushing spatial and temporal resolution for functional and diffusion mri in the human connectome project,” *Neuroimage*, vol. 80, pp. 80–104, 2013.
- [96] T. Sarwar, K. Ramamohanarao, and A. Zalesky, “Mapping connectomes with diffusion mri: Deterministic or probabilistic tractography?” *Magnetic resonance in medicine*, vol. 81, no. 2, pp. 1368–1384, 2019.
- [97] B. Mišić, R. F. Betzel, A. Griffa, M. A. de Reus, Y. He, X.-N. Zuo, M. P. van den Heuvel, P. Hagmann, O. Sporns, and R. J. Zatorre, “Network-based asymmetry of the human auditory system,” *Cerebral Cortex*, vol. 28, no. 7, pp. 2655–2664, 2018.
- [98] M. Bastiani, N. J. Shah, R. Goebel, and A. Roebroeck, “Human cortical connectome reconstruction from diffusion weighted mri: The effect of tractography algorithm,” *Neuroimage*, vol. 62, no. 3, pp. 1732–1749, 2012.

- [99] M. Rubinov and O. Sporns, “Complex network measures of brain connectivity: Uses and interpretations,” *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.
- [100] D. J. Schaeffer, K. D. Johnston, K. M. Gilbert, J. S. Gati, R. S. Menon, and S. Everling, “In vivo manganese tract tracing of frontal eye fields in rhesus macaques with ultra-high field mri: Comparison with dwi tractography,” *NeuroImage*, vol. 181, pp. 211–218, 2018.
- [101] C. Thomas, Q. Y. Frank, M. O. Irfanoglu, P. Modi, K. S. Saleem, D. A. Leopold, and C. Pierpaoli, “Anatomical accuracy of brain connections derived from diffusion mri tractography is inherently limited,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 46, pp. 16 574–16 579, 2014.
- [102] K. H. Maier-Hein, P. F. Neher, J.-C. Houde, M.-A. Côté, E. Garyfallidis, J. Zhong, M. Chamberland, F.-C. Yeh, Y.-C. Lin, Q. Ji, *et al.*, “The challenge of mapping the human connectome based on diffusion tractography,” *Nature communications*, vol. 8, no. 1, p. 1349, 2017.
- [103] R. Gămănuț, H. Kennedy, Z. Toroczka, M. Ercsey-Ravasz, D. C. Van Essen, K. Knoblauch, and A. Burkhalter, “The mouse cortical connectome, characterized by an ultra-dense cortical graph, maintains specificity by distinct connectivity profiles,” *Neuron*, vol. 97, no. 3, pp. 698–715, 2018.
- [104] K. Shen, A. Goulas, D. S. Grayson, J. Eusebio, J. S. Gati, R. S. Menon, A. R. McIntosh, and S. Everling, “Exploring the limits of network topology estimation using diffusion-based tractography and tracer studies in the macaque cortex,” *NeuroImage*, vol. 191, pp. 81–92, 2019.
- [105] A. Zalesky, A. Fornito, L. Cocchi, L. L. Gollo, M. P. van den Heuvel, and M. Breakspear, “Connectome sensitivity or specificity: Which is more important?” *Neuroimage*, vol. 142, pp. 407–420, 2016.
- [106] H. Azadbakht, L. M. Parkes, H. A. Haroon, M. Augath, N. K. Logothetis, A. de Crespigny, H. E. D’arceuil, and G. J. Parker, “Validation of high-resolution tractography against in vivo tracing in the macaque visual cortex,” *Cerebral cortex*, vol. 25, no. 11, pp. 4299–4309, 2015.
- [107] K. Shen, A. Goulas, D. Grayson, J. Eusebio, J. S. Gati, R. S. Menon, A. R. McIntosh, and S. Everling, “Exploring the limits of network topology estimation using diffusion-based tractography and tracer studies in the macaque cortex,” *BioRxiv*, p. 356 576, 2018.
- [108] M. P. van den Heuvel, M. A. de Reus, L. Feldman Barrett, L. H. Scholtens, F. M. Coopmans, R. Schmidt, T. M. Preuss, J. K. Rilling, and L. Li, “Comparison of

diffusion tractography and tract-tracing measures of connectivity strength in rhesus macaque connectome,” *Human brain mapping*, vol. 36, no. 8, pp. 3064–3075, 2015.

