

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Parsimonious estimates of functional connectivity and biomarkers of cognitive development in early childhood

Permalink

<https://escholarship.org/uc/item/3jj801h0>

Author

Madhow, Sylvia

Publication Date

2024

Peer reviewed|Thesis/dissertation

Parsimonious estimates of functional connectivity and biomarkers of cognitive
development in early childhood

by

Sylvia Madhow

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Kristofer Bouchard, Co-chair

Professor Frederic Theunissen, Co-chair

Professor Michael Yartsev

Professor Jack Gallant

Spring 2024

Parsimonious estimates of functional connectivity and biomarkers of cognitive
development in early childhood

Copyright 2024
by
Sylvia Madhow

Abstract

Parsimonious estimates of functional connectivity and biomarkers of cognitive development in early childhood

by

Sylvia Madhow

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Kristofer Bouchard, Co-chair

Professor Frederic Theunissen, Co-chair

Mapping the functional connectome underlying recorded time-series of brain activity can reveal meaningful pathways of large-scale neural computation. However, it remains commonplace to estimate functional connectomes using pairwise correlative methods which are prone to spurious correlations and vulnerable to structured noise. Multivariate statistical estimates of functional connectivity offer new possibilities for generating parsimonious networks. We generated functional connectomes from resting state EEG (electroencephalography) and fMRI (functional magnetic resonance imaging) collected from infants and children under the age of 3. We generated directed multivariate functional connectomes using the ensemble statistical framework Union of Intersections (UoI) to perform regularized multivariate linear regression. We also generated pairwise connections using the Pearson correlation coefficient between each pair of time series in the recording. We found that multivariate estimates of functional connectivity from EEG are sparse and small-world, while pairwise connectomes are lattice-like and spatially correlated. We found a significant difference in small-worldness ω ($p \ll 0.0001$) and proximity dependence of coupling strength ($p \ll 0.0001$) between pairwise and multivariate EEG functional connectomes. Functional connectomes generated from fMRI were spatially distributed for both pairwise and multivariate estimates, but we observed a significant difference in ω ($p \ll 0.001$), with multivariate connectomes clustered tightly around 0, indicating small-worldness. We observed lateralized structure in the multivariate fMRI connectomes that remained stable across age groups. In particular,

coupling between the language network posterior superior temporal gyrus (pSTG) and salience network supramarginal gyrus (SMG) showed strong positive ipsilateral connections and strong negative contralateral connections, reinforcing the known lateralization of the language network. The same grid-like structure was observed in sensorimotor lateral fields coupling with the dorsal attention network intraparietal sulcus (IPS). This lateralized structure was not found in pairwise estimates.

We used fMRI functional connectomes to predict cognitive development scores assessed with the Mullen Scales of Early Learning (MSEL). We partitioned subjects into three groups using unsupervised clustering based on their raw MSEL scores. Multiple feature sets were extracted from the functional connectomes, by principal components analysis (PCA) and sparse PCA, and by selecting task-relevant functional network pairs. Random forest classifiers were used to predict the Mullen groups from neural feature vectors. We found that sparse PCs predicted raw scores significantly better than chance for multivariate connectomes, while both PCs and sparse PCs scored significantly better than chance for pairwise connectomes. For multivariate connectomes, functional connectivity between language and salience (SN) scored better than chance. For pairwise connectomes, coupling between the default mode network (DMN) and the frontoparietal network (FPN), and between SN and DMN scored better than chance.

Our understanding of distributed representations of complex sounds in auditory cortex remains incomplete. This is in part due a lack of experimental data for neural responses to complex naturalistic stimuli. Here, we describe the development of a large, diverse dataset of natural sounds. The sources and semantic content of the sounds are described, and acoustic features are calculated for a hand-selected set of 99 high-quality sounds and used to predict semantic labels using a support vector machine. We briefly discuss the intended use case of the sound database as a stimulus set for the characterization of neural representations of natural sound statistics, and selection for features that discriminate semantic categories.

Contents

Contents	ii
List of Figures	iv
List of Tables	xiii
1 Introduction	1
1.1 Description of Chapters	3
2 Background	6
3 Multivariate functional connectivity estimation in fMRI and EEG	12
3.1 Introduction	12
3.2 Methods	15
3.3 Results	23
3.4 Conclusion	40
4 Prediction of cognitive development from multivariate functional connectomes	43
4.1 Introduction	43
4.2 Methods	44
4.3 Results and Discussion	61
4.4 Conclusion	67
5 Acoustic features of complex sounds	69
5.1 Introduction and Motivation	69
5.2 Methods	72
5.3 Results and Discussion	82
5.4 Conclusion	84

6 Conclusion	87
Appendix	90
A Correlations of Mullen Scores to Connectomes	90
B Principal components for multivariate and pairwise connectomes . . .	90
Bibliography	138

List of Figures

3.1	Functional connectivity is vulnerable to spurious correlations that may misattribute connectivity. Ground truth connectivity is shown in blue on the left, estimated connectivity is shown in orange on the right. The top panel shows that due to a lack of explaining away, spurious correlation may be found for an unconnected pairs of nodes that each share variability with a third intermediate node. The middle panel shows that if multiple nodes are interconnected, shared variability may be attributed to the wrong pair of nodes. Finally, the bottom panel shows that methods such as purely pairwise approaches do not capture directionality, potentially missing crucial asymmetric structure in the connectome. . . .	14
3.2	The layout of the 128 channel sensor cap used for collecting the EEG dataset. The channels highlighted in blue are part of the 10-20 channel montage (most frequently used subset of channels that covers headspace with limited overlap in signal recording). Electrodes highlighted in yellow ones were also processed by HAPPE.	16
3.3	Resting state fMRI networks, figure courtesy of Dr. Bruchhage[39]. a. cerebellar: anterior, posterior. b. default mode: MPFC, medial prefrontal cortex; left and right LP, lateral parietal cortex; PCC, posterior cingulate cortex. c. dorsal attention: left and right FEF, frontal eye field; left and right IPS, inferior frontal gyrus. d. fronto-parietal: left and right LPFC, lateral prefrontal cortex; left and right PPC, posterior parietal cortex. e. salience: left and right anterior insula, left and right RPFC, rostral prefrontal cortex; left and right SMG, supramarginal gyrus. f. sensorimotor: superior, left and right lateral. g. visual: medial, occipital, left and right lateral. h. language: left and right IFG, inferior frontal gyrus; left and right pSTG, posterior superior temporal gyrus. . . .	18
3.4	Connectomes were generated by setting $A_{ij} = \beta_{ij}$, where A_{ij} is the adjacency matrix of the graph and β_{ij} is the coupling coefficient between electrodes i and j	20

- 3.5 **Multivariate coupling models produce sparse, small-world, spatially distributed functional connectivity networks in resting state EEG data.** **a.** Comparison of a weight matrix for a representative example of the high-gamma frequency band component of resting-state EEG data collected from a 6 month old subject. Pairwise FC (top) is dense, and coupling strength is heavily dependent on the physical proximity between each pair of electrodes. **b.** high-gamma resting state, 12 month old subject. **c.** high-gamma resting state, 36 month old subject. **d.** Small worldness, ω vs. sparsity for multivariate and thresholded pairwise graphs. **e.** Coupling weight as a function of electrode distance for high-gamma networks for all subjects aged 6-36 months. **f.** Median Pearson correlation coefficient ρ of coupling strength to electrode-distance for all high gamma subjects. **g.** Median ω for all high gamma subjects. 24
- 3.6 **Multivariate EEG functional connectomes are sparse, small-world, spatially distributed across frequency bands.** **a-c.** Functional connectivity network (pairwise top, multivariate bottom) for 36-month old example subject shown for different frequency bands: **a.** theta **b.** beta **c.** high gamma. **d.** Functional coupling vs. node distance for each frequency band. **e.** ω vs network sparsity for each frequency band. **f.** bar graphs of median ω and ρ for each frequency band. 26
- 3.7 **fMRI functional connectomes are small-world and spatially distributed for both pairwise and multivariate estimates.** **a.** Median functional connectivity template for subjects aged 3-6 months. $n = 12$. **b.** Template for 12-15 months. **c.** Template for 19-22 months. **d.** Coupling strength vs. ROI distance for subjects in corresponding age group. **e.** Median Pearson correlation coefficient ρ between coupling strength and ROI distance. **f.** Small worldness, ω vs. sparsity for multivariate and thresholded pairwise graphs. **g.** Distribution of ω , where pairwise ω is calculated on a network thresholded to the level of sparsity of the corresponding UoI_{LASSO} network. 27
- 3.8 **Multivariate fMRI connectome templates show a large fraction of stable couplings, but also some highly unstable couplings.** **a-c** Standard deviation of bootstrapped aggregate templates for pairwise (top) and multivariate (bottom) functional networks, with $n = 12$ subselected from N available subjects. **a.** 120-210 days, $N = 23$, **b.** 300-390 days, $N = 20$, **c.** 360-450 days, $N = 25$. **d-f** kde-smoothed distribution of standard deviations for listed age ranges. 28

3.9	Standard deviation of edge weights scales with mean edge weight. a-d: standard deviation plotted against mean edge weight for age groups a. 90-180 days. b. 180-270 days. c. 270-360 days. d. 360-450 days. e-h: standard deviation plotted against median edge weight for age groups e. 90-180 days. f. 180-270 days. g. 270-360 days. h. 360-450 days.	30
3.10	Contributions of dorsal attention network to sensorimotor network show lateralized connectivity: Multivariate functional connectivity networks showcase strong negative contralateral weights from dorsal attention intraparietal sulcus (IPS) to lateral sensorimotor during periods of intensive motor learning.	32
3.11	Contributions of salience network to language network show lateralized structure for IFG and pSTG: Multivariate functional connectivity shows positive ipsilateral connectivity and negative contralateral connectivity between salience network supramarginal gyrus and language network pSTG (Wernicke’s area). Anterior insula contributes strongly to IFG (Broca’s area).	33
3.12	Contributions of frontoparietal network to language network show evolving contributions to pSTG. Both ipsilateral and contralateral couplings can be seen, but pSTG-PPC appears more likely to show positive ipsilateral coupling and negative contralateral coupling. . .	34
3.13	Contributions of default mode network to language network show lateralized contributions from lateral parietal default mode regions to right pSTG. Default mode network seems to contribute more consistently to Wernicke’s area (STG) than to Broca’s area (IFG).	36
3.14	Fine motor raw scores are correlated with select sensorimotor, dorsal attention, cerebellar, and salience network connections. Scatter plots of Fine Motor raw score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black.	37
3.15	Receptive language scores are correlated with frontoparietal contributions to the language network. Scatter plots of Receptive Language raw score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black.	38
3.16	Early learning composite score is correlated with frontoparietal-default mode couplings. Scatter plots of Early Learning Composite score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black. .	39

3.17	Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations. a. Spearman correlation coefficients and b. p-values between MSEL Expressive Language raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. c. Spearman correlation coefficients and d. p-values between MSEL Expressive Language t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	41
4.1	Raw Mullen scores for patients with available fMRI data can be subdivided into three groups. a. Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 75$. b. Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.	47
4.2	Grouping of all available MSEL scores, including those for which neural data was not available, replicates the three groups from Figure 4.1, but adds a fourth group with older subjects a. Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 200$. b. Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.	50
4.3	Subjects between the ages of 90 and 1000 days fall into 3 consistent bands that are preserved across expressive language, receptive language, fine motor, gross motor, and visual receptive scores. Between the ages of 90 and 1000 days, raw scores do not vary significantly with age. When older subjects between 1000 and 2000 days are added into the dataset, a fourth group, mostly comprising older children, shows strong correlation with age in raw scores. For group 4, the early learning composite score provides a metric that is decorrelated from age. Groups 1-4 are not significantly separated for the early learning composite score.	51
4.4	T-score groups separate Early Learning Composite into two distinct distributions. a. Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 75$. b. Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.	52

4.5	PCs are dense and difficult to interpret, but some lateralized structure is discernible. a. Coefficients of first 10 principal components plotted in the original feature space. b. Distribution of PC values within Mullen groups.	53
4.6	Sparse PCs show selection for lateralized connections and are highly concentrated in the visual functional network. a. Sparse coefficients of first 10 sparse principal components plotted in the original feature space. b. Distribution of PC values within Mullen groups.	54
4.7	Excluding visual network from sparse PCA does not improve separability of Mullen groups. a. Sparse coefficients of first 10 sparse principal components plotted in the original feature space. b. Distribution of PC values within Mullen groups.	56
4.8	Functional connectome principal components do not consistently predict MSEL scores. 10-fold cross-validated confusion matrices for optimized random forest classifiers for 10-component PCA, sparse PCA, and sparse PCA with vision excluded.	59
4.9	Functional connectivity between default-mode, salience, and frontoparietal networks slightly improves classifier performance. a. Schematic of the interconnected “big three” functional networks: default mode (DMN), frontoparietal (FPN) and salience (SN). b. 10-fold cross-validated confusion matrices for optimized random forest classifiers for each pair of networks.	60
4.10	Functional connectivity between other selected network pairs are little better than chance at predicting MSEL groupings. 10-fold cross-validated confusion matrices for optimized random forest classifiers for other selected network pairs. Language-salience, default mode-visual and visual-default mode, and sensorimotor-dorsal attention perform slightly better than chance (33%).	61
4.11	Mullen raw scores are better predicted by pairwise connectome principal components than by multivariate connectome principal components. Confusion matrices for 10-fold cross-validated RFCs, with feature sets comprising 20 components each.	62
4.12	T-scores are better predicted by sparse PCs of multivariate connectomes than pairwise connectomes. Confusion matrices for 10-fold cross-validated RFCs predicting t-score Mullen groupings from 20-dimensional sparse PCs. a. Multivariate connectomes b. Pairwise connectomes	63

5.1	Figure courtesy of Baratham et al.[119]. Tonotopic map generated from μ ECoG recording from rat A1.	70
5.2	Histogram of full sound database by number of discrete number of files (top) and total number of available seconds (bottom) in a given category.	73
5.3	Spectrograms of different types of natural vocalizations drawn from the sound database. Note complex harmonic structure and temporal modulation.	74
5.4	Distribution of semantic labels for primary analysis dataset of hand-selected sounds, Set A.	75
5.5	Log distribution of the total power contained in each 200ms sound frame, as calculated by the \mathcal{L}_1 -norm of the waveform. 25th, 50th, and 75th percentiles are designated.	77
5.6	Example biosound of a 200ms frame of a fox cackling. a. Sound pressure waveform (black) with temporal envelope (red) b. Power spectral density (PSD). Quadrants 1, 2, 3 are superimposed in black, mean formants 1 (red) and 2 (cyan). c. STFT spectrogram. e. Frequency of fundamental frequencies and formants over time.	78
5.7	Visualization of selected acoustic features shows some separability between example semantic categories. In particular, foxes show high variability in their first formant, but tight distribution for their second. Vocalizations (fox and human) have high variability in modulation separability compared to music, which shows consistently high separability.	81
5.8	Prediction of semantic labels from acoustics performs well above chance, but semantic categories with high levels of acoustic variability are incorrectly assigned to other sound categories. Confusion matrices with predicted semantic labels on the x-axis and true labels on the y axis. Left: first 15 PCA components. Right: Five UMAP dimensions.	83
5.9	Figure courtesy of Vitto Resnick in the Bouchard lab. Example of stimulus set for presentation to anesthetized rats.	86

.1	Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations. a. Spearman correlation coefficients and b. p-values between MSEL Fine Motor raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. c. Spearman correlation coefficients and d. p-values between MSEL Fine Motor t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	91
.2	Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations. a. Spearman correlation coefficients and b. p-values between MSEL Gross Motor raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. c. Spearman correlation coefficients and d. p-values between MSEL Gross Motor t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	92
.3	Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations. a. Spearman correlation coefficients and b. p-values between MSEL Receptive Language raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. c. Spearman correlation coefficients and d. p-values between MSEL Receptive Language t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	93
.4	Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations. a. Spearman correlation coefficients and b. p-values between MSEL Visual Reception raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. c. Spearman correlation coefficients and d. p-values between MSEL Visual Reception t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	94
.5	Neither multivariate or pairwise connectomes show statistically significant correlates to Early Learning Composite score. a. Spearman correlation coefficients and b. p-values between MSEL Early Learning Composite scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.	95

.6	Cumulative variance captured for multivariate connectome PCA. Blue lines show total variance captured by 20 PCs, while red lines show 80% variance captured.	96
.7	Cumulative variance captured for pairwise connectome PCA. Blue lines show total variance captured by 20 PCs, while red lines show 80% variance captured.	97
.8	Principal components 1 (top) and 11 (bottom) of multivariate connectomes.	98
.9	Principal components 2 (top) and 12 (bottom) of multivariate connectomes.	99
.10	Principal components 3 (top) and 13 (bottom) of multivariate connectomes.	100
.11	Principal components 4 (top) and 14 (bottom) of multivariate connectomes.	101
.12	Principal components 5 (top) and 15 (bottom) of multivariate connectomes.	102
.13	Principal components 6 (top) and 16 (bottom) of multivariate connectomes.	103
.14	Principal components 7 (top) and 17 (bottom) of multivariate connectomes.	104
.15	Principal components 8 (top) and 18 (bottom) of multivariate connectomes.	105
.16	Principal components 9 (top) and 19 (bottom) of multivariate connectomes.	106
.17	Principal components 10 (top) and 20 (bottom) of multivariate connectomes.	107
.18	Sparse principal components 1 (top) and 11 (bottom) of multivariate connectomes.	108
.19	Sparse principal components 2 (top) and 12 (bottom) of multivariate connectomes.	109
.20	Sparse principal components 3 (top) and 13 (bottom) of multivariate connectomes.	110
.21	Sparse principal components 4 (top) and 14 (bottom) of multivariate connectomes.	111
.22	Sparse principal components 5 (top) and 15 (bottom) of multivariate connectomes.	112
.23	Sparse principal components 6 (top) and 16 (bottom) of multivariate connectomes.	113
.24	Sparse principal components 7 (top) and 17 (bottom) of multivariate connectomes.	114
.25	Sparse principal components 8 (top) and 18 (bottom) of multivariate connectomes.	115
.26	Sparse principal components 9 (top) and 19 (bottom) of multivariate connectomes.	116
.27	Sparse principal components 10 (top) and 20 (bottom) of multivariate connectomes.	117
.28	Principal components 1 (top) and 11 (bottom) of pairwise connectomes. .	118
.29	Principal components 2 (top) and 12 (bottom) of pairwise connectomes. .	119

.30	Principal components 3 (top) and 13 (bottom) of pairwise connectomes. .	120
.31	Principal components 4 (top) and 14 (bottom) of pairwise connectomes. .	121
.32	Principal components 5 (top) and 15 (bottom) of pairwise connectomes. .	122
.33	Principal components 6 (top) and 16 (bottom) of pairwise connectomes. .	123
.34	Principal components 7 (top) and 17 (bottom) of pairwise connectomes. .	124
.35	Principal components 8 (top) and 18 (bottom) of pairwise connectomes. .	125
.36	Principal components 9 (top) and 19 (bottom) of pairwise connectomes. .	126
.37	Principal components 10 (top) and 20 (bottom) of pairwise connectomes.	127
.38	Sparse principal components 1 (top) and 11 (bottom) of pairwise connectomes.	128
.39	Sparse principal components 2 (top) and 12 (bottom) of pairwise connectomes.	129
.40	Sparse principal components 3 (top) and 13 (bottom) of pairwise connectomes.	130
.41	Sparse principal components 4 (top) and 14 (bottom) of pairwise connectomes.	131
.42	Sparse principal components 5 (top) and 15 (bottom) of pairwise connectomes.	132
.43	Sparse principal components 6 (top) and 16 (bottom) of pairwise connectomes.	133
.44	Sparse principal components 7 (top) and 17 (bottom) of pairwise connectomes.	134
.45	Sparse principal components 8 (top) and 18 (bottom) of pairwise connectomes.	135
.46	Sparse principal components 9 (top) and 19 (bottom) of pairwise connectomes.	136
.47	Sparse principal components 10 (top) and 20 (bottom) of pairwise connectomes.	137

List of Tables

2.1	Traditional frequency bands [59] of brain waves.	9
3.1	Standard deviation of connectomes edge weights shows significant correlation to magnitudes of corresponding mean and median weight values for multivariate connectomes. Spearman correlation coefficients of standard deviation with absolute value of mean and median edge weights for multivariate and pairwise connectomes within 90 day age windows.	31
3.2	Spearman correlation coefficients and p-values for selected Mullen scores and ROI-ROI coupling weights for all plots show in Figures 3.14, 3.15, 3.16.	35
4.1	Mullen scores, group size, and ages for clustering-derived groups for the subset of subjects for whom neural data was provided. Means and standard deviations (in parentheses) are reported.	48
4.2	Mullen scores, group sizes, and ages for clustering-derived groups for larger dataset with expanded age range. Means and standard deviations (in parentheses) are reported.	49
4.3	Raw Mullen score groupings are predicted by pairwise connectome features better than by multivariate connectomes. Tenfold cross-validated classification scores of grid-search optimized random forest classifiers for feature set used for prediction of Mullen raw score groupings. Feature sets that performed significantly better than chance (33.3%) are marked with an asterisk.	58
4.4	T-score Mullen groupings are predicted by multivariate connectomes better than by pairwise networks. Tenfold cross-validated classification scores of grid-search optimized random forest classifiers for feature set used for prediction of Mullen raw score groupings. Feature sets that performed significantly better than chance (50%) are marked with an asterisk.	59

5.1 Components of acoustic feature vectors calculated for all 200 ms sound frames. 80

Acknowledgments

This dissertation could never have been written without the tireless efforts of my advisor, Professor Kristofer Bouchard. The breadth and depth of his understanding of neuroscience and its intersection with data science, combined with his impeccable attention to detail, have challenged and expanded my perspective immeasurably during my time in his lab. On a personal level, his support and empathy, and his dogged determination to maneuver me through the complexities, setbacks, and self-doubt of graduate school, were above and beyond anything I could have imagined.

I benefited immensely from the generosity of my committee chairman, Professor Michael Yartsev, in offering his time to the development of my doctoral research. He was, incidentally, a delight as a member of my qualifying examination committee as well—particularly his enthusiasm in asking me questions and immediately answering them himself. Additionally, members of his lab, including Dr. Julie Elie and Dr. Daria Genzel were immensely helpful in my efforts to curate the sound library and calibrate laboratory equipment to handle ultrasonic sounds.

Ever since rotations in my first year of grad school, Professor Frederic Theunissen has offered both academic and moral support. For the guidance and expertise in auditory neuroscience he has provided as a member of my thesis committee, and for his help in welcoming me back to the Helen Wills Neuroscience graduate program after my leave of absence, I am deeply grateful.

Professor Jack Gallant was an essential member my thesis committee. First, because of his unique insight into the field of fMRI, and second because his commentary makes everything far more fun. Professor Gallant's keen observations and blunt delivery are unique and memorable, and responsible for 45% of the anecdotes I'll take with me from my time at Helen Wills Institute.

The members of my qualifying committee, Professor Robert Knight, Professor Bruno Olshausen, and Professor Kevin Weiner,

To Candace Groskreutz and Leleña Avila, who have supported countless students through the doctoral gauntlet: the program would not be the same without you. I would not be the same without you.

My cohort comprised the most brilliant motley crew of Renaissance scholars I have ever had the privilege to know, and I enjoyed every moment spent among them. To these most passionate, welcoming, sometimes inexplicable people, Dariya Bakshinska, Hayley Bounds, Christiane Voufo, Molly LaPoint, Ellen Zippi, Ryan Schultz, Lily Xue Gong, and Celia Ford: thank you for being the way that you are.