- [109] B. E. Stein and M. A. Meredith, *The merging of the senses*. The MIT Press, 1993.
- [110] R. L. Seilheimer, A. Rosenberg, and D. E. Angelaki, “Models and processes of multisensory cue combination,” *Current opinion in neurobiology*, vol. 25, pp. 38–46, 2014.
- [111] M. A. Meredith, “On the neuronal basis for multisensory convergence: A brief overview,” *Cognitive brain research*, vol. 14, no. 1, pp. 31–40, 2002.
- [112] J. M. Shine, M. Breakspear, P. T. Bell, K. A. E. Martens, R. Shine, O. Koyejo, O. Sporns, and R. A. Poldrack, “Human cognition involves the dynamic integration of neural activity and neuromodulatory systems,” *Nature neuroscience*, vol. 22, no. 2, p. 289, 2019.
- [113] A. A. Ghazanfar and C. E. Schroeder, “Is neocortex essentially multisensory?” *Trends in cognitive sciences*, vol. 10, no. 6, pp. 278–285, 2006.
- [114] A. Barrat, M. Barthelemy, and A. Vespignani, *Dynamical processes on complex networks*. Cambridge university press, 2008.
- [115] J. C. Worrell, J. Rumschlag, R. F. Betzel, O. Sporns, and B. Mišić, “Optimized connectome architecture for sensory-motor integration,” *Network Neuroscience*, vol. 1, no. 4, pp. 415–430, 2017.
- [116] B. N. Mathur, “The claustrum in review,” *Frontiers in systems neuroscience*, vol. 8, p. 48, 2014.
- [117] K. M. Sabrin and C. Dvornik, “The hourglass effect in hierarchical dependency networks,” *Network Science*, vol. 5, no. 4, pp. 490–528, 2017.
- [118] L. K. Gallos, H. A. Makse, and M. Sigman, “A small world of weak ties provides optimal global integration of self-similar modules in functional brain networks,” *Proceedings of the National Academy of Sciences*, vol. 109, no. 8, pp. 2825–2830, 2012.
- [119] B. Mota and S. Herculano-Houzel, “How the cortex gets its folds: An inside-out, connectivity-driven model for the scaling of mammalian cortical folding,” *Frontiers in neuroanatomy*, vol. 6, p. 3, 2012.
- [120] T. Sun and R. F. Hevner, “Growth and folding of the mammalian cerebral cortex: From molecules to malformations,” *Nature Reviews Neuroscience*, vol. 15, no. 4, p. 217, 2014.

- [121] C. Stringer, M. Pachitariu, N. Steinmetz, C. B. Reddy, M. Carandini, and K. D. Harris, “Spontaneous behaviors drive multidimensional, brain-wide population activity,” *BioRxiv*, p. 306 019, 2018.
- [122] G. T. Meijer, P. E. Mertens, C. M. Pennartz, U. Olcese, and C. S. Lansink, “The circuit architecture of cortical multisensory processing: Distinct functions jointly operating within a common anatomical network,” *Progress in neurobiology*, 2019.
- [123] G. Fagiolo, “Clustering in complex directed networks,” *Physical Review E*, vol. 76, no. 2, p. 026 107, 2007.
- [124] M. Bota, O. Sporns, and L. W. Swanson, “Architecture of the cerebral cortical association connectome underlying cognition,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 16, E2093–E2101, 2015.
- [125] P Taylor, J. Hobbs, J Burroni, and H. Siegelmann, “The global landscape of cognition: Hierarchical aggregation as an organizational principle of human cortical networks and functions,” *Scientific reports*, vol. 5, p. 18 112, 2015.
- [126] M. Granovetter, “Threshold models of collective behavior,” *American journal of sociology*, vol. 83, no. 6, pp. 1420–1443, 1978.
- [127] P. Sanz-Leon, S. A. Knock, A. Spiegler, and V. K. Jirsa, “Mathematical framework for large-scale brain network modeling in the virtual brain,” *Neuroimage*, vol. 111, pp. 385–430, 2015.
- [128] M. Newman, *Networks*. Oxford university press, 2018.
- [129] J. M. Beggs and N. Timme, “Being critical of criticality in the brain,” *Frontiers in physiology*, vol. 3, p. 163, 2012.
- [130] C. Cappe and P. Barone, “Heteromodal connections supporting multisensory integration at low levels of cortical processing in the monkey,” *European Journal of Neuroscience*, vol. 22, no. 11, pp. 2886–2902, 2005.
- [131] A. Falchier, S. Clavagnier, P. Barone, and H. Kennedy, “Anatomical evidence of multimodal integration in primate striate cortex,” *Journal of Neuroscience*, vol. 22, no. 13, pp. 5749–5759, 2002.
- [132] R. J. Morrill and A. R. Hasenstaub, “Visual information present in infragranular layers of mouse auditory cortex,” *Journal of Neuroscience*, pp. 3102–17, 2018.
- [133] D. Plenz and T. C. Thiagarajan, “The organizing principles of neuronal avalanches: Cell assemblies in the cortex?” *Trends in neurosciences*, vol. 30, no. 3, pp. 101–110, 2007.