To members of my lab, past and present, there are too many individual thanks to list. But in particular Dr. Jesse Livezy, Dr. Pratik Sachdeva, Dr. Charles Frye,

and Max Dougherty set the gold standard for creating a work environment founded in cooperation, rigor, and ethics.

I have been privileged to have as role models and guides many brilliant, passionate, and kind coworkers. My primary collaborators on the work presented in Chapters 3 and 4, Professor Muriel Bruchhage and Professor Laurel Gabard-Durnham, are exceptional scientists who go above and beyond in supporting newer researchers, and I will carry their example with me throughout my career. Likewise my supervisors from my very first job, Nicholas Hogasten and Stephanie Lin, set a standard of patience and mentorship that I can only strive to emulate.

Thank you to my parents, Alejandra Folguera and Upamanyu Madhow, for being so loving, so supportive, believing in me so enthusiastically that I frankly wanted to hide under the table. Thank you to my little brother, Sunil Madhow, for being so funny, and secretly rather sweet.

Writing down the ways in which my husband Ian has helped me, supported me, and uplifted me would require a separate, longer dissertation. He has been there for every sleepless night putting together figures, every elusive bug in my code, every wild-eyed monologue. His earnest good cheer and steady optimism are, candidly, incomprehensible—but I could not have made it to this point without them.

To my dog, Ogden: thank you for keeping my feet warm. They were very cold.

To my daughter, Miranda: thank you for lighting up the house with your cackling and soliloquies while I work, for greeting me with an impish grin at the end of each day, and for (sometimes) sharing your sour grass.

Chapter 1

Introduction

The study of neuroscience seeks to find the neural substrates underlying cognition. Neural computation operates on many spatial and temporal scales[1], while measurement and interpretation of neural activity is often constrained to a single spatiotemporal scale[2]. With the rapid technological and scientific advance of neuroimaging, there are many modalities by which human brain activity can be measured, including functional near-infrared spectroscopy (fNIRS)[3, 4], magnetoencephalography (MEG)[5, 6], electroencephalography (EEG)[7, 8], and diffusion and functional MRI[9, 10, 11]. Each neuroimaging modality is characterized by strengths and weaknesses of spatial and temporal coverage, resolution, and signal-to-noise considerations[4, 5, 8, 10, 11]. Furthermore, the signal captured by each modality is not necessarily originating from the same underlying brain activity and functions[11, 12]. Furthermore, brain activity is highly responsive to environmental perturbation, and small differences in acquisition conditions may introduce confounds that complicate the analysis of already complex datasets[12, 13]. The work of finding the neural mechanism underlying a particular behavioral function is thus often the work of many decades, scientists, and overlapping measurement modalities and experimental paradigms[13].

One question of particular interest in human neuroscience is the discovery of neuro-biomarkers of behavioral or neurological phenotypes[14, 15, 16, 17]. Identifying characteristics of neural activity that correlate with various psychopathologies (e.g., executive disorders[18, 19], autism spectrum disorder[20], and insomnia[21] to name a few) and other neurological conditions (e.g., epilepsy[22, 23], cerebral palsy[24]) is desirable, both for the potential clinical utility and for the scientific insight such discoveries might offer into the underlying workings of the brain. When first identified, a novel neuro-biomarker might be modality-specific—for example, fMRI regions of the default mode network shown to predict insomnia[21]—only for subsequent studies to

uncover the neural mechanisms at play. Consider the theory of hippocampal indexing that was first advanced in 1986[25], nearly 40 years ago, based upon the apparent importance of the hippocampus in episodic memory retrieval[26]. Much early interest in the mechanisms underlying human memory stemmed from one patient, widely known as H.M., who underwent an extensive surgical resection of hippocampus and related areas, and subsequently lost the ability to form new episodic memories[27]. Since the original proposal of hippocampal indexing theory in 1986, behavioral and neural data ranging from human fMRI to rat single-neuron studies, have provided supporting evidence[28, 26]. This example is offered here to illustrate that the path from experimental neural data to cohesive theory is a long and recursive one that requires the integration of many experimental paradigms, neuroimaging modalities, and analysis methods[29].

Neural correlates of human pathologies have been studied across many different recording and imaging modalities, such as EEG and fMRI. Frequency bands of local field potentials have been characterized according to correlated behaviors such as sleep (alpha and delta bands) and working memory (theta band)[30]. Specific event-related potentials (ERPs) such as the “Negative central”, Nc, wave in the frontal brain regions[31] differs between neurotypical and ASD (autism spectrum disorder) children during facial recognition tasks[32]. EEG is often used in clinical settings due to the relative ease of acquisition, and its utility in diagnosing conditions that impact brain waves, such as narcolepsy[33]. fMRI biomarkers of behavior and cognition have been extensively characterized by both correlative and predictive studies. Many of the major canonical networks of fMRI were developed based on tasks, behaviors, and functions that correlate with activation of regions of interest (ROIs). For example, the default mode network, known as the “resting state” network[34], has been shown to increase activity during periods of quiet wakefulness, and decrease activity during cognitively intensive tasks. Because the default mode network is thought to be responsible for self-referential thoughts and inward-oriented cognition, it is thought that reduced resting-state functional default mode network (DMN) in autistic subjects[35] may suggest a deficit in self-referential thought[36]. The frontoparietal network, also known as the executive control network, is located in brain regions such as the prefrontal cortex, which is known to mediate attention and decision-making. Disruption of frontoparietal activation is widespread in many psychopathological disorders including depression, schizophrenia, and autism[37]. The frontoparietal functional network shows strong activation during cognitively intensive tasks, and is thought to “switch off” with the default mode network. Connectivity of the frontoparietal network has been correlated with general or “fluid” intelligence[38]. fMRI functional networks have also been linked to measures of cognitive development in infancy and early childhood[39], suggest-

ing that information about individual differences in cognition could be gleaned from fMRI-derived biomarkers. These neural correlates of high-level cognitive function provide biomarkers that can be used in concert with other measures to predict or understand the neural substrates of complex behaviors or pathologies.

Beyond correlations to observed behavior, fMRI data has also been used to predict or “decode” the stimuli that evoked the recorded activity. For example, using machine learning algorithms to predict the contents of natural movie stimuli presented during recording sessions[40][41]. Connectivity derived from resting state fMRI recordings have also been used to predict brain maturation[42] using simple multivariate machine learning techniques. This prediction analysis was then used to identify features (i.e., biomarkers) of the brain data that were most predictive of maturity. These predictive neuro-biomarkers can then be interpreted in the larger context of existing knowledge. Thus, prediction of cognitive function from brain biomarkers is both a useful tool for diagnostics, and a powerful method for gaining a deeper understanding of the brain.

In addition to the analysis of signals recorded directly from brain activity, it is a crucial undertaking of neuroscience to understand the stimuli that evoke brain activity. The analysis of sensory stimuli has been of particular interest to theoretical neuroscientists, who have developed neural-inspired algorithms to decompose natural stimuli (i.e., visual, acoustic) into sparse features. For example, sparse coding algorithms have been applied to both images[43] and sounds[44], producing learned receptive fields that resemble the receptive fields of sensory neurons early in the sensory processing pathways of the brain. Further analysis has sought to identify the salient acoustic features of human speech[45][46][47] and other vocalizations with complex semantic and spectrotemporal structure such as birdsong[48][49]. These insights have been instrumental in mapping neural representations of acoustic features[50][51][52]. Thus, the analysis of complex or naturalistic sensory stimuli is instrumental, both for the development of normative models of sensory processing in the brain and for the interpretation of recordings from sensory areas.

1.1 Description of Chapters

This dissertation details my work to contribute to the mapping of neural computation across different modalities. I approach the open problem of characterizing human brain connectivity using neural data by applying statistical machine learning methods to EEG and fMRI data collected from infant and young child subjects. Specifically, I estimated functional connectivity using a regularized regression ensemble algorithm, UoI_{LASSO} , to learn sparse functional coupling between the neural

populations measured in the data. I showed that functional connectomes estimated using this method are sparse, stable, and small-world for both fMRI and multiple frequency bands in EEG. I assessed the interpretability of the functional connectomes and their utility in the search for neuro-biomarkers of cognitive development in early childhood.

In another project, I collected a large and diverse database of complex natural sounds. I segmented the sounds into short, standardized samples and characterized their spectrotemporal structure. In addition, I worked to find acoustic features that enhance discriminability of semantic categories by classifying sounds using acoustic features as observation vectors.

Chapter 2 covers background topics relevant to this work, ranging from data science methods for the analysis and interpretation of high-dimensional datasets, to components of EEG and fMRI data. I briefly describe Union of Intersections (UoI), an ensemble machine learning framework that will see extensive use in this dissertation. I summarize the neural underpinnings of EEG and the functions of its frequency bands. I describe the major functional networks of fMRI and regions of interest (ROIs) that will be used in chapters 3 and 4.

In Chapter 3 I generated functional connectomes from EEG and fMRI data collected from infants and young children. I compared the properties of multivariate functional connectomes estimated using a regularized regression ensemble method, UoI_{LASSO}, with functional connectomes generated using pairwise correlation, the standard method. I showed that multivariate estimates of functional connectomes in EEG data produce sparse, small-world networks that cannot be replicated by thresholding pairwise networks. I also discuss structural differences in the features of multivariate and pairwise functional connectomes in fMRI, and showed that pairwise connectomes show greater correlation to standardized evaluation scores.

in Chapter 4 I used the fMRI multivariate functional connectomes from Chapter 3 as predictors of cognitive development in early childhood. I used unsupervised clustering to group children based on their scores on the Mullen Scales of Early Learning (MSEL), a measure of early childhood development. I calculated several feature sets from the functional connectome data and use these as the input observation variables to a Random Forest Classifier in order to predict Mullen group labels.

In Chapter 5 I describe the collection and characterization of a large database of natural acoustic objects, developed with the intention of characterizing distributed representations of complex sounds in rat auditory cortex. I calculated acoustic feature vectors for standardized sound files in order to identify acoustic features that predicted semantic categories. I briefly discuss initial attempts to classify sounds based on acoustic feature vectors. Finally, I examined the predictive power of these acoustic features by using them to predict the assigned semantic labels.

Chapter 6 concludes this dissertation and offers closing remarks and reflections on the potential future directions of the work presented in this thesis.

Chapter 2

Background

Parametric models in neuroscience

In scientific analyses of brain data, parametric models are used in diverse ways to explain neural activity. Encoding models map stimuli to brain signals. For example, an encoding model could be used to model auditory cortex activation when presented with natural sound stimuli, producing spectro-temporal receptive fields. In the other direction, decoding models are used to “translate” brain activity to identify the presented stimulus. Decoding models are often used in fMRI, reconstructing the perceived images using regression or deep learning models[53]. Functional coupling maps the statistical dependency of the activation of one brain region onto activation of other brain regions activated simultaneously. Normative models are trained without brain data, but are used to test a theorized neural mechanism. Examples of these models abound in cognitive neuroscience and psychology[29][54].

In scientific contexts, predictive power is not the singular purpose of a parametric model. Models must also be evaluated on the basis of their explanatory power, and thus their interpretability. While the definition of “interpretability” is both broad and situational, generally in order to be interpretable in the context of neuroscience, a model must be predictive, parsimonious, and stable[55]. Parsimonious feature selection—that is, selecting the minimal number of features needed to predict the target variable—while balancing adequate predictive power is necessary for interpretable scientific modeling. Furthermore, both feature selection and estimation must be stable under perturbation and robust to noise. These rigorous requirements place a heavy burden on statistical modeling in neuroscience.

Big Data and L_1 Regularization

Advances in neural imaging technology have led to a proliferation of high-dimensional recording modalities. These rich, high-dimensional datasets permit the investigation of elusive questions, such as neural representation of complex sensory stimuli, or the prediction of cognition and behavior from brain data. However, as the dimensionality of neural data increases, statistical analysis becomes more difficult, as expressed by the “curse of dimensionality[56].” Fortunately, it has been repeatedly demonstrated that high-dimensional brain data is often highly constrained to a small subset of activation patterns, such that observed activity can be expressed by a far smaller number of features[57]. It is thus often desirable to find a sparse representation, such that the L_0 norm of the data in that basis \ll the dimensionality of the data. Consider a target signal $y \in \mathbb{R}$ that is to be modeled by predictive variables $x \in \mathbb{R}^p$, mapped by parameters $\beta \in \mathbb{R}^p$ such that:

$$y = \beta^T x + \epsilon$$

$$\epsilon \sim N(0, \sigma^2)$$

In cases where a sparse β is assumed, feature compression, the selection of a subset of salient features to explain a target signal, may be induced through regularization. For example, Lasso is a linear regression algorithm that utilizes an L_1 norm regularization penalty to minimize the total magnitude of the model parameters. An estimate of β constrained by an L_1 penalty would then take the form of optimization problem:

$$\hat{\beta} \in \operatorname{argmin}_{\beta} \|y - \beta x\| + \lambda \|\beta\|_1$$

Where λ represents the regularization weight. For sufficiently high values of λ this L_1 constraint induces feature compression by driving components of β to zero. Thus, a Lasso fit with strong regularization will return a sparse model.

However, regularization alone does not ensure a stable or accurate model. Particularly in cases where features greatly outnumber data samples, or when features exhibit high levels of collinearity, an incorrect subset of features may be selected. While the resulting model may predict the outcome variable with a high level of accuracy, the model will not be reliably interpretable[58]

The Union of Intersections framework

Regularized regression models such as LASSO are used to fit sparse, predictive parametric models. As discussed previously, the utility of a scientific model lies not only in predictive accuracy, but in its interpretability as a stable, parsimonious predictive

model. The statistical-learning framework Union of Intersections (UoI) is a modular, flexible method to balance and enhance model selection and estimation. UoI has been tested on numerous neural data modalities, parametric models, and statistical algorithms. The UoI framework divides model selection and model estimation into separate steps, operating over bootstraps of the neural data. A brief summary of the procedure of UoI_{LASSO} is laid out as described in Bouchard et al., 2017[58] and Sachdeva et al., 2021[55]:

1. **Selection:** An intersection operation is used to select model supports (S_j). Regularization parameter λ is varied over a given range, and for each λ_j the Lasso optimization is solved for N resamples of the data to generate parameter estimates. Only variables that are consistently selected in a specified fraction of resamples for λ_j are included in S_j —this is the “intersection” of supports, which imposes a more rigorous, stable selection criterion for supports.
2. **Estimation:** Over N resamples of the data, fit each resample for each support set S_j without regularization (i.e., $\lambda = 0$). For each resample, the fit that performs best according to a chosen criterion (e.g., R^2 , BIC) is selected. These fits are then unionized (i.e., averaged), resulting in the final model.

The resulting model is an aggregate fit that preserves predictive power while improving stability and reducing estimation bias [58][55].

EEG and canonical frequency Bands

EEG, or electroencephalography, is a neural recording modality that captures macro-scale electrical potentials[30]. These potentials are often decomposed into frequency bands that have been characterized based on their behavioral correlates and neural functions. Table 2.1 summarizes the canonical frequency bands used in the the analysis of EEG signals.

fMRI Networks

This section briefly describes the functional brain networks included in the fMRI data analysed in Chapter 3. In particular, I will discuss commonly proposed behavioral and cognitive correlates for each network, and specify regions of particular interest.

The Default Mode Network

The default-mode network (DMN) has been implicated in internally-focused tasks, such as social cognition, self-reflection, and episodic memory[34]. DMN activity has

Wave	Frequency band (Hz)	Function
δ	2-3.99	deep sleep
θ	4-5.99	working memory, spatial navigation
low α	6-8.99	early-life α , begins at lower frequency
α	9-12.99	sleep, rest
β	13-29.99	active concentration, motor feedback
low γ	30-50	attention, perception
high γ	65-90	task-activated

Table 2.1: Traditional frequency bands [59] of brain waves.

been shown to be enhanced during restful or passive periods of wakefulness and suppressed while a subject is occupied with intensive, externally-focused tasks. The DMN has also been shown to be associated with areas related to language and speech comprehension, such as the anterior temporal cortex and middle temporal gyrus. Finally, the DMN is implicated in “daydreaming,” or spontaneous, off-task thoughts. Thus, the DMN can be simplistically described as the network of introspection, integrating multiple cognitive functions to a coherent inner narrative[60]. Some studies have shown reduced resting-state functional DMN in autistic subjects[35] and may suggest a difference in self-referential thoughts[36].

The DMN has been modeled as part of a trio with the salience (SN) and frontoparietal (FPN) networks. Specifically, SN has been interpreted as involved in “switching” from DMN to FPN when behaviorally relevant stimuli are detected[61].

The Frontoparietal Network

The frontoparietal network (FPN), also commonly referred to as the central executive network, has been shown to play a central role in mediating cognitive control—the process by which goals influence behavior[37]. The FPN is strongly linked with goal-oriented behavior and problem solving, and FPN coupling patterns shift significantly during rapid task-switching[38] and object-, feature-, and category-based attentional control[62]. Disruption of FPN activation is widespread in many disorders including depression, schizophrenia, and autism[37]. Furthermore, distributed coupling of the FPN, and in particular resting-state coupling to the DMN, is associated with fluid intelligence[38].

The Salience Network

The salience network, as the name may suggest, has been characterized as a network of brain regions that direct attention to behaviorally relevant or novel stimuli. The salience network shows activation in response to all manner of sensory modalities, including visual “oddball” stimuli[63], task error[64], and pain[65]. It has also been shown in an fMRI study that salience network functional connectivity was suppressed in dyslexic subjects during a narrative comprehension task[66]. In particular, the right dorsal anterior insula (dAI) has been described as a “causal outflow hub” that mediates the activation of other major networks[67]. This includes generating control signals that play a causal role in switching between the central-executive network (also known as the frontoparietal network) and the default-mode network[68].

The Dorsal Attention Network

The dorsal attention network (DAN) plays a role in the orientation of visual attention. In particular, it has been demonstrated that DAN is involved in top-down attention, and the selection of visual stimuli during task-driven behavior[69]. This stimulus selection mechanism has been observed to be anticipatory—that is, the DAN activates when a visual stimulus is expected, and moves visual attention in the expected direction][70]. The primary regions of focus in the DAN are the intraparietal sulcus (IPS) and the FEF (frontal eye field). It has been shown that in children the DAN shows greater intra-network connectivity than in adults, whereas functional coupling of DAN to regions outside the network is stronger in adults than in children[71].

The Cerebellar Network

The cerebellum has long been known to be involved in motor control and learning[72]. The contributions of the cerebellum to sensorimotor tasks is diverse[73][74], including oculomotor control, speech articulation, and hand and limb movement. Real-time prediction, timing of synchronous action, and inhibition of M1 excitability have all been shown to depend upon the cerebellar network[75]. However, more recent work has suggested that the cerebellum plays a greatly expanded role in higher-level tasks such as language and working memory, emotional processing, and executive tasks[76]. In addition, the neocerebellum has been shown to contribute to networks associated with high-level cognitive function such as the default mode, central executive, and salience networks[77]. It has been shown that in subjects with schizophrenia, cerebellar functional connectivity to DMN and FPN is disrupted[78].

The Language Network

The language functional network encapsulates different regions of the brain that are implicated in language processing. In particular, Wernicke's area, located in the superior temporal cortex, is associated with receptive language skills, i.e., speech comprehension. By contrast, Broca's area, in the inferior frontal cortex, is associated with language production, i.e., speech[79]. Functional connectivity in the language functional network has been studied extensively, and it has been demonstrated that language processing shows high levels of lateralization in the brain[80]. In particular, high levels of left-side lateralization have been observed in language learning. In analysis of resting-state functional connectivity, this left-side preference was observed in both Wernicke's and Broca's areas as the language network develops in children 3-5 years of age[81].

The Visual and Sensorimotor Networks

Functional connectivity studies of infant brains have shown that the visual and sensorimotor networks are established and functionally synchronized from birth[82], in contrast to networks such as the DMN[34], which are not temporally synchronized at birth and develop over time. The visual network comprises regions that contribute to processing of visual information. The primary and associative visual cortices are contained in the occipital lobe, which is responsible for visuospatial processing[83]. The sensorimotor network is widely connected in order to coordinate complex motor movement. In addition to connectivity to the visual network, the sensorimotor network is highly synchronized with the frontoparietal and cerebellar networks during moving tasks. Furthermore, for the coordination of accurate movement, connectivity with the intraparietal sulcus (DAN) was observed as a correlate[84].

Chapter 3

Multivariate functional connectivity estimation in fMRI and EEG

3.1 Introduction

Connectivity within the brain has long been studied as a means to predict behavior, cognition, and development[85]. Long- and short-range connections between neural populations integrate local circuit-level computations to generate large-scale, complex brain function. The mapping of these high-level connections generates a connectome based on a given neuroimaging modality or collection of modalities that can measure activity from across the whole brain. The mapping of the human connectome in particular presents extraordinary challenges, both due to the complexity of human brains, and the necessity of non-invasive, ethical techniques for data collection. With advances in neuroimaging has come a wealth of modalities that enable non-invasive human connectomics. These include (but are certainly not limited to): EEG, MEG, fNIRS, diffusion MRI, and functional MRI. Each modality is subject to different constraints of spatial and temporal resolution and coverage. As a result, connectomes from the same subject, taken on the same day with the same environmental conditions, but recorded using different modalities will capture different types of brain structure. Furthermore, the methods by which measurements of neural activity are “mapped” to a connectome vary widely, both between neuroimaging modalities and within them. As the fields of neuroimaging and data science advance in tandem, new avenues for data-driven connectomics become available. The analyses in this chapter and the next were carried out in collaboration with a large

consortium of domain experts specializing in different neuroimaging modalities. The overarching goal of the consortium is to map human brain data to early childhood cognitive developmental trajectories. Our part in the collaboration was to develop a stable multivariate method of estimating a human connectome that can be applied across multiple modalities, leveraging ensemble learning methods and regularized linear regression.

Structural and functional connectivity present complementary methods of defining a brain connectome. Functional connectivity is defined through the temporal association of activation between neural populations[86]. Structural connectivity endeavors to reconstruct the physical, anatomical structure of white matter in the brain. Thus, a structural connectome is likely to remain stable over the course of a single recording session (i.e., minutes). In contrast to a structural connectome, a functional connectome may vary dynamically over time—depending on such factors as the task or alertness of the subject. This variability in functional connectivity can be utilized to identify dynamic relationships between neural populations[87], highlighting stable connections that are used frequently. However, instability in the estimated functional connectome—resulting from factors such as sampling variability [88]—may result in fictitious dynamic variability that impedes interpretation of functional networks. In this chapter, I will discuss a method of estimating functional connectivity that enhances the stability and interpretability of the functional connectome.

Estimation of the functional connectome relies solely upon correlations between time series of recorded neural activation, and generally does not require a model based on assumptions of the underlying neurobiology. Because it is purely correlative in nature, functional coupling cannot be interpreted as a direct, causal connection between two regions. Rather, strong co-activation between two regions, A and B, may be explained by a multitude of underlying structural connectomes. For example, A and B may appear strongly correlated as a result of a third signal, C, driving both A and B in parallel. Alternatively, the correlation may be due to “daisy chaining,” such that A strongly drives C, which in turn strongly drives B, resulting in an estimated connection between A and B (Figure 3.1). A purely correlative approach is also vulnerable to spurious correlations caused by structured noise or other confounds (e.g., signal leakage caused by conductivity in EEG recordings, head movements in fMRI). It is thus desirable to account for noise and spurious variability when estimating the functional connectome. While rigorously pre-processing data can remove known confounds, spurious correlations are more difficult to address. For this reason, I will argue that a multivariate approach to functional connectivity mitigates common failings of functional connectomes based on purely pairwise estimation.

Functional connectome analysis is used for a wide array of neuroimaging modalities, including functional near-infrared spectroscopy (fNIRS), magnetoencephalog-

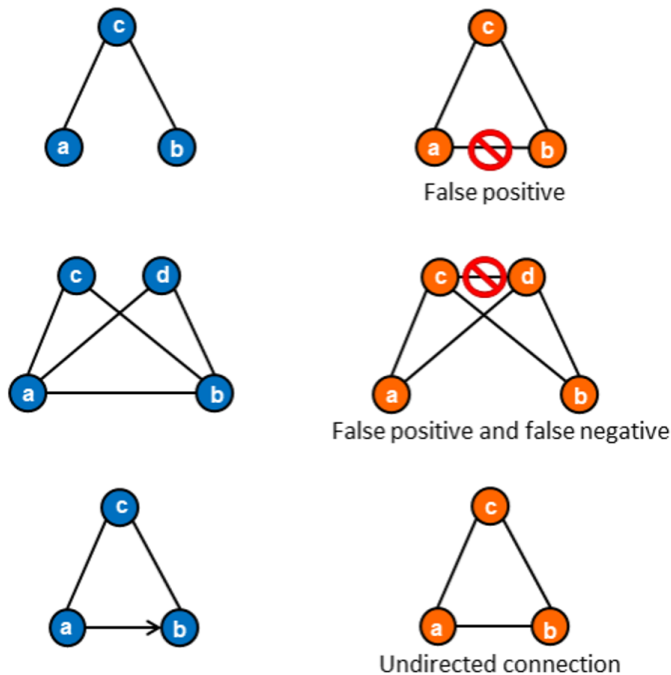


Figure 3.1: **Functional connectivity is vulnerable to spurious correlations that may misattribute connectivity.** Ground truth connectivity is shown in blue on the left, estimated connectivity is shown in orange on the right. The top panel shows that due to a lack of explaining away, spurious correlation may be found for an unconnected pairs of nodes that each share variability with a third intermediate node. The middle panel shows that if multiple nodes are interconnected, shared variability may be attributed to the wrong pair of nodes. Finally, the bottom panel shows that methods such as purely pairwise approaches do not capture directionality, potentially missing crucial asymmetric structure in the connectome.

raphy (MEG), electroencephalography (EEG), and has been particularly prominent in the field of functional MRI [89]. Each of these modalities comes with benefits and drawbacks: EEG is relatively easy to collect and boasts decent temporal resolution (originally sampled up to 512 Hz and subsequently down-sampled to capture canonical frequency bands described in Chapter 2, up to 90kHz), but is limited to scalp-level signals and vulnerable to structured noise from conductivity leakage and muscular movement artifacts; fMRI measures the full brain volume with decent spatial resolution ($3 \times 3 \times 3\text{mm}^3$ voxels), but has very low temporal resolution (c.1Hz)[90][91]. In this chapter, I will use a novel statistical machine-learning framework, Union of Intersections (UoI, see Section 2), to generate multivariate estimates of functional connectivity for resting state EEG and fMRI data, and characterize the graph-theoretic properties of the resulting networks. In this chapter I will argue that UoI_{LASSO} , a regularized ensemble learning algorithm developed[58] and rigorously tested [55] in the Bouchard lab, provides a promising avenue for functional connectivity estimation that can be applied to diverse neuroimaging modalities.

3.2 Methods

Neural data

Neural data was provided by collaborators through the Gates Developmental Imaging Consortium. This section will outline the parameters of recording and pre-processing of data as reported by our collaborators.

EEG

EEG data was provided by Professor Laurel Gabard-Durnam and Dr. Carol Wilkinson. The work for the EEG section used data from the Infant Sibling Project (ISP), a longitudinal study of infants for the purposes of studying autism spectrum disorder (ASD). All participants were born full-term (i.e., at least 36 gestational weeks) to families with no known history of genetic disorders, and had at least one elder sibling at home. Participants were sorted into two groups: the low-risk control group (LRC) and those at high risk of autism (HRA). The basis of this designation was family history—LRC participants had no siblings or first- or second-degree family relations diagnosed with autism. Our analysis in this chapter exclusively utilized data from the LRC group.

The data was collected at the Boston Children’s Hospital in a light-attenuated and sound-isolated experiment room. Participants were instructed to sit on their caregivers’ laps and looked straight to the experimental screen in front of them,




Figure 3.2: The layout of the 128 channel sensor cap used for collecting the EEG dataset. The channels highlighted in blue are part of the 10-20 channel montage (most frequently used subset of channels that covers headspace with limited overlap in signal recording). Electrodes highlighted in yellow ones were also processed by HAPPE.

approximately 65 cm in between. Recording sessions lasted for between 2-5 minutes. Data was collected using NetStation 4.5 software (Electrical Geodesics, Inc., Eugene OR) with NetAmps 300 Amplifier and 128-channel Hydrocel Geodesic Sensor Nets (Electrical Geodesics, Inc., Eugene OR). For a better fit to infant patients, electrodes 125-128 were removed physically from the cap. The layout of the recording cap is shown in Figure 3.2.

The data was recorded at 500Hz and run through the Batch EEG Automated Processing Platform (BEAPP)[92] and the Harvard Automated Processing Pipeline for Electroencephalography (HAPPE)[93] to clean up artifacts and generate power spectral densities (PSD). The full HAPPE pipeline is discussed in full elsewhere, but the broad steps and specific parameters are listed here. A band-pass filter of 1Hz-

249Hz is applied to baseline EEG data. Channels to be processed by HAPPE are then sub-selected based on the layout of the electrode cap. The selected channels for ISP are highlighted in yellow (the International 10-20 electrodes, which achieve near-total coverage of the scalp) and blue (additional electrodes). Electrical noise (i.e., 60Hz noise) and high-impedance channels are removed from the dataset. Several stages of components analysis are then used to remove artifacts (e.g., scalp or muscle movements). Data is then evaluated for remaining artifacts in 2-second segments, and all segments containing artifacts are rejected. Recording sessions were excluded from the dataset if more than 20% of the selected channels were rejected, or if there were fewer than 40 seconds of usable (i.e., artifact-free) data.

Data was divided into canonical frequency bands theta (4-5.99Hz), low alpha (6-8.99Hz), high alpha (9-12.99Hz), beta (13-29.99Hz), low gamma (30-50Hz), and high gamma (65-90Hz). Our analysis focused primarily on the high-gamma band, but was also applied to theta, beta, low gamma, and low alpha.

fMRI

All fMRI data used in this chapter was drawn from the BAMBAM (Brown university Assessment of Myelination and Behavioral development Across Maturation) study, pre-processed and provided by Dr. Muriel Bruchhage. A lengthier description of the dataset is included in her paper[39]. BAMBAM is an ongoing longitudinal study of neurodevelopment in neurotypical, healthy children. Participating children were assessed approximated every 6-12 months. Each visit included multi-modal MRI (including the fMRI recordings that are the basis of our analysis), performance assessments, and parental reports. Children with risk factors such as in-utero alcohol exposure, preterm or otherwise medically complex births, or familial history of psychiatric or learning disorders were excluded from the study. MRI data was acquired from sleeping patients, without sedatives.

All neuroimaging data were acquired on a 3 T Siemens Trio scanner with a 12-channel head RF array. rsfMRI data were acquired during natural sleep with the following parameters: TE=34 ms, TR=2.5 s, flip angle=80°, field of view = 24 × 24cm², imaging matrix = 80 × 80, and 32 interleaved 3.6 mm slices (for a voxel resolution: 3×3×3.6mm³), BW=751 Hz/pixel, and GRAPPA acceleration factor of 2. 132 volumes were acquired for a total acquisition time of approximately 7:00 min. A total of 32 ROIs were defined using the CONN software toolbox for functional connectivity analysis[94]. For a visualization of the placement of the 32 ROIs used, see Figure 3.3[39].

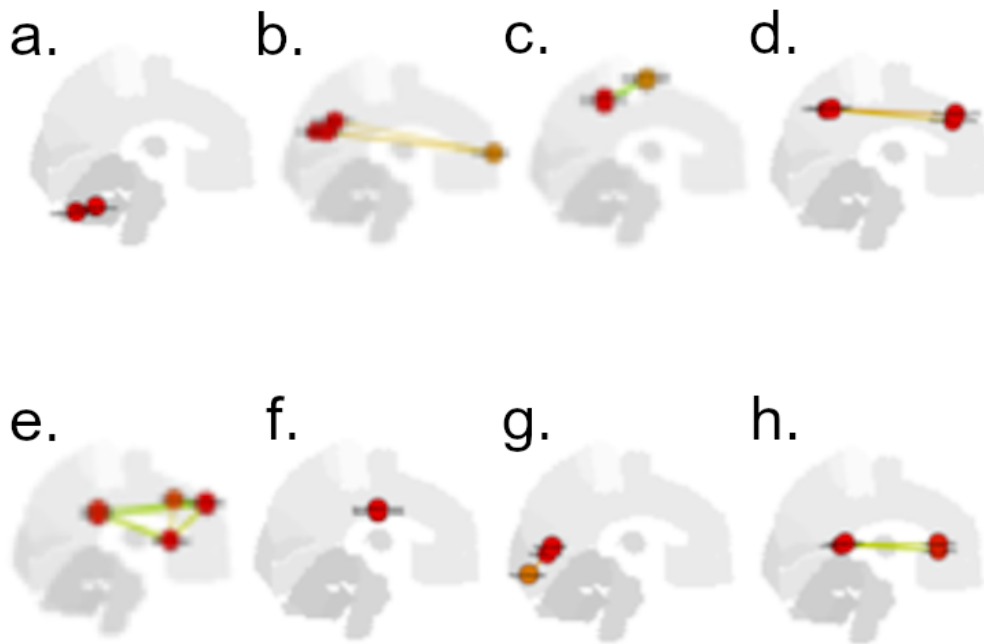


Figure 3.3: Resting state fMRI networks, figure courtesy of Dr. Bruchhage[39]. **a.** cerebellar: anterior, posterior. **b.** default mode: MPFC, medial prefrontal cortex; left and right LP, lateral parietal cortex; PCC, posterior cingulate cortex. **c.** dorsal attention: left and right FEF, frontal eye field; left and right IPS, inferior frontal gyrus. **d.** fronto-parietal: left and right LPFC, lateral prefrontal cortex; left and right PPC, posterior parietal cortex. **e.** salience: left and right anterior insula, left and right RPF, rostral prefrontal cortex; left and right SMG, supramarginal gyrus. **f.** sensorimotor: superior, left and right lateral. **g.** visual: medial, occipital, left and right lateral. **h.** language: left and right IFG, inferior frontal gyrus; left and right pSTG, posterior superior temporal gyrus.

Functional connectivity estimation

Multivariate coupling models

To fit a multivariate coupling model, I used the statistical framework, UoI_{LASSO} due to its stability, selectivity, and accuracy (see Background, Chapter 2, for a

discussion of Lasso and the UoI framework). Multivariate functional connectomes were estimated using the same algorithm for both EEG and fMRI.

I will define the time series of each neural signal, as measured by an EEG electrode or an fMRI ROI, as $n_i(t)$. Multivariate coupling was calculated independently for each time series in the recording session.

$$n_i(t) = \beta_{i0} + \sum_{j \neq i}^N \beta_{ij} n_j(t)$$

Pairwise functional connectivity was calculated as the Pearson correlation between the activity of each pair of signals in the recording session. Figure 3.4 shows the correspondence between the matrix of coupling coefficients and the resulting connectome graph.

Functional connectome creation and analysis

For the multivariate connectomes, I created both symmetrized and directed graphs. To create symmetrized graphs, edge weight w_{ij} between node i and node j was set as the mean $\beta_{ij} = \frac{\beta_{ij} + \beta_{ji}}{2}$. I analyzed the resulting graphs in order to quantify key comparisons between multivariate and pairwise FC. I quantified graph spatial distribution as Pearson correlation ρ of correlation coefficient w_{ij} to distance d_{ij} .

Small worldness ω and sparsity

I calculated small worldness, $\omega = \frac{L_r}{L} - \frac{C}{C_l}$, where L is the characteristic path length of the network, L_r is the characteristic path length for an equivalent random network, C is the clustering coefficient, and C_l is the clustering coefficient of an equivalent lattice network. $\omega \in [-1, 1]$, where ω close to -1 is a lattice graph, 1 is a random graph, and ω close to 0 is a small world graph. I also calculated sparsity, the fraction of zero edges in the graph for FC graphs. For pairwise graphs, I applied a series of thresholds $T \in [0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9]$ such that $|\beta_{ij}| < T$ were set to zero. Sparsity and ω were calculated for each thresholded graph.

One core strength of UoI_{LASSO} is its stable selectivity. In particular, the stability of support selection in UoI_{LASSO} is ensured by the intersection step of UoI. The algorithm is not constrained to a predetermined level of sparsity, instead discarding supports that are not consistently predictive across multiple folds of the data. Thus, the sparsity of a multivariate UoI_{LASSO} network is a parameter that is inferred from the data itself, and can be examined as a potentially salient feature of the network. In order to compare ω values between multivariate and pairwise graphs, I considered the




Figure 3.4: Connectomes were generated by setting $A_{ij} = \beta_{ij}$, where A_{ij} is the adjacency matrix of the graph and β_{ij} is the coupling coefficient between electrodes i and j .

$\omega_{pairwise}$ for the thresholded network closest in sparsity to the multivariate network for the same subject, and conducted a paired-sample t-test for each filter-band and age group.

Correlation of coupling strength to node distance

In order to quantify the dependence of coupling strength between two regions on spatial proximity, for each pair of nodes n_i and n_j I calculated the Pearson correlation coefficient ρ_{ij} between distance d_{ij} between the sensors and coupling weight β_{ij} .

Connectome templates

In order to identify trends in functional coupling, I created aggregate connectome templates by taking a median of the functional connectomes for participants within a given 90-day age range. Because the ages of participants were unevenly distributed, templates were calculated only for ranges with 12 recordings or more. The $n = 12$ was chosen to balance the goal of producing templates spanning the full age range of available subjects while encompassing data from as many subjects as possible. If more than 12 subjects fell within a given age range, the appropriate number of subjects were selected at random. Templates were calculated as the median of the weight matrices, so that $\beta^{template} = M(\{\beta\}^{subjects})$. The median was selected rather than the mean in order to preserve the selectivity of the networks. Using the median β means that if more than half of the values in $\{\beta_{ij}\}^{subjects}$ are zero, then the aggregate template $\beta_{ij}^{template}$ will also be set to zero. Conversely, if more than half of the values in $\{\beta_{ij}\}^{subjects}$ are high, then the aggregate template $\beta_{ij}^{template}$ will also be high. In this way, I preserve sparsity in the aggregate connectome, while also preserving high-power functional connections.

To examine the stability of this method of generating aggregate templates, I analyzed age ranges which contained $n > 12$ subjects. For each such age range, I aggregated templates for all n choose 12 sets of 12 subjects, and calculated the standard deviation for each weight β_{ij} .

I further examined variance of edge weights within 90 day age window by calculating the standard deviation, mean, and median of all edge weights within all age windows with $n \geq 12$ subjects. Within each age group, I calculated the Spearman correlation coefficient between the standard deviation of each edge weight and the absolute value of the mean and median of the edge weight. The results are reported in Table 3.1.

fMRI network interpretation and Mullen score comparison

I consulted fMRI functional connectivity literature[63] and collaborated with fMRI domain expert Dr. Muriel Bruchhage to identify pairs of ROI pairs of particular interest in the prediction of early childhood development. The Mullen Scales of Early Learning (MSEL) provide scores in the areas of expressive language, receptive language, fine motor, gross motor, and visual reception, as well as the verbal composite, nonverbal composite, and early learning composite (ELC). I visually examined correlations between Mullen scores and ROI pairs identified from visual inspection of the aggregate template networks.

I examined β_{ij} for ROI pairs that have been shown in the literature[39][61][34] to be relevant to a given area of cognition. The functional couplings examined for each of the areas measured by the MSEL are listed below.

For receptive language, I examined functional connectivity of posterior superior temporal gyrus (pSTG) in the language network (Wernicke’s area), both incoming and outgoing. I examined contributions of the salience network (SN), particularly the anterior insula, which is considered a hub of outgoing top-down signals for salient signal selection, and the supramarginal gyrus (SMG), which is implicated in phonological processing[95]. I also visualized contributions of the frontoparietal network (FPN) to the pSTG and IFG[96], especially lateralized contributions. For expressive language, I examined functional connectivity of inferior temporal gyrus (ITG) in the language network (Broca’s area). For visual reception, I examined functional connectivity of the visual network, particularly contributions from the dorsal attention network (DAN) and the SN. For fine and gross motor scores, I examined the functional connectivity of the cerebellar network and the sensorimotor network. I also examined contributions of the visual network, to consider the possibility of sensory feedback, and the frontoparietal network, for top-down influence of task-oriented behavior. I also considered functional connectivity between sensorimotor and dorsal attention networks due to evidence in the literature implicating the dorsal attention network in sensorimotor inhibition[97].

In order to find specific ROI pairs that showed functional coupling consistently over time, I visualized our pre-selected pairs of networks for 3-month aggregate templates. **Note:** for clarity of interpretation, it is worthwhile to reiterate that in all included matrix visualizations in this chapter, row i represents the contributions to the i th ROI by all other ROIs $j \neq i$. That is, for multivariate networks, the ROIs on the x-axis of coupling matrices are the regressors, while the ROIs on the y-axis are the targets. For this reason, I will discuss the ROIs on the x-axis as “contributing to” those on the y-axis. Since the pairwise coupling networks are symmetric, this distinction does not apply to them.

Functional connectivity correlates of MSEL scores

I visualized the correlations of functionally relevant ROI pairs for the areas of the Mullen assessment. The primary purpose of this analysis was exploratory in order to identify features of the functional connectome that might subsequently be used for Mullen score prediction. Some ROI pairs of interest are shown in Figure 3.14 for fine motor scores, Figure 3.15 for receptive language scores, and Figure 3.16 for early learning composite scores, along with the linear regression fit. Table 3.2 reports

accompanying Spearman correlations with all zero-valued weights removed before the correlation is calculated.

I calculated the Spearman correlation for each MSEL score to all edge weights in the connectome, dropping self-connections (the diagonal) and zero-valued datapoints. P-values were corrected for false discovery rate using the Benjamini-Hochberg procedure. Figure 3.17 shows the correlation matrices for Expressive Language scores, while equivalent figures for the other scores can be found in Appendix A.

3.3 Results

Multivariate estimates of functional connectivity for EEG produce sparse, small-world, spatially distributed networks

Functional connectivity in pairwise graphs displayed coupling coefficients that correlated strongly with proximity between electrodes (Fig 3.5e-f). Furthermore, pairwise graphs consistently displayed “daisy chaining” connections, such that if proximate electrodes A and B display powerful correlations, and B shares a powerful correlation to electrode C, electrode A will also appear to couple strongly to electrode C. This effect is an inevitability of pairwise FC estimation, because shared variability is not “explained away” once accounted for by another coupling connection. Thresholding, a common practice in FC estimation, does not address this shortcoming. In fact, thresholding serves only to preserve the strong connections of proximate electrodes, while eliminating smaller, but potentially meaningful coupling. Fig 3.5d demonstrates how pairwise FC graphs grow closer and closer to a lattice structure ($\omega < 0$) as they are thresholded to higher levels of sparsity.

EEG data suffers from significant signal spread, resulting in high multi-collinearity between signals from nearby electrodes. Pairwise estimates of functional connectivity fail to explain away shared variability, resulting in highly redundant spurious connections. Utilizing a multivariate approach, such as UoI_{LASSO} , explains away covariance as a part of the fitting process, reducing the dominance of short-range, redundant connections. As a result, multivariate graphs can reveal weaker but still significant connections that would be lost in thresholding for pairwise networks. Furthermore, UoI_{LASSO} is regularized and selective, and thus resulted in sparse networks without sacrificing accuracy. As shown in Fig 3.5d and Fig 3.6e, these sparse multivariate networks were consistently small-world (ω close to 0). A paired t-test showed a highly statistically significant difference between ω distributions for multivariate versus pairwise connectomes ($p \ll 0.0001$, $n = 42$) for all frequency bands and age groups. The small-worldness of multivariate aligns with previous work characterizing

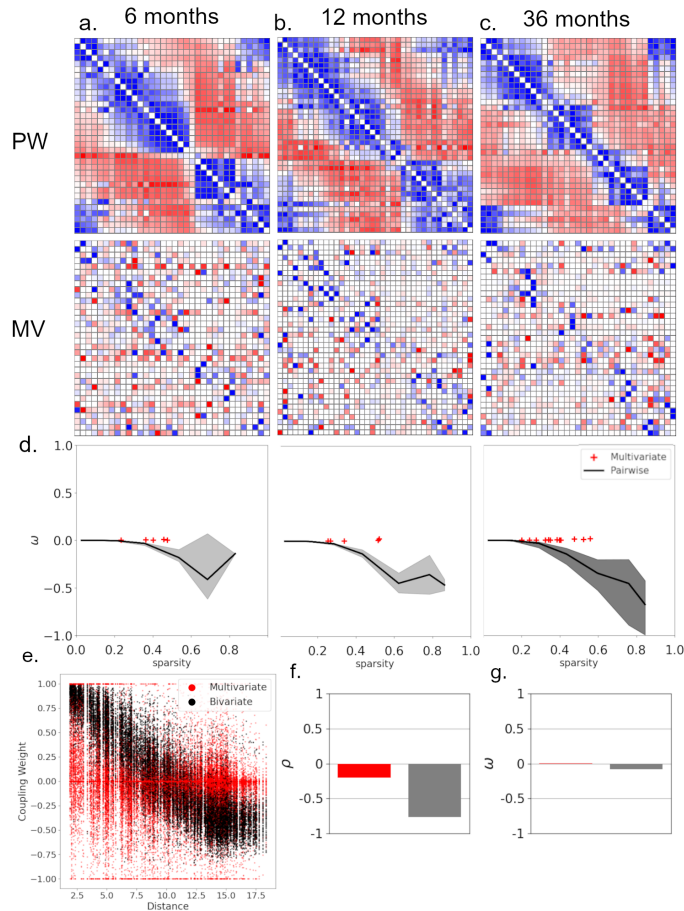


Figure 3.5: **Multivariate coupling models produce sparse, small-world, spatially distributed functional connectivity networks in resting state EEG data.** **a.** Comparison of a weight matrix for a representative example of the high-gamma frequency band component of resting-state EEG data collected from a 6 month old subject. Pairwise FC (top) is dense, and coupling strength is heavily dependent on the physical proximity between each pair of electrodes. **b.** high-gamma resting state, 12 month old subject. **c.** high-gamma resting state, 36 month old subject. **d.** Small worldness, ω vs. sparsity for multivariate and thresholded pairwise graphs. **e.** Coupling weight as a function of electrode distance for high-gamma networks for all subjects aged 6-36 months. **f.** Median Pearson correlation coefficient ρ of coupling strength to electrode-distance for all high gamma subjects. **g.** Median ω for all high gamma subjects.

brain networks as innately small-world[98][99], and encouraging indicator that UoI multivariate EEG connectomes reflect underlying properties of the brain.