- [134] N. T. Markov, M. Ercsey-Ravasz, D. C. Van Essen, K. Knoblauch, Z. Toroczkai, and H. Kennedy, “Cortical high-density counterstream architectures,” *Science*, vol. 342, no. 6158, p. 1 238 406, 2013.
- [135] J. A. Harris, S. Mihalas, K. E. Hirokawa, J. D. Whitesell, J. Knox, A. Bernard, P. Bohn, S. Caldejon, L. Casal, A. Cho, *et al.*, “The organization of intracortical connections by layer and cell class in the mouse brain,” *BioRxiv*, p. 292 961, 2018.
- [136] D. M. Johnson, K. R. Illig, M. Behan, and L. B. Haberly, “New features of connectivity in piriform cortex visualized by intracellular injection of pyramidal cells suggest that primary olfactory cortex functions like association cortex in other sensory systems,” *Journal of Neuroscience*, vol. 20, no. 18, pp. 6974–6982, 2000.
- [137] S. Srinivasan and C. F. Stevens, “A quantitative description of the mouse piriform cortex,” *BioRxiv*, p. 099 002, 2017.
- [138] N. Cayco-Gajic and Y Sweeney, “Delving deep into crossmodal integration,” *The Journal of neuroscience: The official journal of the Society for Neuroscience*, vol. 38, no. 29, pp. 6442–6444, 2018.
- [139] A. Griffa and M. P. Van den Heuvel, “Rich-club neurocircuitry: Function, evolution, and vulnerability,” *Dialogues in Clinical Neuroscience*, vol. 20, no. 2, p. 121, 2018.
- [140] S. Oldham, B. Fulcher, L. Parkes, A. Arnatkeviciute, C. Suo, and A. Fornito, “Consistency and differences between centrality metrics across distinct classes of networks,” *ArXiv preprint arXiv:1805.02375*, 2018.
- [141] M. Rubinov and O. Sporns, “Complex network measures of brain connectivity: Uses and interpretations,” *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.
- [142] F. de Pasquale, M Corbetta, V Betti, and S Della Penna, “Cortical cores in network dynamics,” *Neuroimage*, vol. 180, pp. 370–382, 2018.
- [143] J. J. McAuley, L. da Fontoura Costa, and T. S. Caetano, “Rich-club phenomenon across complex network hierarchies,” *Applied Physics Letters*, vol. 91, no. 8, p. 084 103, 2007.
- [144] L. Harriger, M. P. Van Den Heuvel, and O. Sporns, “Rich club organization of macaque cerebral cortex and its role in network communication,” *PloS one*, vol. 7, no. 9, e46497, 2012.
- [145] M. Senden, G. Deco, M. A. de Reus, R. Goebel, and M. P. van den Heuvel, “Rich club organization supports a diverse set of functional network configurations,” *Neuroimage*, vol. 96, pp. 174–182, 2014.

- [146] P. Rombach, M. A. Porter, J. H. Fowler, and P. J. Mucha, “Core-periphery structure in networks (revisited),” *SIAM Review*, vol. 59, no. 3, pp. 619–646, 2017.
- [147] G. Zamora-López, C. Zhou, and J. Kurths, “Cortical hubs form a module for multisensory integration on top of the hierarchy of cortical networks,” *Frontiers in neuroinformatics*, vol. 4, p. 1, 2010.
- [148] I. Ripp, A.-N. zur Nieden, S. Blankenagel, N. Franzmeier, J. N. Lundström, and J. Freiherr, “Multisensory integration processing during olfactory-visual stimulation fmri graph theoretical network analysis,” *Human brain mapping*, vol. 39, no. 9, pp. 3713–3727, 2018.
- [149] A. Destexhe and T. J. Sejnowski, “The wilson–cowan model, 36 years later,” *Biological cybernetics*, vol. 101, no. 1, pp. 1–2, 2009.
- [150] K. M. Sabrin, Y. Wei, M. van den Heuvel, and C. Dvornik, “The hourglass organization of the *C. elegans* connectome,” *BioRxiv*, p. 600 999, 2019.
- [151] B. Mišić, R. F. Betzel, A. Nematzadeh, J. Goni, A. Griffa, P. Hagmann, A. Flammini, Y.-Y. Ahn, and O. Sporns, “Cooperative and competitive spreading dynamics on the human connectome,” *Neuron*, vol. 86, no. 6, pp. 1518–1529, 2015.
- [152] J.-P. Noel, M. Łukowska, M. Wallace, and A. Serino, “Multisensory simultaneity judgment and proximity to the body,” *Journal of Vision*, vol. 16, no. 3, pp. 21–21, 2016.
- [153] G. Zhou, G. Lane, T. Noto, G. Arabkheradmand, J. A. Gottfried, S. U. Schuele, J. M. Rosenow, J. K. Olofsson, D. A. Wilson, and C. Zelano, “Human olfactory-auditory integration requires phase synchrony between sensory cortices,” *Nature communications*, vol. 10, no. 1, p. 1168, 2019.
- [154] J. X. Maier, M. L. Blankenship, J. X. Li, and D. B. Katz, “A multisensory network for olfactory processing,” *Current Biology*, vol. 25, no. 20, pp. 2642–2650, 2015.
- [155] J. S. Baizer, C. C. Sherwood, M. Noonan, and P. R. Hof, “Comparative organization of the claustrum: What does structure tell us about function?” *Frontiers in systems neuroscience*, vol. 8, p. 117, 2014.
- [156] J. D. Van Horn, *What is old is new again: Investigating and analyzing the mysteries of the claustrum*, 2019.
- [157] F. C. Crick and C. Koch, “What is the function of the claustrum?” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 360, no. 1458, pp. 1271–1279, 2005.