Multivariate connectomes were highly spatially distributed compared to pairwise estimates. While multivariate FC does select strong supports from proximate electrodes, it also shows more strong, long-distance connections (Fig 3.5e-f). I calculated the correlation coefficient of coupling weights β_{ij} to node distance d_{ij} for each subject. Paired t-tests for each frequency band and age group shows a highly significant difference between pairwise and multivariate groups ($p < 0.0001, n = 42$).

Multivariate fMRI functional connectomes show sparse, asymmetric network activations

Pairwise functional connectivity estimation for fMRI did not display the same proximity-dependence or lattice structure as EEG. Pairwise fMRI connectomes show spatial dependence of coupling similar to multivariate networks. Correlation ρ between ROI distance and coupling strength did not vary significantly between multivariate and pairwise networks (t-statistic = 0.252, $p = 0.8$).

Furthermore, thresholding pairwise networks did not consistently produce lattice (ω close to -1) networks as seen in the EEG data (Fig 3.5d). As seen in Fig 3.7f, at levels of sparsity close to their multivariate counterparts, pairwise graphs are relatively small-world (ω close to 0). However, while the difference in small-worldness between multivariate and pairwise connectomes was not as striking as for EEG connectomes, a paired t-test shows that the difference is statistically significant (t-statistic=6.03, $p=2.11e-08$).

Multivariate fMRI templates show stable selectivity for many nodes, with some highly variable couplings

Figure 3.8 demonstrates a core difference between multivariate and pairwise estimates of connectomes. The standard deviation σ of bootstrapped aggregate age templates for multivariate networks showed a high percentage of connections with highly stable coupling values (σ close to 0), with a long tail showing some highly variable connections (Fig 3.8). In particular, the coupling to visual functional network appeared to vary strongly across target networks, with particularly variable couplings to the lateral visual fields. A Kolmogorov-Smirnov test (KS test) between the standard deviation distributions shown in Figure 3.8d-f show a statistically significant difference between pairwise and multivariate templates for all eligible age groups ($p < 0.0001, n = 992$).

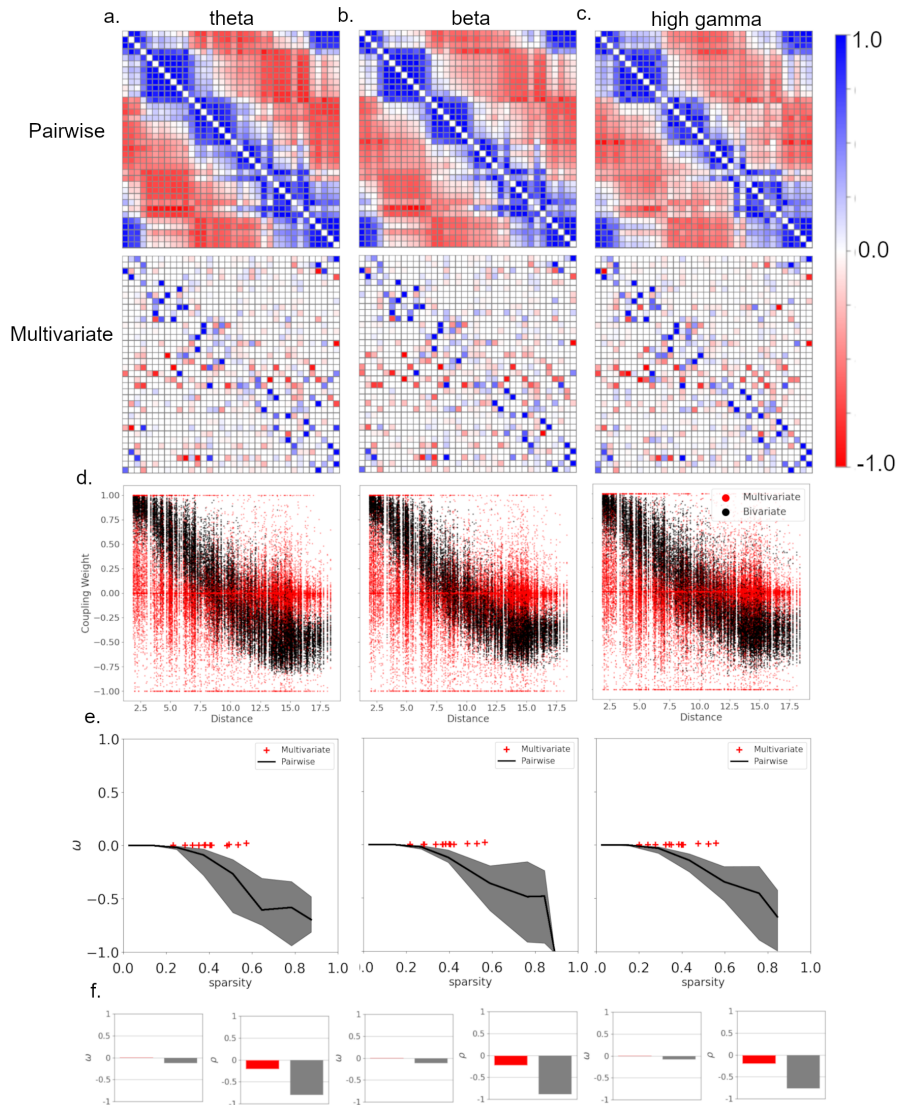


Figure 3.6: **Multivariate EEG functional connectomes are sparse, small-world, spatially distributed across frequency bands.** **a-c.** Functional connectivity network (pairwise top, multivariate bottom) for 36-month old example subject shown for different frequency bands: **a.** theta **b.** beta **c.** high gamma. **d.** Functional coupling vs. node distance for each frequency band. **e.** ω vs network sparsity for each frequency band. **f.** bar graphs of median ω and ρ for each frequency band.

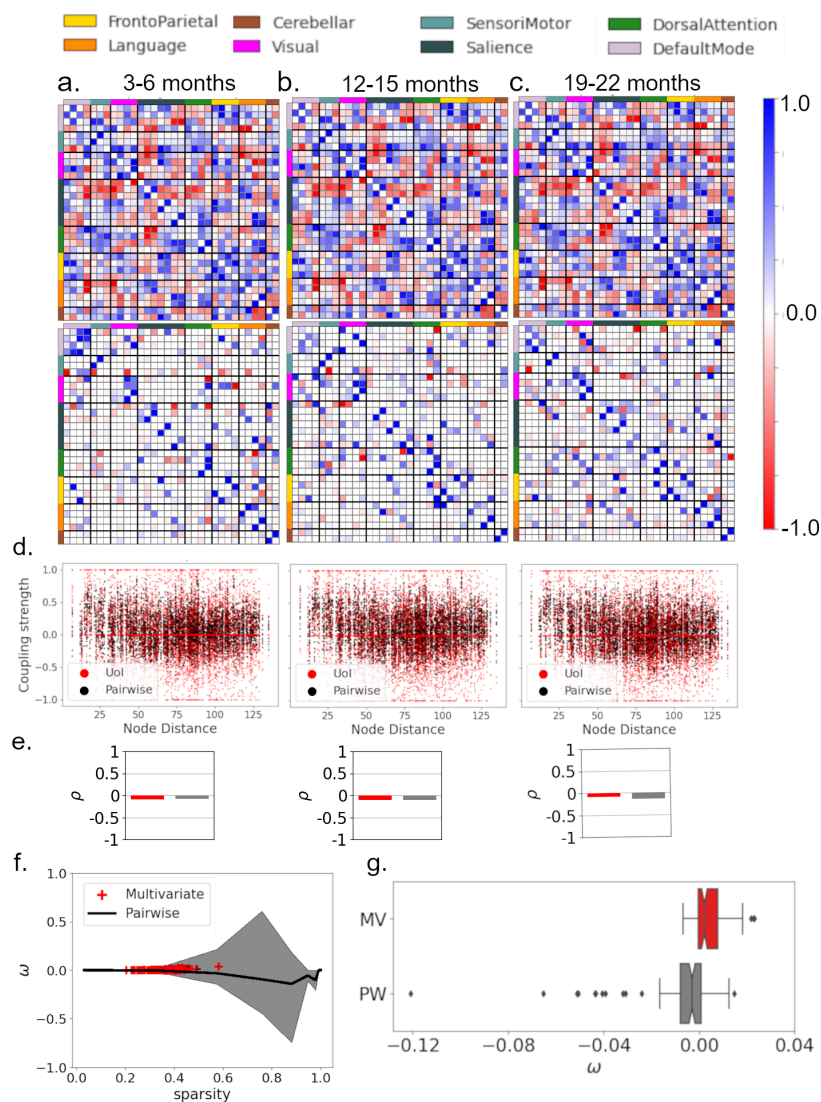


Figure 3.7: **fMRI functional connectomes are small-world and spatially distributed for both pairwise and multivariate estimates.** **a.** Median functional connectivity template for subjects aged 3-6 months. $n = 12$. **b.** Template for 12-15 months. **c.** Template for 19-22 months. **d.** Coupling strength vs. ROI distance for subjects in corresponding age group. **e.** Median Pearson correlation coefficient ρ between coupling strength and ROI distance. **f.** Small worldness, ω vs. sparsity for multivariate and thresholded pairwise graphs. **g.** Distribution of ω , where pairwise ω is calculated on a network thresholded to the level of sparsity of the corresponding UoI_{LASSO} network.

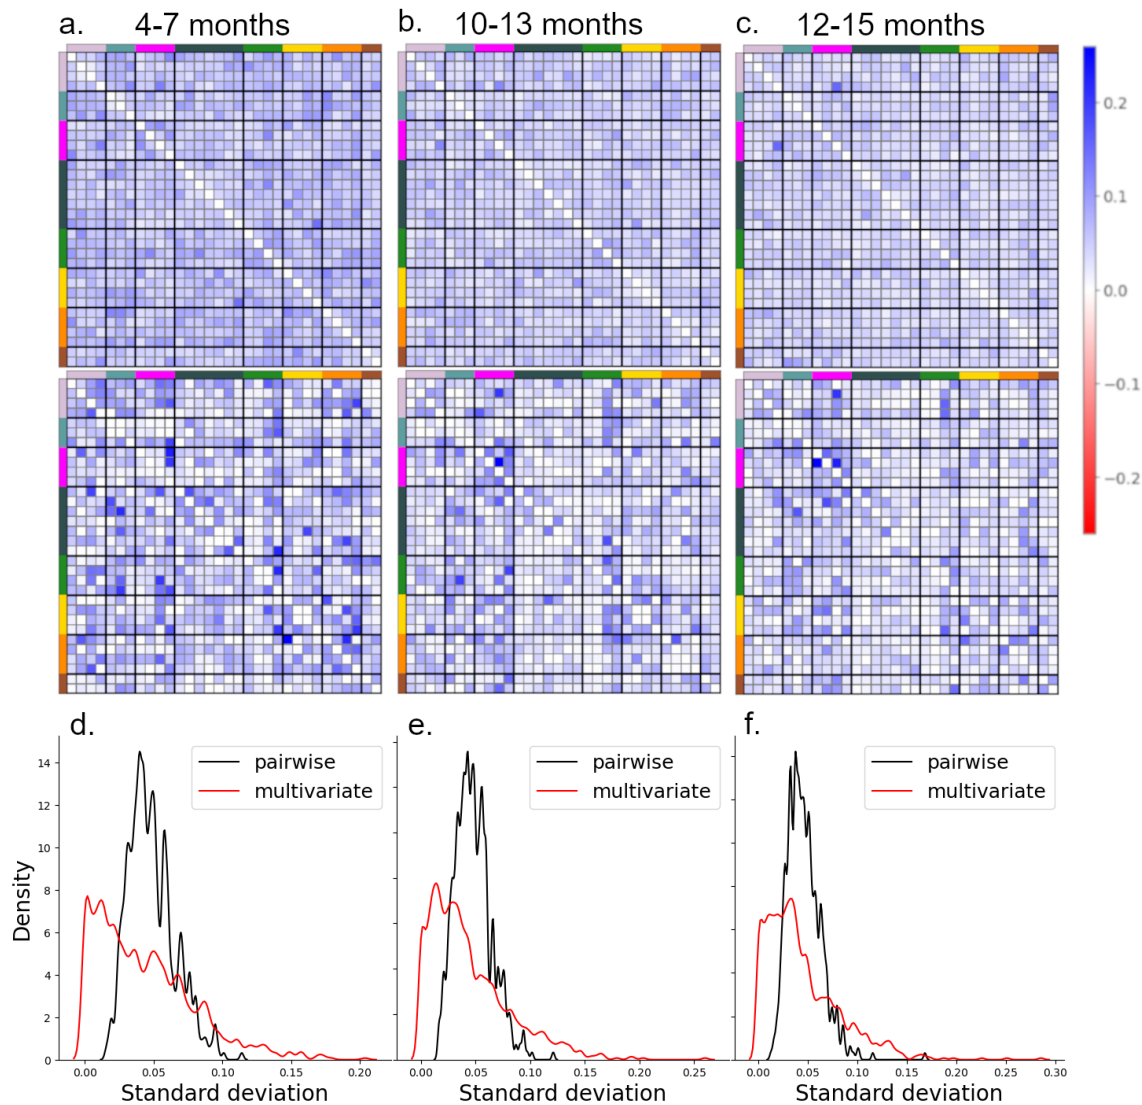


Figure 3.8: **Multivariate fMRI connectome templates show a large fraction of stable couplings, but also some highly unstable couplings.** **a-c** Standard deviation of bootstrapped aggregate templates for pairwise (top) and multivariate (bottom) functional networks, with $n = 12$ subselected from N available subjects. **a.** 120-210 days, $N = 23$, **b.** 300-390 days, $N = 20$, **c.** 360-450 days, $N = 25$. **d-f** kde-smoothed distribution of standard deviations for listed age ranges.

Furthermore, Figure 3.9 and Table 3.1 show that the standard deviation of multivariate connectome weights was significantly correlated with the magnitude of the mean and median value of that edge weight. This suggests that highly weighted ROI-ROI connections are also prone to greater variability.

Multivariate fMRI functional connectomes show stably sparse connections and reveal lateralization through negative couplings

While graph theoretic measures indicate that thresholded pairwise connectomes have similar small worldness and spatial distribution to multivariate connectomes (Fig 3.7), the structure of coupling between functional networks differed greatly between pairwise and multivariate estimates. In multivariate aggregate age-group templates, positive intrahemisphere connectivity was a common theme. In particular, strong positive connections on ipsilateral pathways and conversely strong negative connections on contralateral pathways.

Consider Fig 3.10, which showed sensorimotor coupling to dorsal attention. The multivariate connectomes showed a locus of contributions from intraparietal sulcus (IPS) to lateral sensorimotor fields, with negative contralateral coupling and positive ipsilateral connections. The relative consistency of this locus between age groups may reflect the relatively early functional synchronization of the sensorimotor compared to other networks[82]. By contrast, pairwise connectomes lose the negative contralateral connections altogether, while also placing less total power in the ipsilateral connections.

Figure 3.11 shows another example between pSTG, which contains Wernicke’s area, and supramarginal gyrus, which has been shown to support word processing[100]. Between 3-6 months and 19-22 months, the lateral connections develop further, potentially suggesting maturation of the salience network. Anterior insula also showed positive ipsilateral coupling to IFG, a region that contains Broca’s area, which is important in speech production[79]. As discussed in Chapter 2 background, the dorsal anterior insula in the salience network acts as a hub that selectively exerts top-down control on other major networks[67]. These strong, consistent connections are promising for potential prediction of language skills.

Figure 3.12 is challenging to interpret due to highly inconsistent connections over time. For example, the positive ipsilateral coupling of left-side pSTG to posterior parietal cortex (PPC) appeared in the first two age brackets (so 3 months to 9 months), but not in older children 19-22 months. PPC has been shown to mediate manipulation of working memory in cognitively demanding tasks[101]. The nature of this connection is difficult to assess in sleeping infants, but may merit further examination. The only consistent connection in the language-frontoparietal network

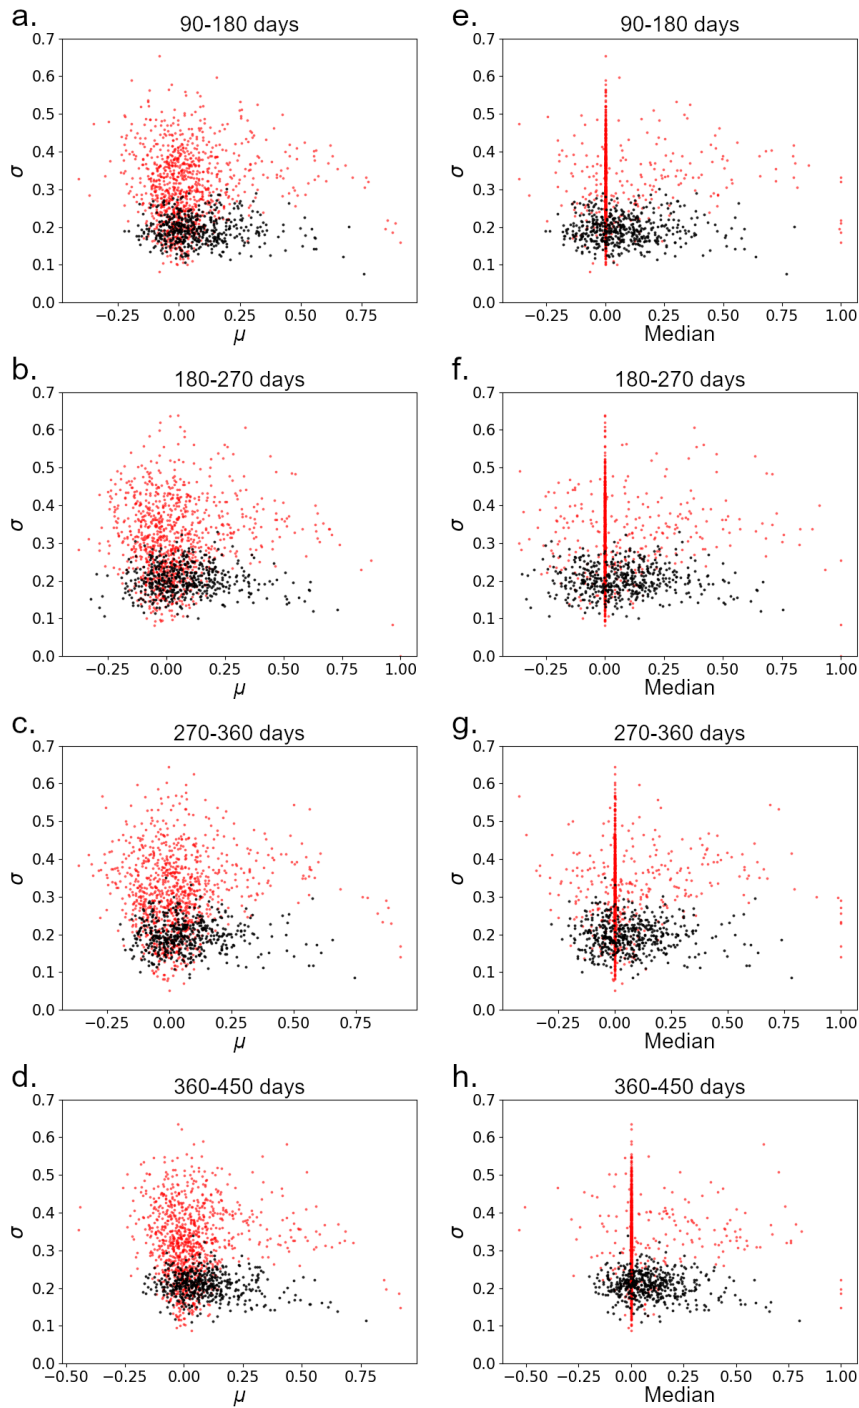


Figure 3.9: **Standard deviation of edge weights scales with mean edge weight.** **a-d:** standard deviation plotted against mean edge weight for age groups **a.** 90-180 days. **b.** 180-270 days. **c.** 270-360 days. **d.** 360-450 days. **e-h:** standard deviation plotted against median edge weight for age groups **e.** 90-180 days. **f.** 180-270 days. **g.** 270-360 days. **h.** 360-450 days.

<i>Ages</i>	<i>n</i>	Multivariate		<i>mean</i> ρ	<i>mean</i> <i>p-val</i>
		<i>median</i> ρ	<i>med</i> <i>p-val</i>		
90-180	19	0.064	0.0449	0.267	1.2E-17
120-210	23	0.092	0.0039	0.264	3.1E-17
150-240	16	0.097	0.0022	0.277	7.2E-19
180-270	13	0.098	0.0020	0.261	5.7E-17
210-300	12	0.123	0.0001	0.273	2.1E-18
240-330	14	0.129	0.0000	0.293	3.8E-21
270-360	14	0.091	0.0041	0.256	2.6E-16
300-390	20	0.087	0.0064	0.289	1.6E-20
330-420	25	0.063	0.0482	0.287	3.2E-20
360-450	25	0.051	0.1109	0.303	1.9E-22
390-480	15	0.071	0.0249	0.259	1.2E-16
570-660	12	0.182	0.0000	0.345	4.4E-29

<i>Ages</i>	<i>n</i>	Pairwise		<i>mean</i> ρ	<i>mean</i> <i>p-val</i>
		<i>median</i> ρ	<i>med</i> <i>p-val</i>		
90-180	19	0.005	0.8840	-0.018	0.5658
120-210	23	-0.056	0.0806	-0.082	0.0101
150-240	16	0.050	0.1177	0.014	0.6704
180-270	13	0.055	0.0862	-0.040	0.2083
210-300	12	0.006	0.8562	-0.059	0.0628
240-330	14	0.004	0.9070	-0.008	0.7912
270-360	14	0.019	0.5502	0.025	0.4398
300-390	20	-0.114	0.0003	-0.132	0.0000
330-420	25	-0.075	0.0177	-0.064	0.0440
360-450	25	-0.110	0.0005	-0.079	0.0125
390-480	15	-0.024	0.4533	-0.114	0.0003
570-660	12	-0.037	0.2486	-0.045	0.1525

Table 3.1: **Standard deviation of connectomes edge weights shows significant correlation to magnitudes of corresponding mean and median weight values for multivariate connectomes.** Spearman correlation coefficients of standard deviation with absolute value of mean and median edge weights for multivariate and pairwise connectomes within 90 day age windows.

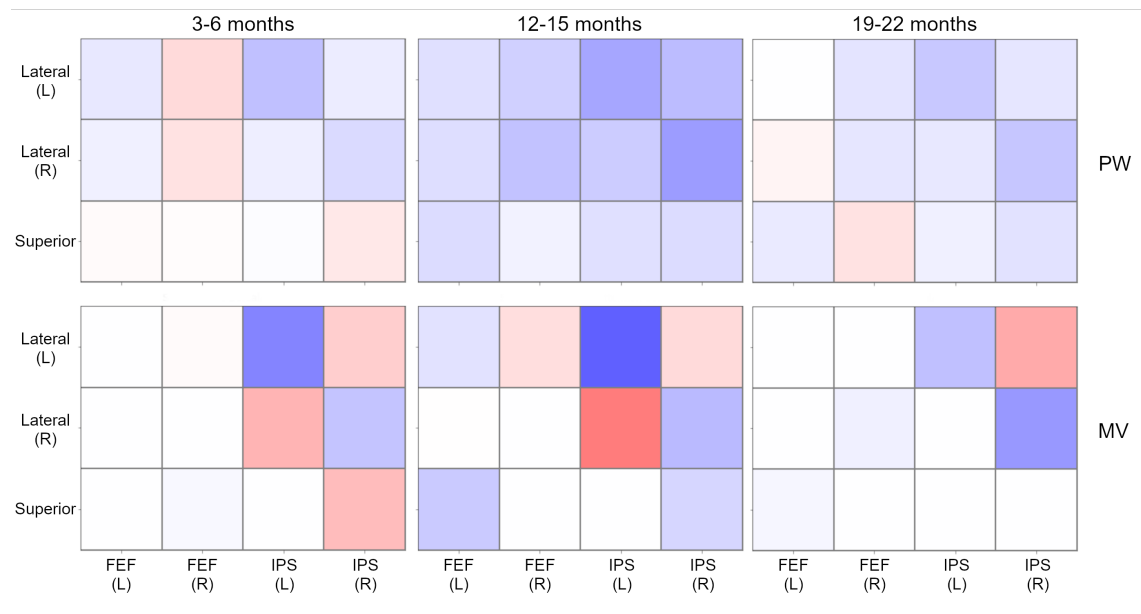


Figure 3.10: **Contributions of dorsal attention network to sensorimotor network show lateralized connectivity:** Multivariate functional connectivity networks showcase strong negative contralateral weights from dorsal attention intraparietal sulcus (IPS) to lateral sensorimotor during periods of intensive motor learning.

pair was a negative ipsilateral right-side coupling between the pSTG and the lateral prefrontal cortex (LPFC). LPFC is widely considered to be a key region in cognitive control tasks, and in some cases right-side LPFC activation has been observed to correlate negatively with task performance[102]. I identified this consistent negative weight as a candidate for Mullen score prediction.

Finally, Figure 3.13 once again showed strong connectivity of the pSTG, specifically contributions from lateral parietal default mode network. Positive ipsilateral and negative contralateral structure show a consistent right-lateralized connections that persist across age groups. This finding appears to be at odds with the known left-side lateralization of the language network[81].

In all the examples discussed above, pairwise connectomes shared the same one or two strongest positive ipsilateral connections, but not the nearby contralateral negative correlations that appeared in multivariate connectomes. The contralateral connections that are distinctly negative in multivariate estimates tend to be inconsistent in pairwise estimates, often near zero or faintly positive. The hemispheric structure




Figure 3.11: **Contributions of salience network to language network show lateralized structure for IFG and pSTG:** Multivariate functional connectivity shows positive ipsilateral connectivity and negative contralateral connectivity between salience network supramarginal gyrus and language network pSTG (Wernicke’s area). Anterior insula contributes strongly to IFG (Broca’s area).

of the multivariate connectomes seems in line with published literature[103].

Selectivity of multivariate functional connectivity estimates improves visual interpretability of connectome, but may not improve correlation with behavioral measures

In examining the estimated functional connectomes, I found that the multivariate and pairwise networks showed very different core structures, as discussed above. The sparse feature selection of the multivariate connectomes made it significantly easier to identify dominant edges in the network, resulting in easier interpretation during visual inspection of functional networks.

However, as shown by Figure 3.17, pairwise connectomes show much stronger correlation with raw Mullen scores, while multivariate connectomes show no statistically significant correlations. In particular, significantly correlated edges appear in clusters for pairs of functional networks, such as language-default mode and sensorimotor-

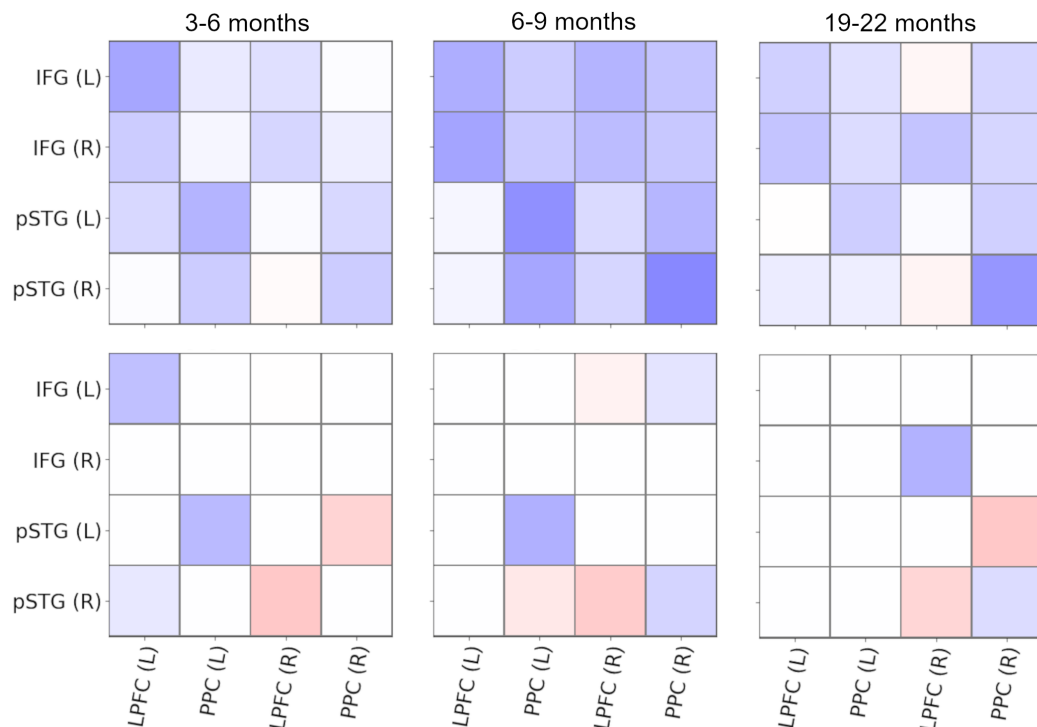


Figure 3.12: **Contributions of frontoparietal network to language network show evolving contributions to pSTG.** Both ipsilateral and contralateral couplings can be seen, but pSTG-PPC appears more likely to show positive ipsilateral coupling and negative contralateral coupling.

visual. This structure is replicated across the other Mullen areas (fine motor, gross motor, receptive language, and visual). By contrast, MSEL t-scores are not significantly correlated with either multivariate or pairwise connectomes.

I plotted Mullen scores against connectome edges that stood out during visual inspection of aggregate connectome templates (as in Figures 3.10-3.12). Figure 3.14 shows examples of connectome edges that correlate with the Fine Motor raw score, particularly between the sensorimotor network and the right-side dorsal attention network, and the cerebellar network to the salience network. Figure 3.15 shows consistently higher correlations for pairwise connectomes than for multivariate connectomes, particularly for language-default mode connections. Figure 3.16 shows that left-side lateralized correlations of connections between frontoparietal and default mode networks with Early Learning Composite for multivariate connectomes.

Multivar.							
	Network	Region	Connect. network	Connect. region	corr	<i>p-val</i>	n
<i>(a)</i> <i>Fine motor raw</i>	Cereb.	Ant.	SN	A Ins. (L)	0.32	0.0103	63
	SMN	Sup.	DAN	IPS (R)	-0.37	0.0018	70
	SMN	Sup.	DAN	IPS (L)	-0.11	0.3599	72
	SMN	Lat. (L)	DAN	FEF (L)	-0.21	0.0676	77
<i>(b)</i> <i>Receptive language raw</i>	Lang.	IFG (L)	SN	A Ins.(R)	-0.13	0.2791	71
	Lang.	pSTG (R)	FPN	LPFC (R)	-0.29	0.0110	75
	Lang.	pSTG (L)	DMN	MPFC	0.26	0.0283	71
	DMN	LP (L)	Lang.	pSTG (R)	0.04	0.7059	80
<i>(c)</i> <i>Early learning comp.</i>	FPN	PPC (L)	DMN	LP (L)	-0.14	0.1559	98
	FPN	PPC (R)	DMN	LP (R)	0.07	0.4916	89
	FPN	LPFC (L)	DMN	LP (L)	-0.27	0.0244	68
	FPN	LPFC (R)	DMN	LP (R)	0.19	0.1233	65
Pairwise							
	Network	Region	Connect. network	Connect. region	corr	<i>p-val</i>	n
<i>(a)</i> <i>Fine motor raw</i>	Cereb.	Ant.	SN	A Ins. (L)	0.15	0.1168	109
	SMN	Sup.	DAN	IPS (R)	-0.22	0.0245	109
	SMN	Sup.	DAN	IPS (L)	-0.15	0.1209	109
	SMN	Lat. (L)	DAN	FEF (L)	-0.14	0.1529	109
<i>(b)</i> <i>Receptive language raw</i>	Lang.	IFG (L)	SN	A Ins.(R)	-0.18	0.0618	109
	Lang.	pSTG (R)	FPN	LPFC (R)	-0.18	0.0546	109
	Lang.	pSTG (L)	DMN	MPFC	0.37	0.0001	109
	DMN	LP (L)	Lang.	pSTG (R)	0.27	0.0048	109
<i>(c)</i> <i>Early learning comp.</i>	FPN	PPC (L)	DMN	LP (L)	-0.07	0.4668	109
	FPN	PPC (R)	DMN	LP (R)	-0.08	0.4121	109
	FPN	LPFC (L)	DMN	LP (L)	0.00	0.9632	109
	FPN	LPFC (R)	DMN	LP (R)	-0.12	0.2066	109

Table 3.2: Spearman correlation coefficients and p-values for selected Mullen scores and ROI-ROI coupling weights for all plots show in Figures 3.14, 3.15, 3.16.




Figure 3.13: **Contributions of default mode network to language network show lateralized contributions from lateral parietal default mode regions to right pSTG.** Default mode network seems to contribute more consistently to Wernicke’s area (STG) than to Broca’s area (IFG).

Table 3.2 summarizes the correlation values and (un-adjusted) p-values for the scatter plots displayed. Based on these plots, strongly activated edges in the multivariate connectomes do not signify a higher correlation with MSEL scores.

The weak correlations shown in these plots is further verified by Figure 3.17, and Figures .1, .2, .3, and .4 in Appendix A, which show virtually no statistically significant correlations between multivariate connectomes and MSEL raw scores. Pairwise connectomes, by contrast, are significantly correlated with Mullen raw scores. In particular, these significant edges are tightly clustered within specific network pairs, such as language-DMN, language-DAN, and DAN-visual. There are also numerous intra-network correlations such as DMN and DAN. The clustering of significant edges may be due to the lack of explaining away of variance in the pairwise networks, thus repeating shared significant variability. Further analysis is necessary to better ex-

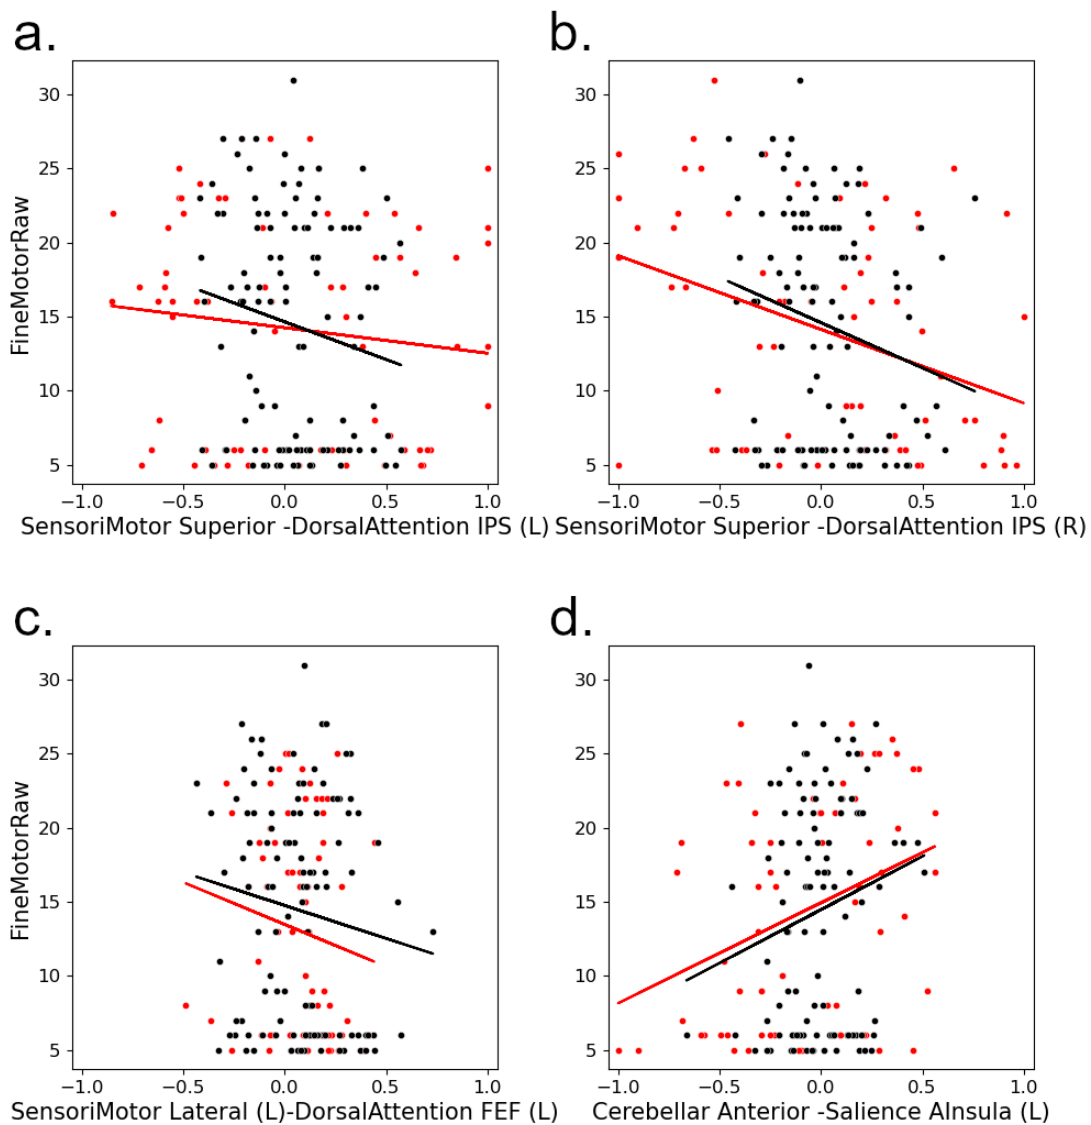


Figure 3.14: **Fine motor raw scores are correlated with select sensorimotor, dorsal attention, cerebellar, and salience network connections.** Scatter plots of Fine Motor raw score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black.

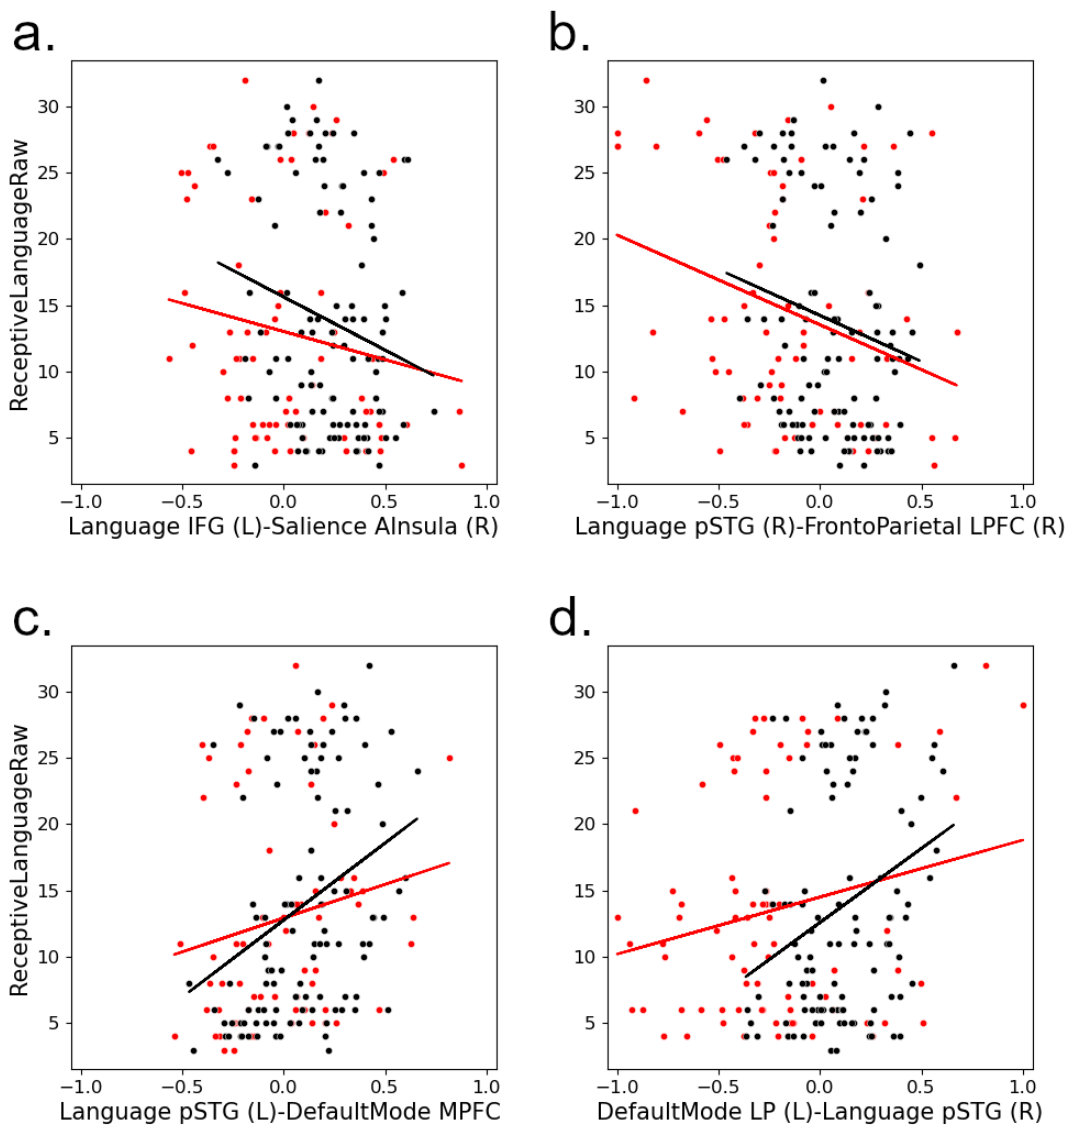


Figure 3.15: **Receptive language scores are correlated with frontoparietal contributions to the language network.** Scatter plots of Receptive Language raw score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black.

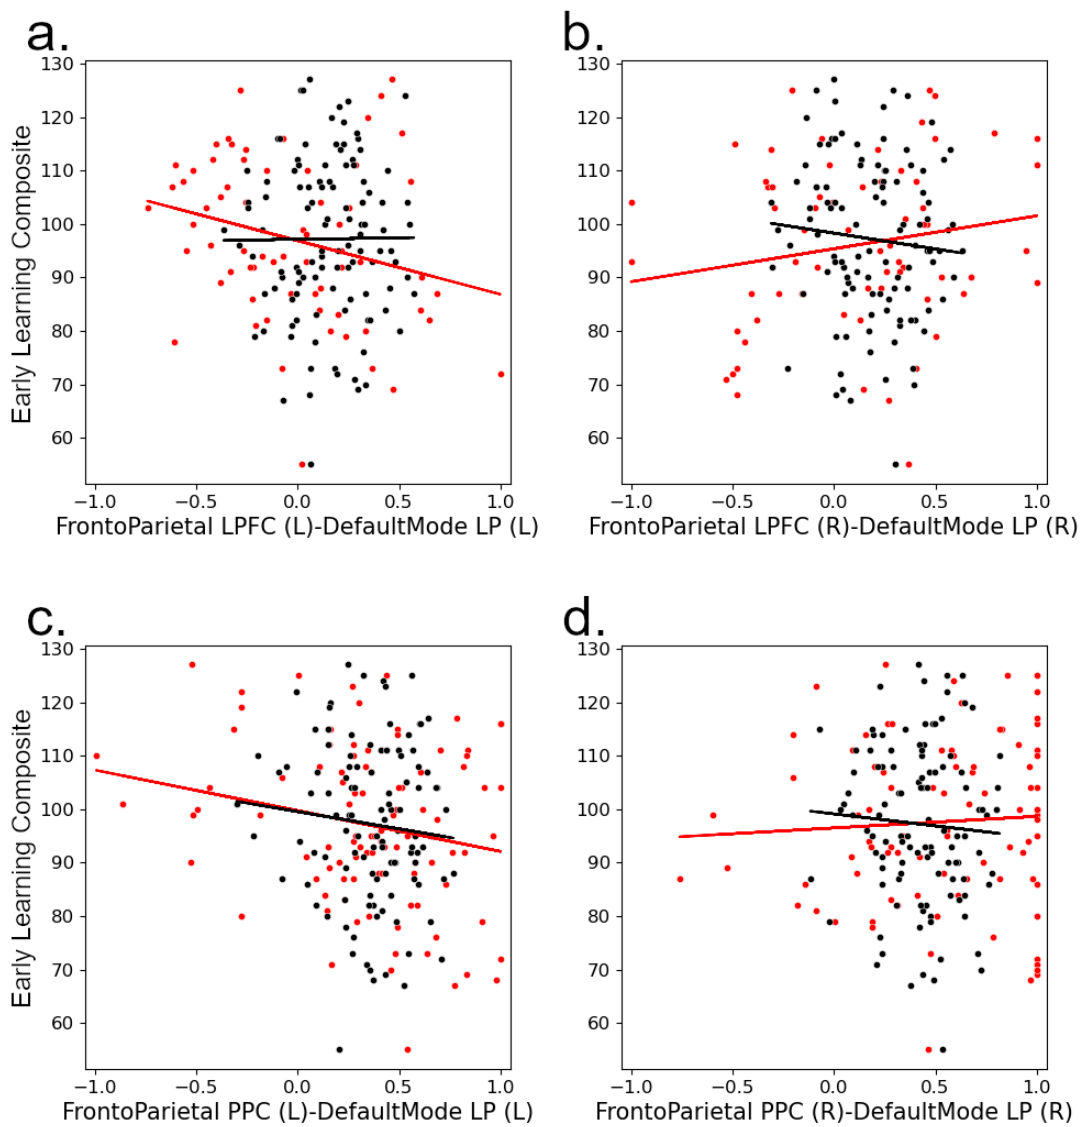


Figure 3.16: **Early learning composite score is correlated with frontoparietal-default mode couplings.** Scatter plots of Early Learning Composite score plotted against example connectome edges, with linear fits. Multivariate connections are shown in red, pairwise connections are in black.

plain the discrepancy between pairwise and multivariate connectomes—whether, for example, it is due to a flaw in the procedure for generating multivariate networks.

3.4 Conclusion

Functional connectomes were generated from resting state EEG and fMRI data collected from young children, using two competing methods: pairwise Pearson correlations and UoI_LASSO multivariate regression. Pairwise and multivariate connectomes were compared for both EEG and fMRI.