- [158] J. R. Whitlock, “Posterior parietal cortex,” *Current Biology*, vol. 27, no. 14, R691–R695, 2017.
- [159] N. Nikbakht, A. Tafreshiha, D. Zoccolan, and M. E. Diamond, “Supralinear and supramodal integration of visual and tactile signals in rats: Psychophysics and neuronal mechanisms,” *Neuron*, vol. 97, no. 3, pp. 626–639, 2018.
- [160] H. Mohan, R. de Haan, H. D. Mansvelder, and C. P. de Kock, “The posterior parietal cortex as integrative hub for whisker sensorimotor information,” *Neuroscience*, vol. 368, pp. 240–245, 2018.
- [161] R. Hishida, M. Kudoh, and K. Shibuki, “Multimodal cortical sensory pathways revealed by sequential transcranial electrical stimulation in mice,” *Neuroscience research*, vol. 87, pp. 49–55, 2014.
- [162] R. R. Menzel and D. S. Barth, “Multisensory and secondary somatosensory cortex in the rat,” *Cerebral Cortex*, vol. 15, no. 11, pp. 1690–1696, 2005.
- [163] R. Romo, A. Hernández, A. Zainos, L. Lemus, and C. D. Brody, “Neuronal correlates of decision-making in secondary somatosensory cortex,” *Nature neuroscience*, vol. 5, no. 11, p. 1217, 2002.
- [164] J. H. Marshel, M. E. Garrett, I. Nauhaus, and E. M. Callaway, “Functional specialization of seven mouse visual cortical areas,” *Neuron*, vol. 72, no. 6, pp. 1040–1054, 2011.
- [165] M. N. Hebart and G. Hesselmann, “What visual information is processed in the human dorsal stream?” *Journal of Neuroscience*, vol. 32, no. 24, pp. 8107–8109, 2012.
- [166] B. E. Stein, T. R. Stanford, and B. A. Rowland, “Development of multisensory integration from the perspective of the individual neuron,” *Nature Reviews Neuroscience*, vol. 15, no. 8, p. 520, 2014.
- [167] W. Jiang, M. T. Wallace, H. Jiang, J. W. Vaughan, and B. E. Stein, “Two cortical areas mediate multisensory integration in superior colliculus neurons,” *Journal of Neurophysiology*, vol. 85, no. 2, pp. 506–522, 2001.
- [168] S. Tyll, E. Budinger, and T. Noesselt, “Thalamic influences on multisensory integration,” *Communicative & integrative biology*, vol. 4, no. 4, pp. 378–381, 2011.
- [169] A. Goulas, K. Zilles, and C. C. Hilgetag, “Cortical gradients and laminar projections in mammals,” *Trends in Neurosciences*, 2018.

- [170] J. E. Knox, K. D. Harris, N. Graddis, J. D. Whitesell, H. Zeng, J. A. Harris, E. Shea-Brown, and S. Mihalas, “High-resolution data-driven model of the mouse connectome,” *Network Neuroscience*, vol. 3, no. 1, pp. 217–236, 2018.
- [171] M. P. Van den Heuvel, E. T. Bullmore, and O. Sporns, “Comparative connectomics,” *Trends in cognitive sciences*, vol. 20, no. 5, pp. 345–361, 2016.
- [172] D. C. Van Essen and M. F. Glasser, “The human connectome project: Progress and prospects,” in *Cerebrum: The Dana forum on brain science*, Dana Foundation, vol. 2016, 2016.
- [173] N. Perra and S. Fortunato, “Spectral centrality measures in complex networks,” *Physical Review E*, vol. 78, no. 3, p. 036 107, 2008.
- [174] A. Hagberg, P. Swart, and D. S Chult, “Exploring network structure, dynamics, and function using networkx,” Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2008.