EEG connectomes were generated for subjects between the ages of 6 months and 36 months, grouped by three month intervals. Connectomes were further generated within canonical frequency bands beta, theta, and high gamma. Across age groups and frequency bands, multivariate EEG connectomes showed marked structural differences from their pairwise counterparts. Multivariate networks showed varying levels of natural sparsity ranging between 20% and 60% of connections being set to zero. All multivariate connectomes showed consistent small-worldness across subjects, age groups, and frequency bands. By contrast, pairwise networks tended towards more lattice-like graphs, suggesting that pairwise connectomes are heavily impacted by volume conduction causing signal leakage. This observation was reinforced by the fact that multivariate connectomes were significantly more spatially distributed than pairwise, showing a more even balance between long- and short-range connections. A regularized multivariate approach to estimating functional connectivity thus seems to be a marked improvement over pairwise correlation for EEG data.

fMRI connectomes were generated for 109 subjects aged between 90 and 903 days (3 to 30 months). Graph analysis of fMRI multivariate connectomes also showed consistent small-worldness, spatial distribution, and 20-60% sparsity. However, thresholded pairwise correlation connectome also showed small-worldness and spatial distribution. Since fMRI data came with the advantage that all time series correspond to well-characterized regions of interest that can be easily compared across subjects, for 3-month-interval age groups I created aggregate connectomes by taking the median of the connectomes of 12 subjects within the specified age range. I then used these age templates to look for interesting connective structures that characterize multivariate and pairwise networks. I found that multivariate connectome templates showed frequently occurring patterns of positive coupling between ipsilateral regions and negative coupling between contralateral regions. I also found that while pairwise connectomes often selected the same ROI pairs for strong positive connections, they often “wash out” the negative couplings to be faint and near zero. I believe that this indicates that multivariate functional connectomes may uncover more spa-

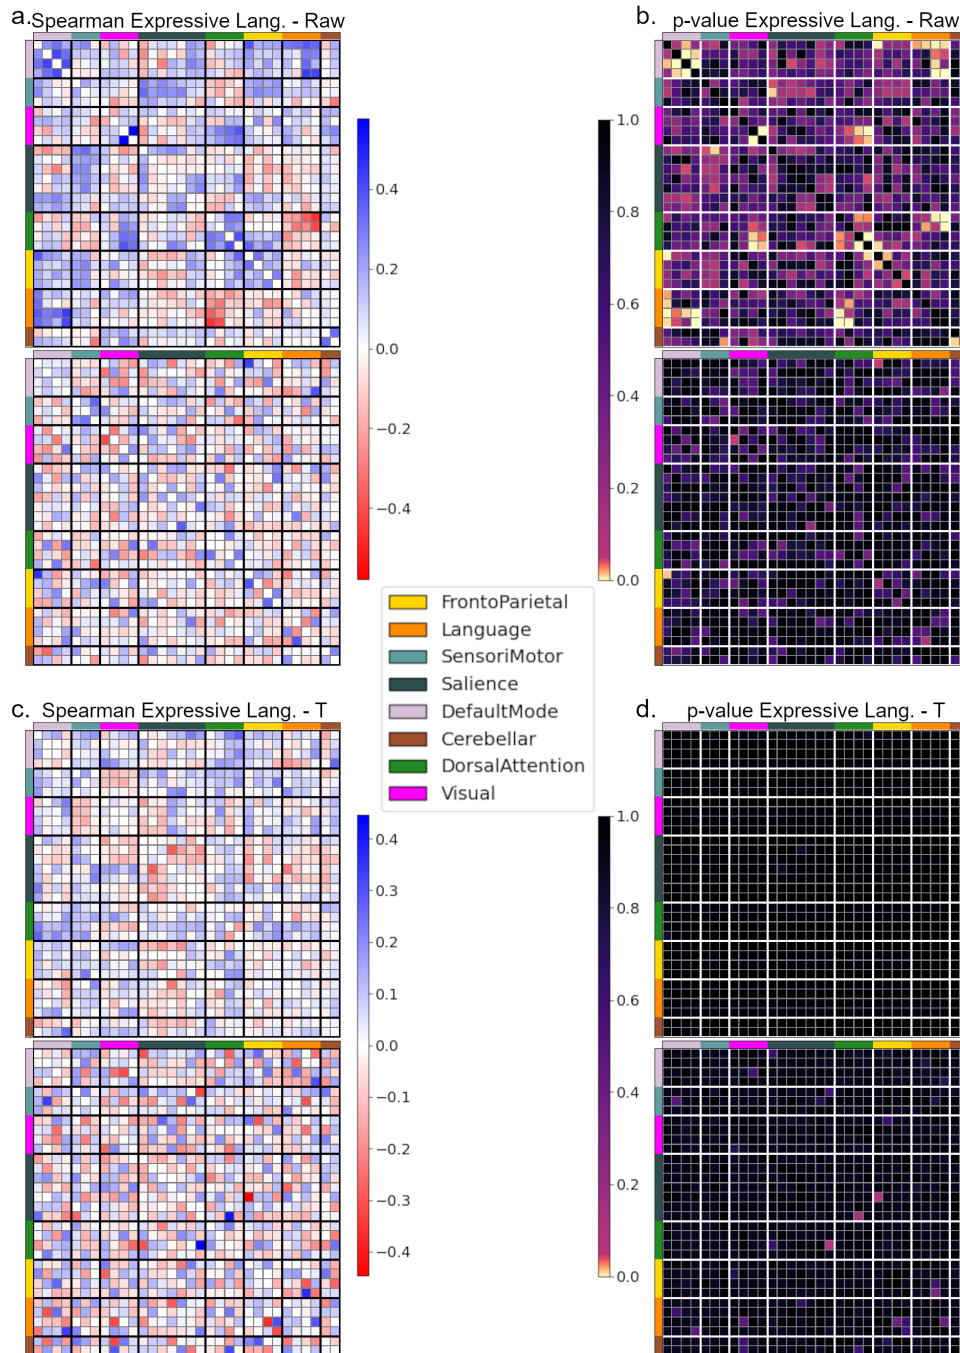


Figure 3.17: **Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations.** **a.** Spearman correlation coefficients and **b.** p-values between MSEL Expressive Language raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. **c.** Spearman correlation coefficients and **d.** p-values between MSEL Expressive Language t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.

tially resolved neural structure than pairwise, due to the “explaining away” of shared variance. On the other hand, pairwise fMRI connectomes appear to correlate much more strongly with Mullen Scales of Early Learning, specifically the raw (i.e., not age-corrected) scores. This difference between pairwise and multivariate connectomes will be further investigated in the next chapter.

Chapter 4

Prediction of cognitive development from multivariate functional connectomes

4.1 Introduction

The study of early childhood developmental trajectories is critical to the improvement of educational outcomes for children with developmental impediments, innate or external. Socioeconomic status is a significant predictor of early childhood success in education[104, 105, 106]. By the time children enter the school environment, gaps have already formed in school readiness, largely along the socioeconomic lines. This is likely due to the disparities in access to resources—for example, nutrition, developmentally appropriate toys, and educational childcare programs. All of these measures are most effective when applied early and consistently[104]. Thus, prediction of developmental trajectories in infancy and early childhood is crucial for effective and timely intervention when a child is at risk of falling behind.

When considered on an international scale, the problem of assessing developmental trajectories becomes thornier. In general, developmental psychologists depend primarily upon behavioral metrics to assess early childhood development[107, 108]. However, the translation of these assessments, linguistically and culturally, is a nontrivial problem and highly vulnerable to cultural bias[108]. In order to equitably and accurately assess child development, a reliable, objective baseline is needed.

One potential approach to objectively assessing early developmental trajectories is to go “straight to the source,” that is, to collect and assess metrics directly from the brain. With modern advances in neuroimaging, the availability of noninvasive, high-

dimensional recording modalities has opened new frontiers in human neuroscience. In order to be viable for widespread deployment, the method of data acquisition must be economical, tolerable for child subjects, and ideally not dependent on a difficult-to-translate task. Resting-state neural recordings, which do not require high levels of patient compliance (of particular interest in young patients), are an ideal target for this kind of analysis.

While correlations have been found between Mullen scores and individual edges in the functional connectome for this dataset[39], the question of whether fMRI-based functional connectomes can predict cognitive developmental measures such as Mullen scores remains open. In this chapter I investigated the predictive power of multivariate functional connectomes created for the BAMBAM fMRI dataset using UoI_{LASSO} . I created multiple neural feature sets by employing dimensionality reduction algorithms such as PCA, as well as hand-selected noteworthy features based on previous fMRI studies. I compared the predictive accuracy of features from multivariate and pairwise connectomes.

With the rising use of predictive algorithms for applications ranging from targeted pages on entertainments sites, to life-changing decisions such as mortgage lending[109] to predicting criminal recidivism[110], it is of the utmost importance to examine use of big data methods with a critical eye. Particularly in applications such as those discussed here, it is desirable that any algorithmically defined diagnostic criteria be human-readable—that is, interpretable and grounded in a broader understanding of the data used to generate predictions. The dataset used in this chapter is as curated as possible for experimental data collected from human children; the subjects are relatively culturally uniform (proficient English-speakers living in the same city), without known genetic conditions, neurodivergence, or gestational complications. The BAMBAM cohort was used as a trial case to study the viability of resting-state fMRI-derived functional connectomes for the prediction of cognitive development.

4.2 Methods

Behavioral and neural data

I used the same BAMBAM cohort discussed in Chapter 3 for prediction of Mullen scores from connectome features. I excluded all subjects for whom either neural or behavioral data was unavailable. The final dataset of usable functional connectomes had $n = 109$ subjects aged between 90 and 903 days (median = 377 days, mean = 397.9 days, standard deviation = 204.4 days). For each of these subjects, I generated

a functional connectome from resting state fMRI data using the UoI_{LASSO} algorithm. For a description of the connectome generation and characterization of the properties of the resulting connectomes, see Chapter 3. The cognitive development of each subject was assessed using the Mullen Scales of Early Learning (MSEL), a standardized method that provides scores for five developmental categories: fine motor, gross motor, visual reception, receptive language, and expressive language. These scores are available in raw form, as well as in age-standardized t-scores (mean 50, standard deviation 7.5). Verbal, nonverbal, and total early learning composite scores are also computed.

The BAMBAM dataset also included subjects for whom neural data was not available at the time of analysis, but with available MSEL assessments. In total (including the 109 subjects described above), the dataset included Mullen scores for 522 subjects between the ages of 64 days and 2041 days (median = 394 days, mean = 560.6 days, standard deviation = 484.7 days). While only the 109 subjects in the neural data set could be included in the prediction dataset, the additional subjects were used to characterize the distribution of Mullen scores within this cohort.

Hierarchical clustering of behavioral groups

I grouped subjects based on their scores on the Mullen scales. I used the raw scores for receptive language, expressive language, gross motor, fine motor, and visual reception assessments. In order to create groupings of maximally similar subjects, I used hierarchical agglomerative clustering (HAC), an unsupervised method of creating . In brief, HAC begins with each data point in a separate cluster, and recursively groups data points by combining one pair of clusters at each step, subject to minimizing a given objective function. I used the Ward objective function, which minimizes in-group variance[111] due to its success in producing groups with cohesive score distributions.

HAC yields a linkage function, which must then be cut at a selected cophenetic distance (i.e., the inter-group dissimilarity at which two clusters are first combined[112]) threshold t_c . I selected the threshold $t_c = 75$ due to the high cophenetic distance between the three resulting groups, as seen in Fig 4.1a. A higher threshold would have resulted in two groups, with groups 2 and 3 combined together. However, as shown in Fig 4.1b, MSEL scores in group 2 are a distinct distribution from group 3, particularly for expressive and receptive language. Conversely, a lower choice of t_c , resulting in four or more groups, would be likely to subdivide group 1, which comprises children who scored low across all five raw scores. For the purposes of predicting developmental trajectories, it is desirable for the purposes of this project to

cast a wide net when flagging for potential developmental difficulties—a false positive is preferable to a false negative.

I investigated the viability of splitting group 1 into two by setting $t_c = 42$, resulting in groups 1a and 1b. I ran a 2-sample t-test comparing the distributions of 1a and 1b for all raw MSEL scores. I found a significant difference in 1a and 1b across all raw scores. However, for the purposes of classification from functional connectivity features. I chose to keep group 1 as a single group.

I quantified the within-group mean and standard deviation of all MSEL scores, including raw scores, t-scores, and age equivalents, as well as verbal, nonverbal, and total early learning composite scores for HAC groups 1, 2, and 3. All statistics are reported in Table 4.1.

I repeated the HAC analysis using all available subjects, including subjects for whom neural data was not available. These additional subjects included older children aged up to 2041 days (or 3.6 years). See Table 4.2 and Figure 4.3 for details of each grouping for the full dataset.

I created an alternate set of groups by clustering all $n=109$ subjects based on the five MSEL t-scores (fine motor, gross motor, receptive language, expressive language, and visual reception) instead of the corresponding raw scores. Based on the cophenetic distance between branches of the dendrogram (Figure 4.4a), I created two groups. These groups were used as alternative target labels for classification.

Connectome feature selection and dimensionality reduction

The brain data consists of 109 multivariate functional connectomes derived from fMRI data with 32 ROIs across 7 different functional networks. Thus, each connectome has 32×32 edges, minus self-connections, resulting in a total of 992 usable dimensions. Considering the low- n dataset, it is necessary to reduce the dimensionality of the data. I generated several sets of features and tested their predictivity for MSEL groups. Each feature set described in this section was generated for both multivariate and pairwise connectomes.

First I used Principal Components Analysis (PCA), a standard method of dimensionality reduction technique in which components, PCs, are the dimensions along which the variance of the data is maximized. PCs are the eigenvectors of the covariance matrix of the data. Thus, PCA supplies a computationally accessible method for reducing the dimensionality of the data.

While PCA is optimal for capturing variance in the data, the resulting components are dense (most elements are nonzero) in the original feature space. This lack of selectivity renders the PCs relatively opaque for purposes of interpretation. In order to extract more parsimonious features, I used sparse PCA (sPCA). Sparse PCA

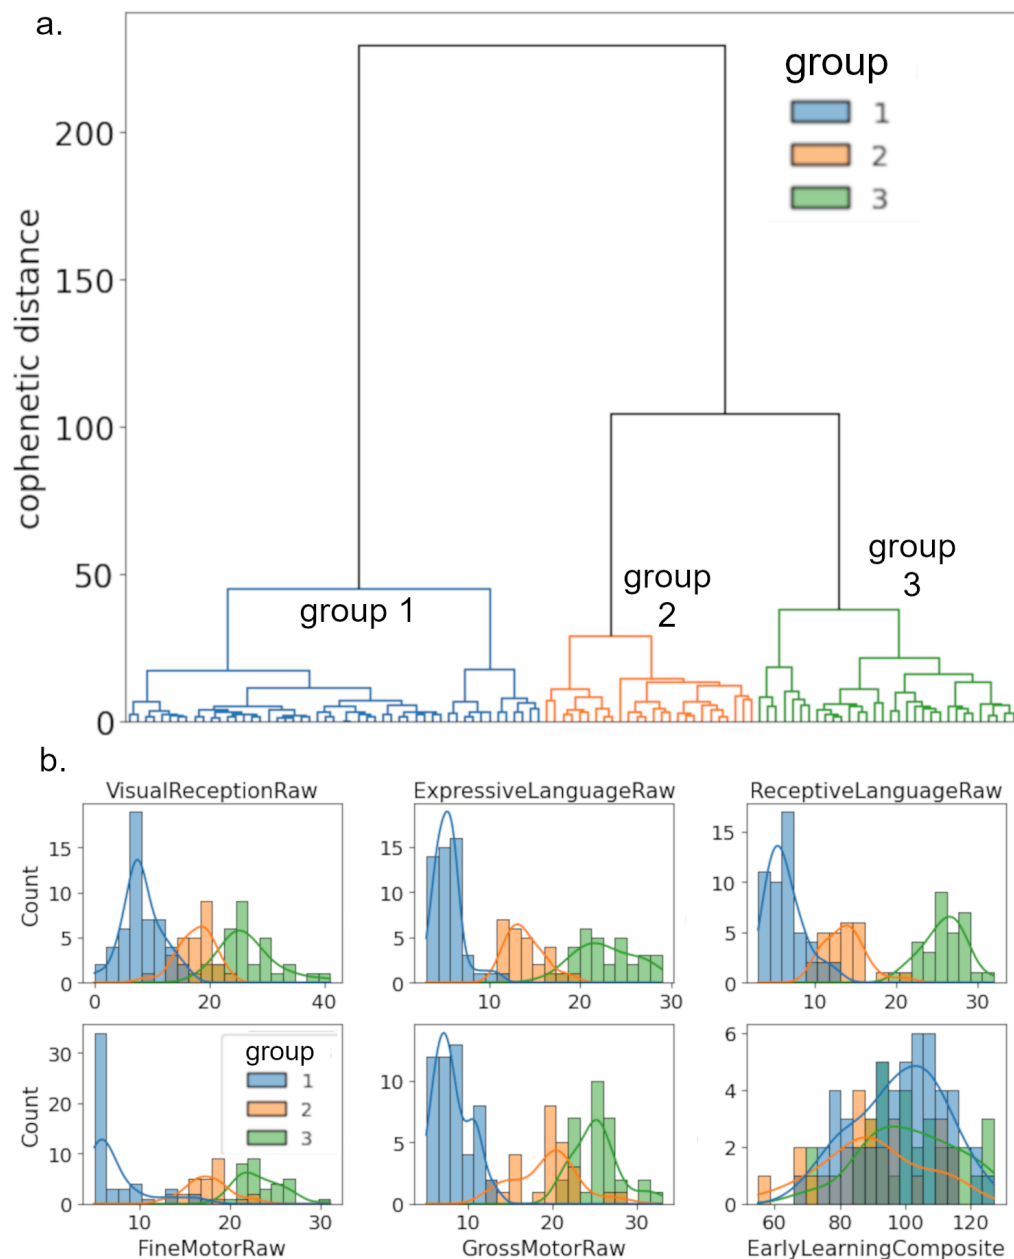


Figure 4.1: Raw Mullen scores for patients with available fMRI data can be subdivided into three groups. **a.** Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 75$. **b.** Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.

		1	2	3
Fine Motor	T-Score	48.6 (9.2)	48.2 (11.4)	44.9 (9.2)
	Age Equiv.	5.3 (3.1)	15.6 (2.6)	22.5 (3.2)
	Raw	7.3 (3.1)	17.5 (2.4)	23.3 (2.6)
Gross Motor	T-Score	49.3 (9.3)	49.7 (9.6)	47.9 (10.6)
	Age Equiv.	5.2 (2.0)	15.7 (3.7)	22.6 (4.1)
	Raw	8.1 (2.1)	19.3 (3.7)	25.3 (2.8)
Expressive Language	T-Score	46.0 (9.2)	45.0 (9.8)	51.5 (9.8)
	Age Equiv.	4.4 (1.6)	13.8 (2.3)	24.3 (4.0)
	Raw	5.4 (1.6)	13.9 (2.0)	22.8 (3.3)
Receptive Language	T-Score	47.1 (10.8)	37.7 (9.1)	55.9 (8.6)
	Age Equiv.	4.4 (2.3)	12.2 (3.0)	26.8 (3.7)
	Raw	6.4 (2.3)	13.7 (2.5)	25.9 (2.7)
Visual Reception	T-Score	54.1 (12.1)	46.6 (11.7)	50.8 (10.6)
	Age Equiv.	6.0 (2.8)	14.6 (3.2)	24.7 (6.3)
	Raw	8.3 (3.4)	17.6 (3.2)	26.6 (4.9)
NonVerbal	Composite	116.5 (27.1)	102.9 (18.5)	100.9 (13.7)
Verbal	Composite	94.6 (24.3)	89.2 (17.3)	110.1 (17.3)
Early Learning	Composite	98.2 (13.6)	89.7 (15.5)	101.8 (14.6)

Age	In Days	421.0 (164.5)	356.5 (191.5)	394.8 (264.6)
Group Size	n Subjects	51	26	32

Table 4.1: Mullen scores, group size, and ages for clustering-derived groups for the subset of subjects for whom neural data was provided. Means and standard deviations (in parentheses) are reported.

		1	2	3	4
Fine Motor	T	46.9 (8.1)	51.3 (10.7)	44.9 (10.4)	52.5 (12.0)
	Age eq.	4.3 (1.6)	15.1 (2.8)	26.0 (5.3)	51.8 (9.9)
	Raw	6.3 (1.6)	17.1 (2.6)	26.0 (4.0)	42.0 (4.9)
Gross Motor	T	48.7 (8.3)	47.9 (11.6)	44.7 (15.1)	20.0 (N/A)
	Age eq.	4.5 (1.7)	14.3 (3.5)	26.6 (5.3)	33.0 (N/A)
	Raw	7.3 (1.9)	17.8 (3.8)	28.7 (4.6)	36.0 (0.2)
Expressive Language	T	48.3 (8.8)	43.5 (11.3)	49.5 (11.1)	52.4 (9.2)
	Age eq.	4.3 (1.2)	12.5 (4.0)	27.4 (7.3)	52.6 (8.4)
	Raw	5.3 (1.2)	12.7 (3.5)	25.2 (5.8)	42.0 (4.0)
Receptive Language	T	47.5 (10.0)	41.9 (11.3)	53.9 (10.6)	57.0 (10.6)
	Age eq.	3.8 (1.7)	12.4 (4.0)	30.1 (6.2)	56.4 (9.7)
	Raw	5.8 (1.7)	13.9 (3.6)	28.1 (4.2)	42.9 (4.4)
Visual Reception	T	51.4 (11.2)	47.7 (11.5)	54.7 (13.2)	59.0 (12.3)
	Age eq.	4.9 (1.9)	13.9 (3.2)	30.8 (9.3)	57.8 (8.4)
	Raw	7.1 (2.4)	16.9 (3.2)	31.1 (6.7)	46.6 (3.0)
Nonverb.	Comp.	109.8 (25.8)	106.7 (18.0)	102.8 (16.6)	109.2 (14.1)
Verb.	Comp.	98.1 (25.2)	90.2 (20.7)	104.8 (19.0)	108.6 (15.4)
Early learning	Comp.	97.3 (12.8)	92.8 (16.0)	101.8 (18.0)	110.5 (17.0)

Age	Days	392 (231)	362 (256)	511 (399)	1561 (278)
Group	n	152	164	141	65

Table 4.2: Mullen scores, group sizes, and ages for clustering-derived groups for larger dataset with expanded age range. Means and standard deviations (in parentheses) are reported.

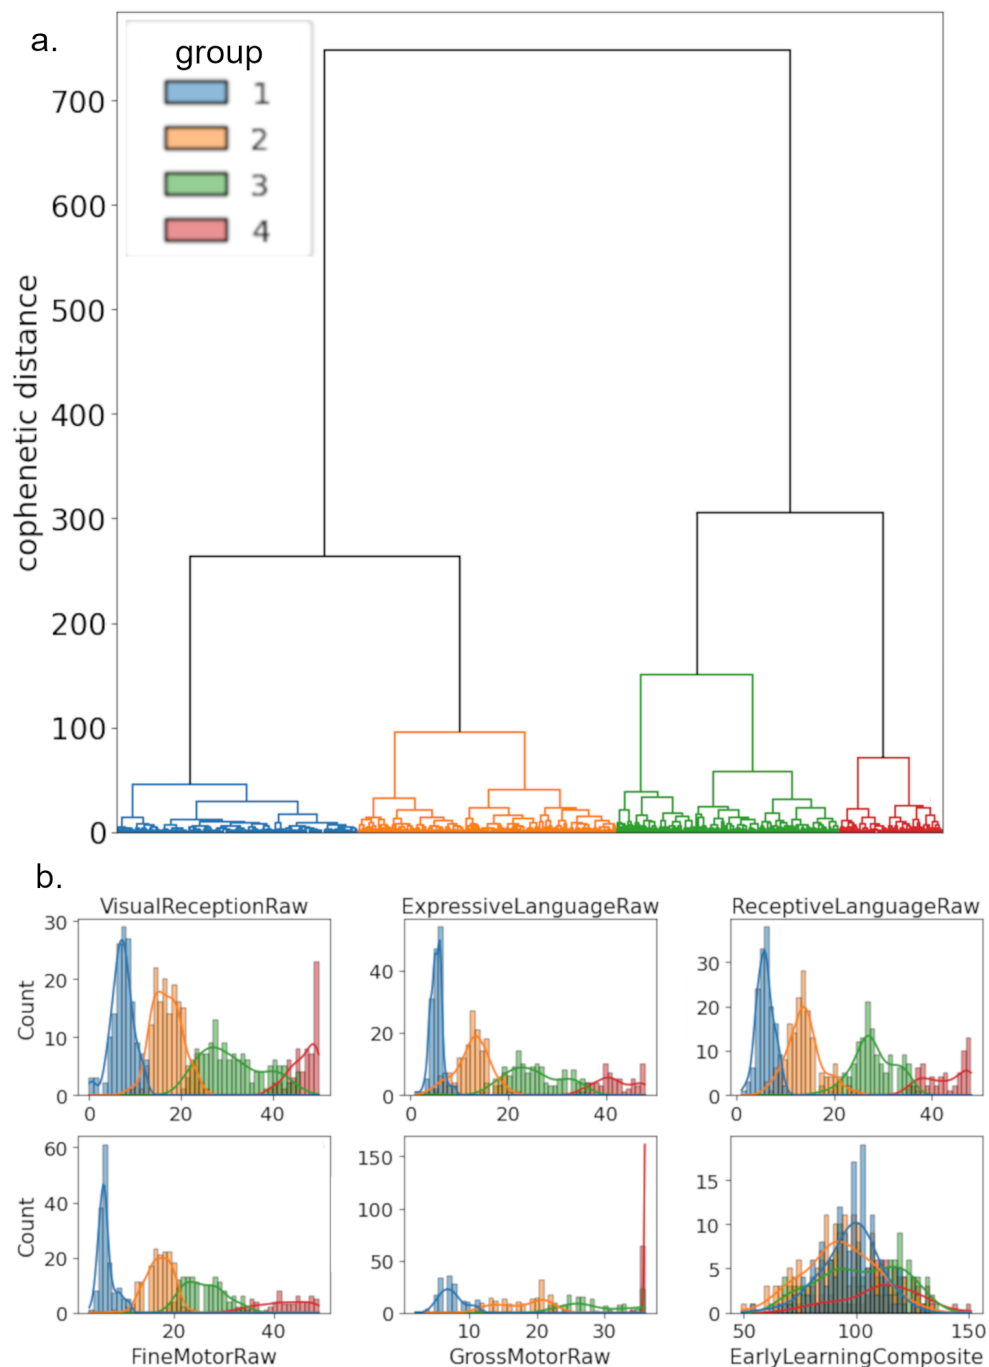


Figure 4.2: **Grouping of all available MSEL scores, including those for which neural data was not available, replicates the three groups from Figure 4.1, but adds a fourth group with older subjects** **a.** Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 200$. **b.** Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.

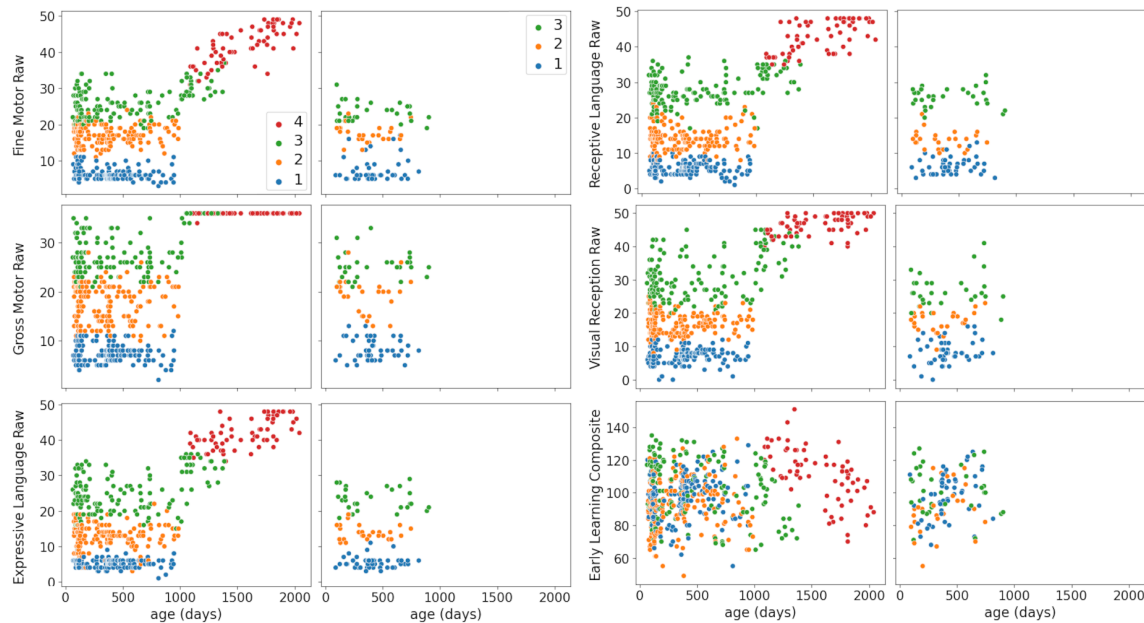


Figure 4.3: **Subjects between the ages of 90 and 1000 days fall into 3 consistent bands that are preserved across expressive language, receptive language, fine motor, gross motor, and visual receptive scores.** Between the ages of 90 and 1000 days, raw scores do not vary significantly with age. When older subjects between 1000 and 2000 days are added into the dataset, a fourth group, mostly comprising older children, shows strong correlation with age in raw scores. For group 4, the early learning composite score provides a metric that is decorrelated from age. Groups 1-4 are not significantly separated for the early learning composite score.

produces component vectors that are linear combinations of a small number of input features. Sparse PCA imposes is based on L_1 regularization placing a penalty on the absolute values of the coefficients of the component, as in LASSO, resulting in more selective component vectors. One drawback of sPCA is that the components it produces are not strictly orthogonal. The full set of sparse PCs therefore do not perfectly capture the full variance of the dataset.

Upon examination of the coefficients of the sparse PCs, I observed that many PCs captured contributions from the visual network. Examination of the standard deviation of networks within 90 day age windows further shows that variance in contributions from the visual lateral ROIs tends to be high. For dimensionality re-

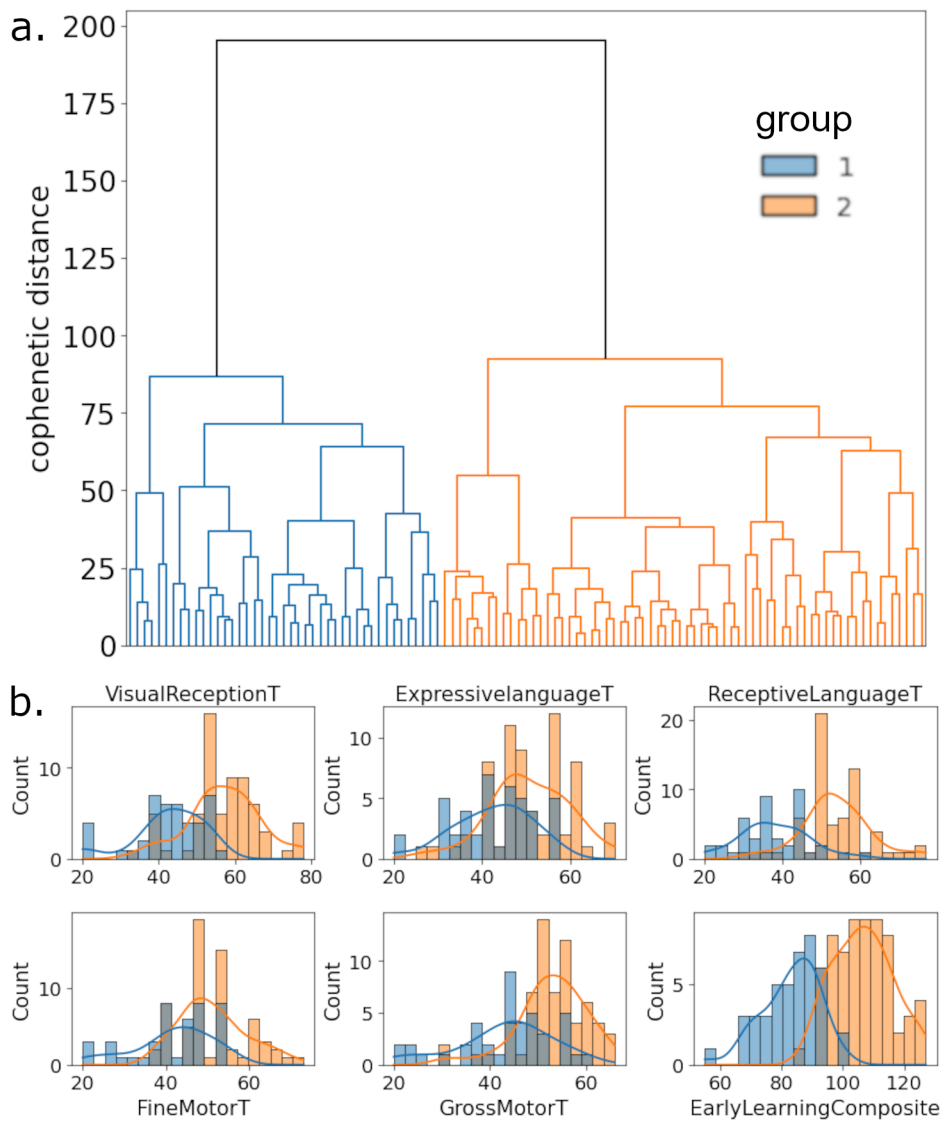


Figure 4.4: **T-score groups separate Early Learning Composite into two distinct distributions.** **a.** Dendrogram of hierarchical clustering based on the five raw MSEL scores. Final cophenetic distance threshold for cluster assignment was set at $t = 75$. **b.** Distributions of raw MSEL scores and early learning composite score for each HAC-assigned group.

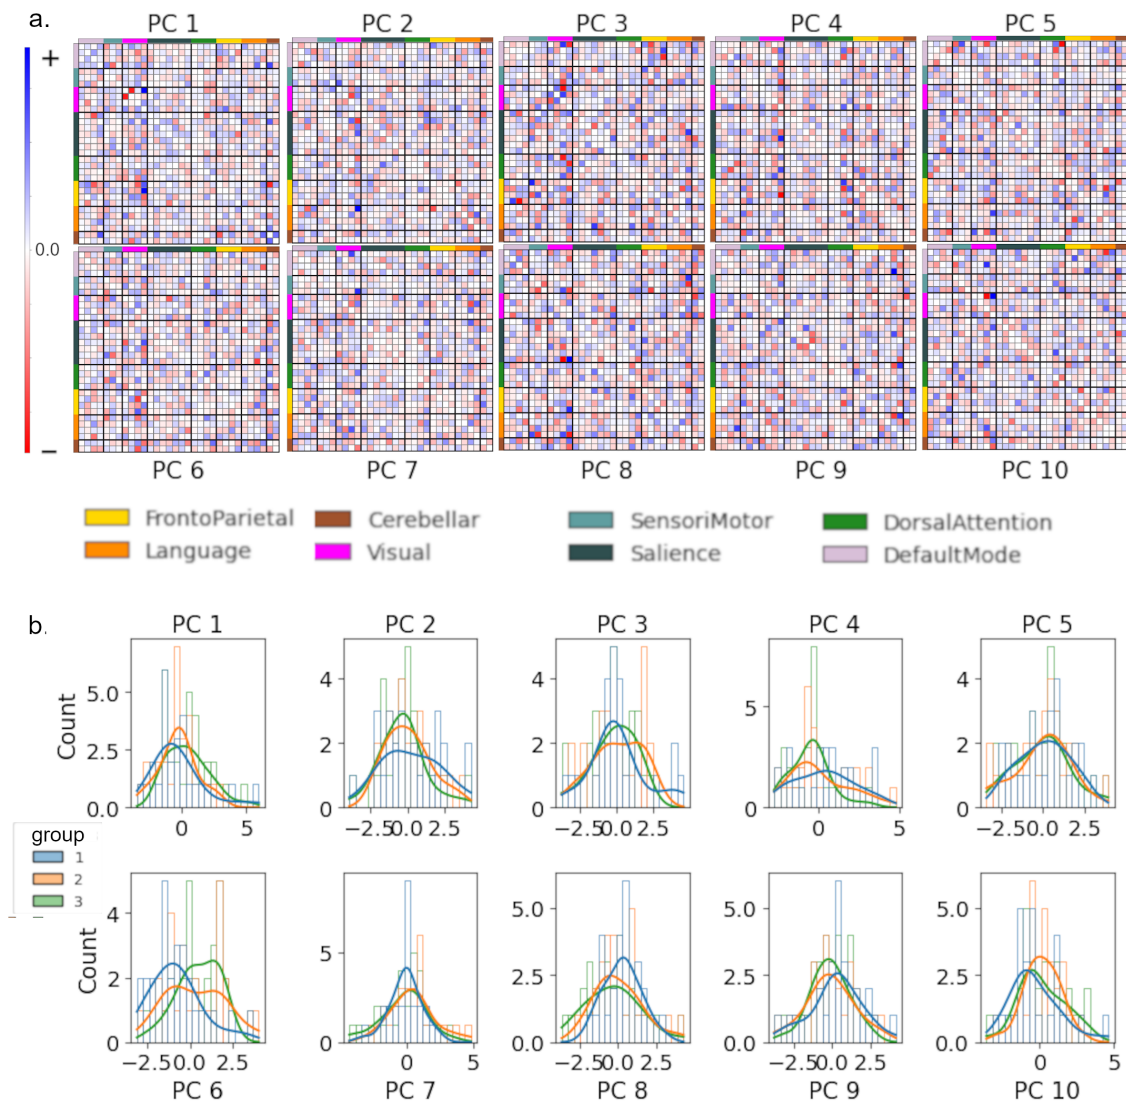


Figure 4.5: **PCs are dense and difficult to interpret, but some lateralized structure is discernible.** **a.** Coefficients of first 10 principal components plotted in the original feature space. **b.** Distribution of PC values within Mullen groups.

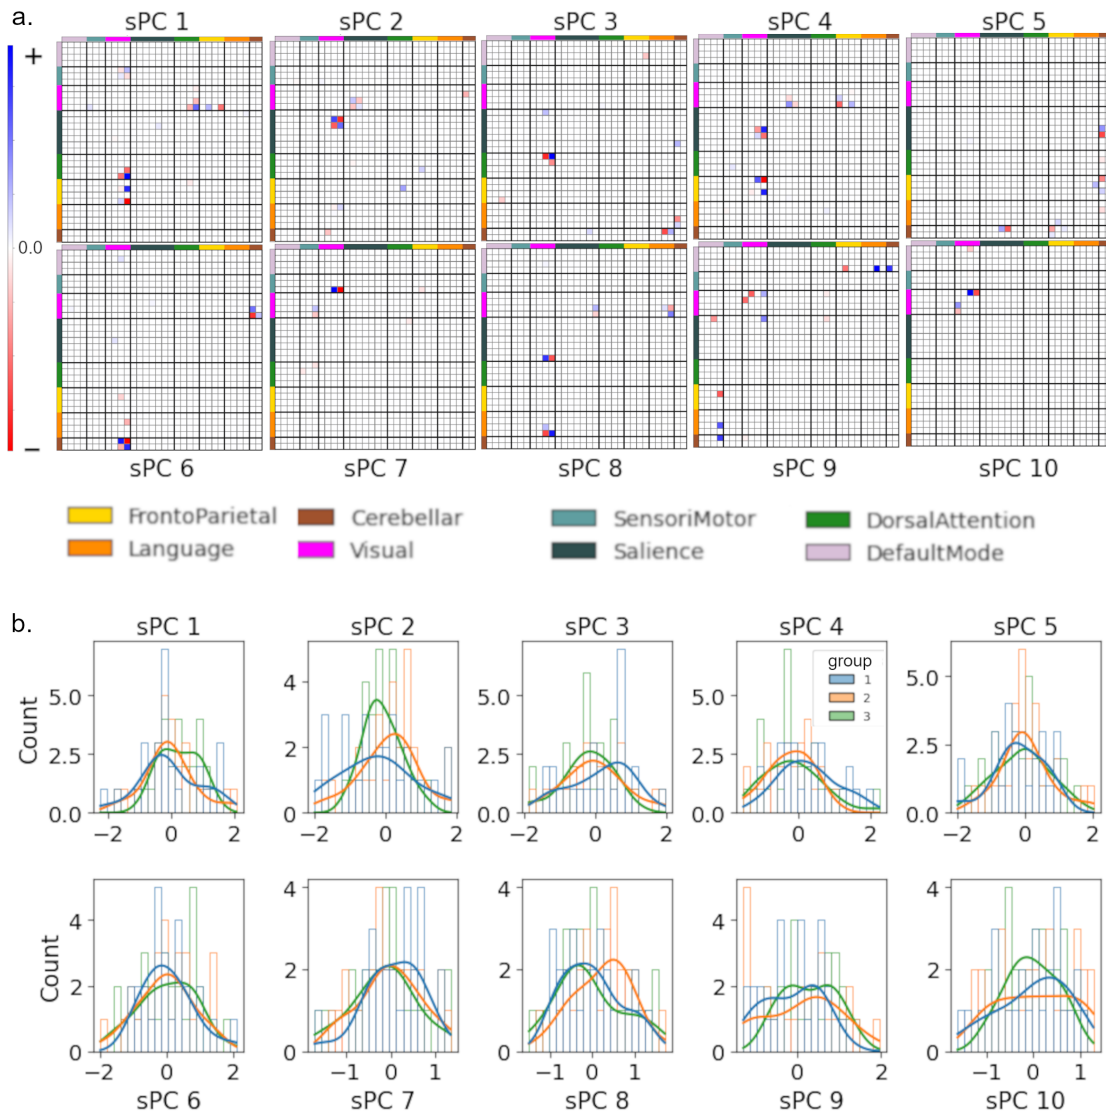


Figure 4.6: **Sparse PCs show selection for lateralized connections and are highly concentrated in the visual functional network.** **a.** Sparse coefficients of first 10 sparse principal components plotted in the original feature space. **b.** Distribution of PC values within Mullen groups.

duction techniques based on PCA, which selects components that capture maximum variance for the dataset, it is natural for the visual network to be heavily represented.

As described in the previous chapter, the fMRI data for the BAMBAM cohort is collected from sleeping subjects. Functional connectivity to the visual network in young, sleeping subjects is not well understood. It has been shown in one study [113] that functional connectivity of the default mode network to the visual network can increase for sleeping adults, but the same pathway was not observed in children. Examination of the first 10 sparse PCs and their value distributions for the Mullen clusters (Figure 4.6b) showed little separation. I hypothesized that it was possible that the highly variable functional connectivity of the visual network may introduce non-meaningful noise, and that excluding the visual system from sPCA may eliminate this noise source. I thus removed all visual network ROIs (Medial, Occipital, Lateral (left), and Lateral (right)) from the functional connectomes and regenerated sparse PCAs using the reduced networks. Figure 4.7a shows the coefficients of first 10 visual-excluded sPCs plotted in functional connectivity feature space. Figure 4.7b shows the sparse PC value distributions for the Mullen clusters. I saw little to no improvement for separability of Mullen groups between sparse PCs including or excluding the visual network. However, sPCs with visual ROIs excluded did seem to select for network pairs of particular interest in measures of general intelligence and language development, such as frontoparietal-default mode (sPC 1), frontoparietal-salience (sPC 2), language-salience (sPC 5), and frontoparietal-language (sPC 6).

I created three sets of PCs, taking the top 10 components for each: regular PCA, sparse PCA with all functional networks, and sparse PCA with the visual network excluded.

In order to more directly assess the predictivity of particular feature pairs, I also hand-selected feature sets directly from the functional connectomes. In this regime, each “edge” (e.g., the coupling between left anterior insula and left superior temporal gyrus) is a potential feature. For this analysis, I subselected pairs of developmentally relevant functional networks and created feature sets comprising all edges between ROIs in those networks. For example, one feature set might contain all edges coupling the frontoparietal network to the default mode network, resulting in the 16 edges between frontoparietal ROIs (LPFC left, LPFC right, PPC left, and PPC right) and default mode ROIs (MPFC, LP left, LP right, and PCC). I created feature sets for network pairs that showed statistically significant correlations to Mullen raw scores (see Chapter 3), as well as pairings between the “big three” functional networks: frontoparietal (FPN), default mode (DMN), and salience (SN).

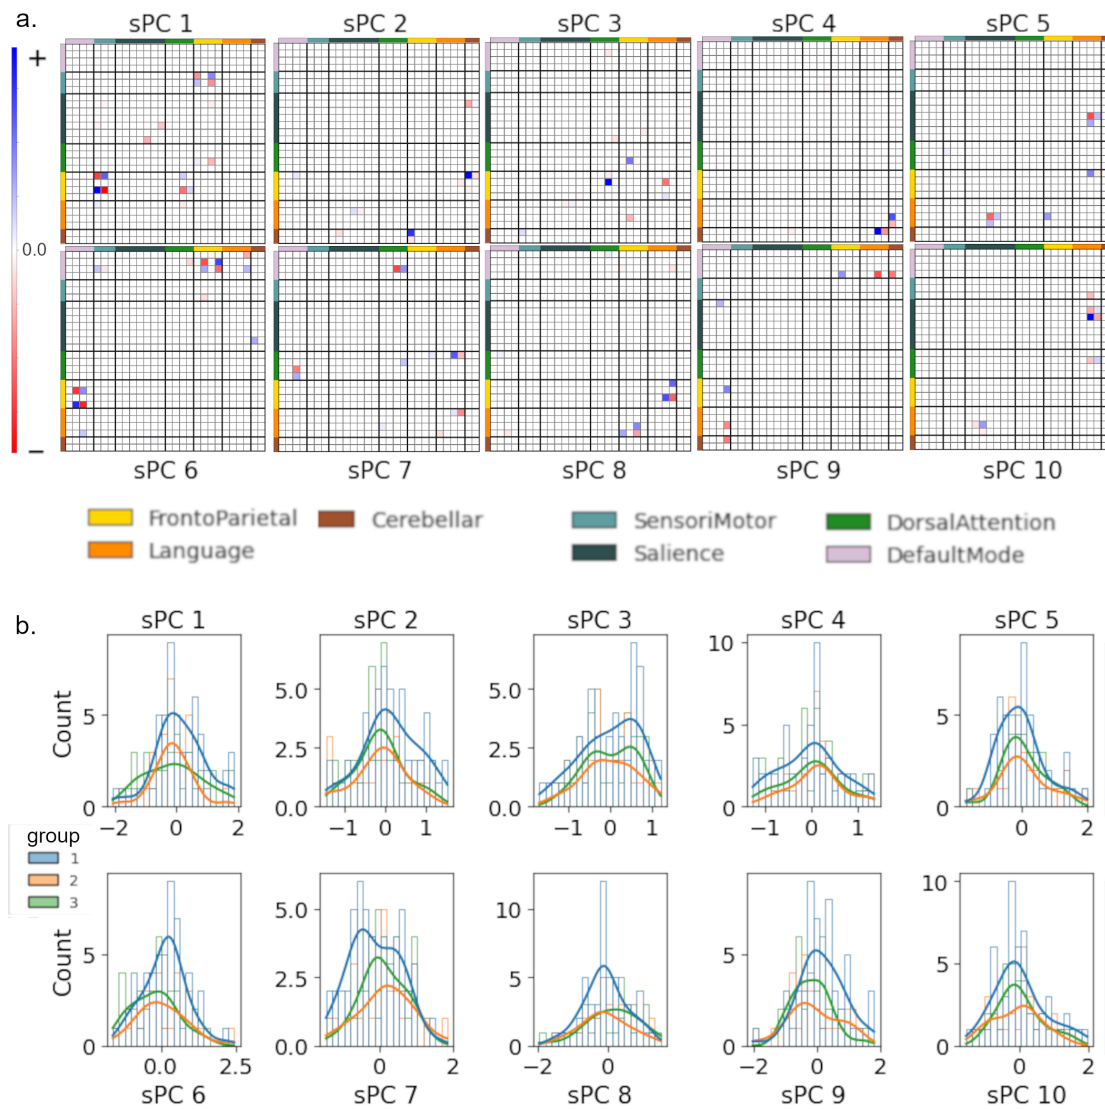


Figure 4.7: **Excluding visual network from sparse PCA does not improve separability of Mullen groups.** **a.** Sparse coefficients of first 10 sparse principal components plotted in the original feature space. **b.** Distribution of PC values within Mullen groups.

Random forest classification

Since the use case of my research is to identify children in need of timely intervention to reach school readiness, it is sufficient for these purposes to classify subjects into broad achievement bands. I thus began my prediction efforts using a Random Forest Classifier (RFC) to predict Mullen labels. Random Forests are an ensemble machine learning method that can be used for both regression and classification. RFCs operate by creating a settable number of decision tree estimators and aggregating over these estimators to provide a final result. The random forest creates decision trees over many bootstrap aggregates, “bags”, of the data. Individual trees also select a subset of features to use in their predictions. The final prediction is provided by aggregating (or “voting” in the case of classification tasks) on the target variables of the datapoints left out of that bag[114].

The input observation vectors \vec{x} are brain features, as described in the previous section. The target variable y is the Mullen clustering-derived group label, so that each subject is assigned a label between 1 and 3, with group 1 comprising relatively low scoring subjects while group 3 comprises high relatively high-scoring students. Group 2 comprises students with intermediate scores, but the distribution of motor scores in group 2 seems to fall closer to those of group 3, while the distribution of receptive language scores falls closer to group 1. I am thus most interested in correctly identifying members of group 1, which comprises the subjects most likely to need help in achieving school readiness.

For each observation vector \vec{x} , I first optimized the RFC hyperparameters by performing a grid search. These hyperparameters are: maximum depth, maximum features, minimum samples for a leaf, minimum samples for a split, and number of estimators. The grid search selected the best hyperparameters to optimize the prediction score for the given observation and target variables. In order to ensure equal group sizes, I randomly subselected data from each group, so that each group included only as many data points as the smallest group ($n = 26$). Thus, the usable dataset had size $n = 78$.

For each RFC, I generated predicted labels y_p and a range of scores using tenfold cross-validation. Confusion matrices for each \vec{x} are shown in Figures 4.8, 4.9, and 4.10. Mean and standard deviation of cross-validated scores are reported in Table 4.3.

Features	connectome	mean	standard error
PCA	multivariate	0.37	0.062
PCA*	pairwise	0.51	0.045
Sparse PCA*	multivariate	0.46	0.029
Sparse PCA*	pairwise	0.50	0.039
Sparse PCA without vision	multivariate	0.38	0.051
Sparse PCA without vision*	pairwise	0.56	0.053
Language-Salienc*	multivariate	0.44	0.039
Language-Salienc	pairwise	0.34	0.057
DMN-Frontoparietal	multivariate	0.48	0.072
DMN-Frontoparietal	pairwise	0.28	0.047
Salienc-Frontoparietal	multivariate	0.40	0.067
Salienc-Frontoparietal*	pairwise	0.45	0.026
DMN-Salienc	multivariate	0.43	0.053
DMN-Salienc	pairwise	0.35	0.042
Sensorimotor-DorsalAttention	multivariate	0.41	0.042
Sensorimotor-DorsalAttention*	pairwise	0.46	0.036
Language-Frontoparietal	multivariate	0.33	0.043
Language-Frontoparietal*	pairwise	0.41	0.023
DorsalAttention-Visual	multivariate	0.30	0.059
DorsalAttention-Visual	pairwise	0.45	0.063

Table 4.3: **Raw Mullen score groupings are predicted by pairwise connectome features better than by multivariate connectomes.** Tenfold cross-validated classification scores of grid-search optimized random forest classifiers for feature set used for prediction of Mullen raw score groupings. Feature sets that performed significantly better than chance (33.3%) are marked with an asterisk.

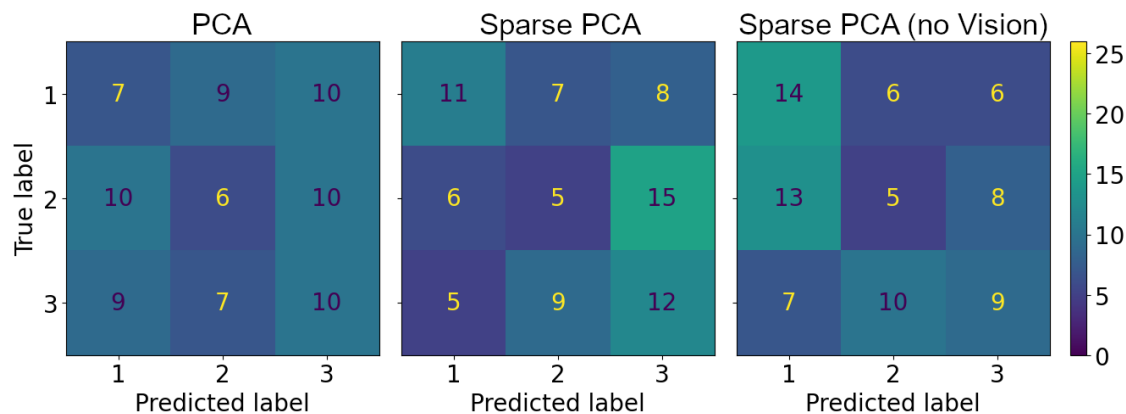


Figure 4.8: **Functional connectome principal components do not consistently predict MSEL scores.** 10-fold cross-validated confusion matrices for optimized random forest classifiers for 10-component PCA, sparse PCA, and sparse PCA with vision excluded.

Features	connectome	labels	mean	standard error
PCA	multivariate	T scores	0.50	0.051
PCA	pairwise	T scores	0.53	0.055
Sparse PCA*	multivariate	T scores	0.62	0.047
Sparse PCA	pairwise	T scores	0.45	0.053
Sparse PCA without vision	multivariate	T scores	0.58	0.046
Sparse PCA without vision	pairwise	T scores	0.47	0.052

Table 4.4: **T-score Mullen groupings are predicted by multivariate connectomes better than by pairwise networks.** Tenfold cross-validated classification scores of grid-search optimized random forest classifiers for feature set used for prediction of Mullen raw score groupings. Feature sets that performed significantly better than chance (50%) are marked with an asterisk.

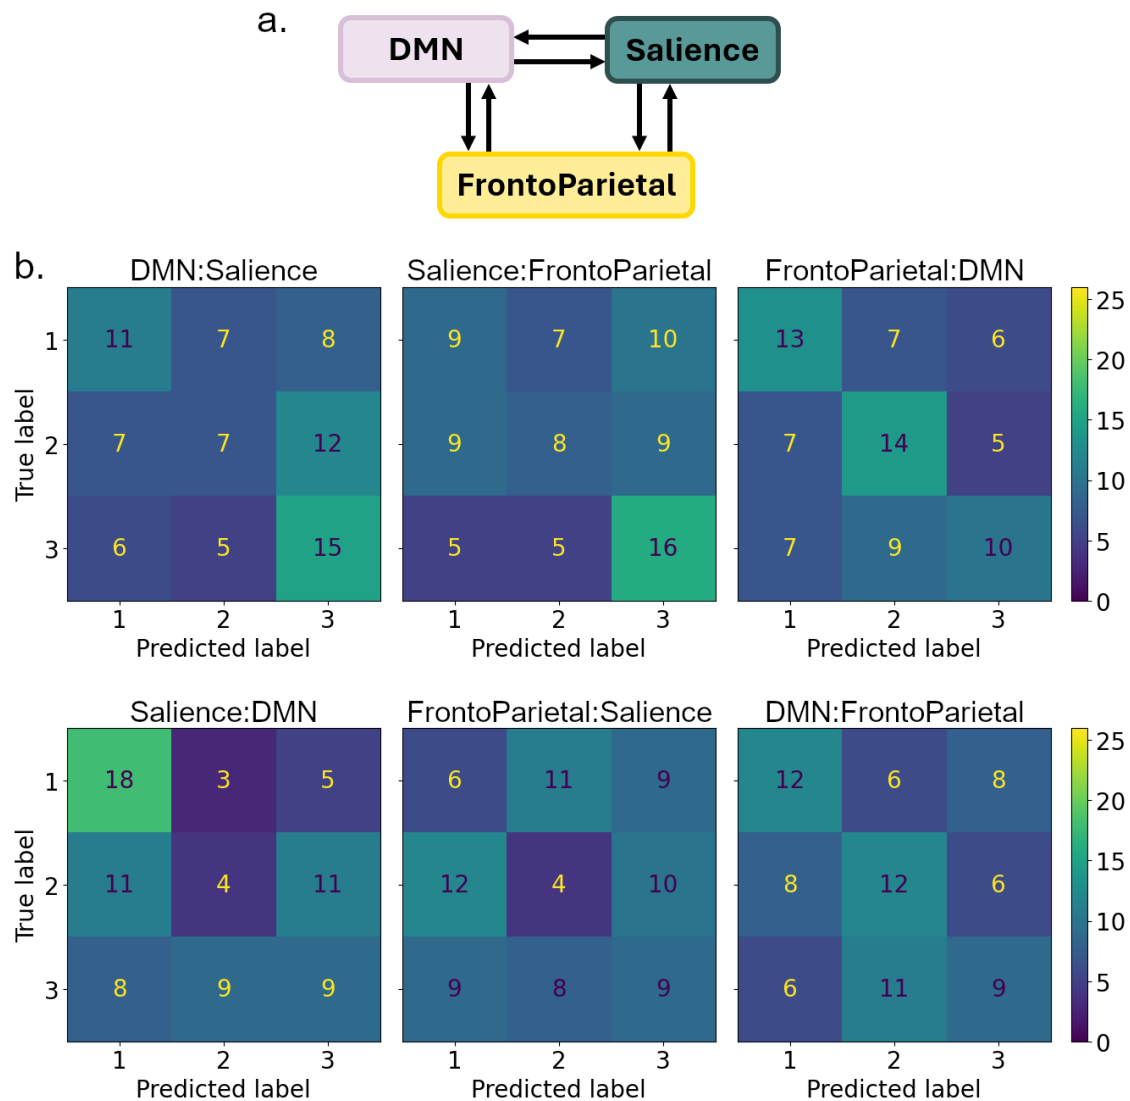


Figure 4.9: **Functional connectivity between default-mode, salience, and frontoparietal networks slightly improves classifier performance.** **a.** Schematic of the interconnected “big three” functional networks: default mode (DMN), frontoparietal (FPN) and salience (SN). **b.** 10-fold cross-validated confusion matrices for optimized random forest classifiers for each pair of networks.

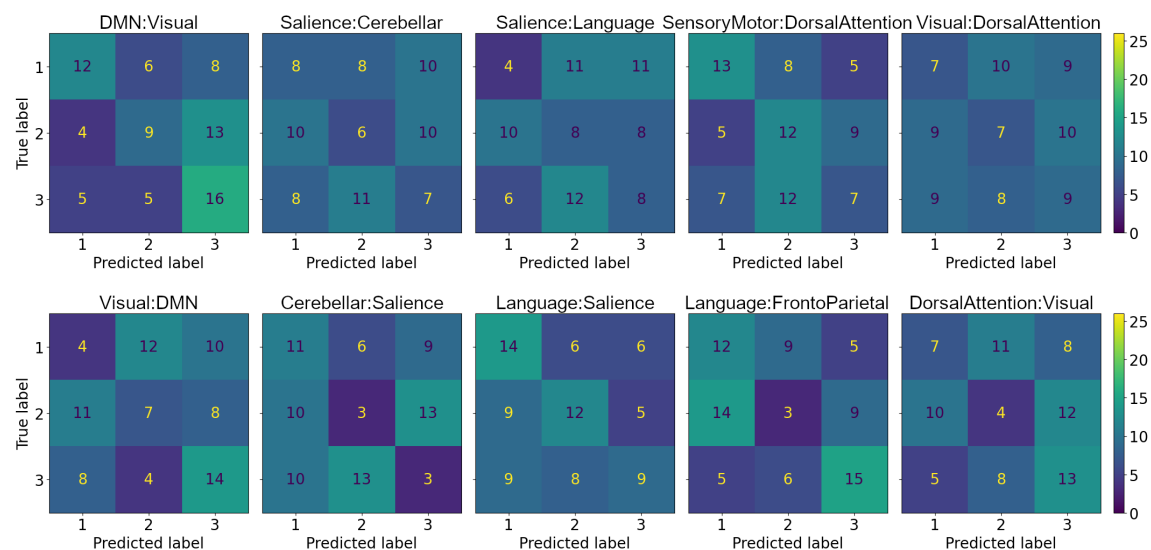


Figure 4.10: **Functional connectivity between other selected network pairs are little better than chance at predicting MSEL groupings.** 10-fold cross-validated confusion matrices for optimized random forest classifiers for other selected network pairs. Language-saliency, default mode-visual and visual-default mode, and sensorimotor-dorsal attention perform slightly better than chance (33%).

4.3 Results and Discussion

Raw Mullen scores show distinct performance groups

Unsupervised clustering of the MSEL scores of the BAMBAM cohort showed marked sub-populations distinguished by their raw scores in all five assessment areas. For children between the ages of 3 months and 2.5 years, Mullen raw scores were not significantly correlated with age. Furthermore, for this age range, subjects tended to achieve similar scores across cognitive areas. Figure 4.1b shows the distributions of Mullen scores color coded by HAC-assigned group. It is notable that the early learning composite, which combines all five assessments and normalizes by age group, is not separated by raw-score HAC groups. By contrast, t-score HAC groups partition the early learning composite into two distinct distributions (Figure 4.4b).

I applied the same unsupervised clustering procedure to a larger dataset with an expanded age range (up to 5 years of age). Per Figure 4.2, groups 1-3 remained in similar scoring bands for the raw scores. However, an additional cluster, group 4,




Figure 4.11: **Mullen raw scores are better predicted by pairwise connectome principal components than by multivariate connectome principal components.** Confusion matrices for 10-fold cross-validated RFCs, with feature sets comprising 20 components each.

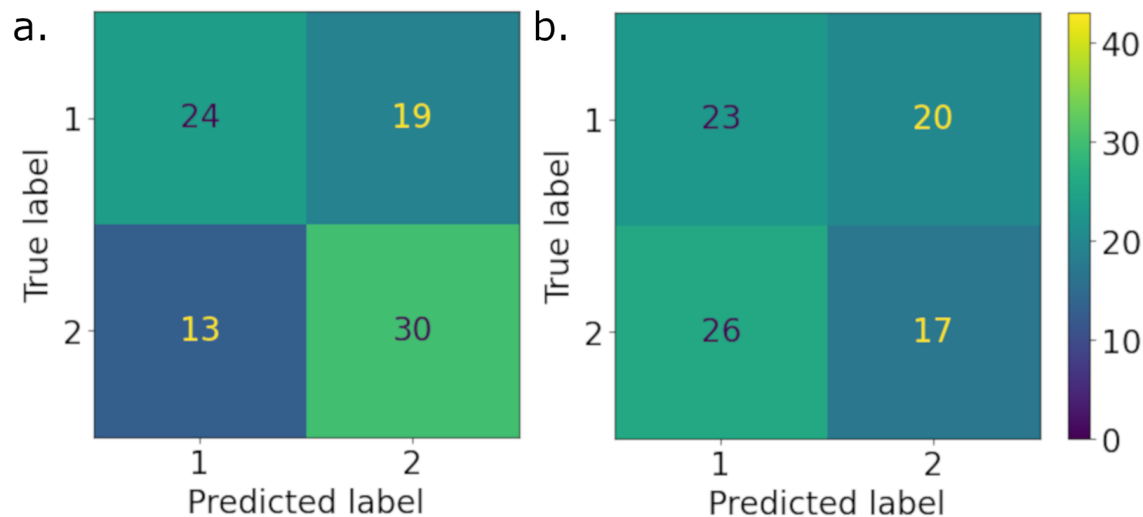


Figure 4.12: **T-scores are better predicted by sparse PCs of multivariate connectomes than pairwise connectomes.** Confusion matrices for 10-fold cross-validated RFCs predicting t-score Mullen groupings from 20-dimensional sparse PCs. **a.** Multivariate connectomes **b.** Pairwise connectomes

with higher scores, was added. Figure 4.3 clearly demonstrates that the addition of children over the age of 2.5 years is the cause of this new grouping. Above 2.5 years, raw scores begin to correlate clearly with age (with the exception of gross motor, which appears to be maximized for almost all children above the cutoff of 2.5 years). Figure 4.3 also demonstrates that the early learning composite successfully decorrelates scores from age for children above the 2.5 year cutoff. This indicates that if I were to apply this prediction analysis to an older cohort, the T-scores and composite scores may be a better target.

Lack of age correlation of raw scores under the age of 2.5 may cast doubt on their validity—does the variability in raw scores really signify meaningful cognitive ability differences of 3 month olds have the same score distribution as 2 year olds? In addition to the correlations in Chapter 3 Figure 3.17, previous work on this dataset has shown that there exist significant correlations between pairwise functional connectivity and Mullen scores[39], suggesting at least some relevant signal. Further analysis, predicting Mullen scores individually within tighter age ranges, would help

interpret the MSEL and its relationship to functional connectivity.

Sparse principal components replicate lateralized structure seen in aggregate age group template connectomes

In all three PC-based datasets (Figures 4.5, 4.6, 4.7), individual PCs selected for grouped ROI-ROI couplings that displayed lateralized structure—most commonly, negative contralateral weights and positive ipsilateral weights. This indicates that the variability of multivariate functional connectomes tends to be tightly clustered in related ROI-ROI edges.

As mentioned in the Methods section, it is notable that coupling from the visual network is heavily selected for in several of the PCs. This visual network dominance may be explained by two different reasons. First, the visual system is established and functionally synchronized from birth, while other networks such as default mode and frontoparietal are established over months. Thus, the visual network may be unusually dominant in infant and early childhood connectomes. If this is the root cause of the dominant visual network, then exclusion of the visual network would remove meaningful signal from my analysis. Second, the rsfMRI data used in this analysis is collected from sleeping subjects, rather than passively alert subjects. The neurodevelopmental significance of visual network activation in sleeping children is not well understood, and thus the contributions of the visual network in this dataset may amount to little more than noise. The next section evaluates the contributions of the visual network by their predictive power.

Mullen-grouped distributions show little separation in functional connectivity PC space

In order to parse out the importance of the visual network, I first removed all visual ROIs from the connectome dataset and repeated the sparse PCA, generating a second sPC feature set from the remaining 28 ROIs. Comparison of Figure 4.6 to Figure 4.7 shows that while the removal of the visual ROIs led to greater diversity in the networks featured in the sparse PCs—for example, language and salience were almost absent from the first 10 sPCs for the original feature set, but were the primary features making up sPC 8 in the visual-free neural feature set—it does not meaningfully improve separability of Mullen groups in sPC space. The distributions of sPCs of the Mullen groups remained largely overlapping.

This lack of separability is carried over into the performance of RFCs trained on these principal components. All three PC sets generated networks that performed

no better than chance (see Fig 4.8 for confusion matrices and Table 4.3 for cross-validated scores). This null result suggests that the largest sources of variability in the connectome dataset are not significant to cognitive measures (at least insofar as I accept Mullen raw scores as an adequate metric of cognition).

The next step in algorithmic dimensionality reduction of the connectome features would be to use a sparse linear classifier, trained to predict Mullen scores from the full functional connectomes. Of course, considering the low sample size of this dataset compared to the dimensionality of the connectomes, this classifier would likely be over-fit, and cannot be used as a prediction method in its own right. Rather, this method would be used as a kind of feature discovery—the features learned by the classifier could then be used as components of a connectome feature vector. This method would allow me to reduce the dimensionality of the connectome data as desired, with features that are preselected for their salience in predicting Mullen scores.

Some multivariate network pairs show tentatively better-than-chance classification accuracy

Of all feature sets used as inputs to a RFC, only 2 pairings of functional networks scored significantly better than chance: Frontoparietal-Default Mode and Language-Salience. However, even for comparatively high-performing networks, with mean cross-validation scores between 40% and 47.5%, perusal of the confusion matrices suggested that these feature sets are often better at identifying one Mullen, while failing to discriminate the other two. For example, the Salience-DMN network pair shown in Figure 4.9 appears to correctly identify the majority of group 1 (that is, relatively low-scoring) subjects. However, further examination shows that the Salience-DMN classifier seems to overzealously classify subjects into group 1 in general, placing many group 2 and group 3 subjects in the same category. Conversely, the DMN-Visual classifier shown in Figure 4.10 seemed generally biased towards group 3 (relatively high-scoring students), pressing half of intermediate group 2 subjects and almost a third of group 1 subjects into the group 3 label. This effect is not due to uneven group sizes, since groups were equalized to each have a sample size of 26 before use in the RFC.

I noticed certain cases in which classification accuracy was not symmetric between network pairs—that is, the couplings $\vec{\beta}_{AB}$ of network A to network B was more predictive of Mullen scores than the coupling $\vec{\beta}_{BA}$ of network B to network A. In particular, Language-Salience is significantly more predictive than Salience-Language ($p < 0.02$, $n = 10$). This observation aligns with the putative function of the salience network,

which is to exert top-down attentional control over other networks for the selection of salient stimuli[61]. It is thus intuitively plausible that contributions from salience to language are more significant than vice versa. Similarly, Salience-Frontoparietal connections are significantly more predictive ($p < 0.03$, $n = 10$) than Frontoparietal-Salience. This asymmetry may reflect the role of the frontoparietal network as a central nexus of executive control[62].

Pairwise connectome features are generally better predictors of Mullen raw scores than their multivariate counterparts

Figure 4.11 shows the marked difference in performance between RFCs trained on multivariate connectome PCs and pairwise connectome PCs. In particular, the sparse PCA excluding vision feature set is significantly more predictive for pairwise than multivariate connectomes ($p < 0.02$, $n = 10$). There are also several feature sets for which only the pairwise feature set is significantly more predictive than chance: PCA, Salience-Frontoparietal, Sensorimotor-DorsalAttention, and Language-Frontoparietal. This result was consistent with the findings in Chapter 3, specifically Figure 3.17, which showed that pairwise connectomes were more significantly correlated with raw Mullen scores. The only feature set for which this pattern was reversed was Language-Salience, for which the multivariate feature set was significantly more predictive than chance ($p < 0.03$, $n = 10$), while the equivalent pairwise set was not. This may be a case in which the asymmetry of the multivariate networks provided an advantage; as discussed in the previous subsection, Language Salience was significantly more predictive than Salience-Language. The pairwise networks were symmetric, and thus did not capture the differences in connectivity.

These results also suggested that I simply generated faulty multivariate connectomes, potentially by choosing sub-optimal hyper-parameters. For example, I set model estimation scoring criterion to R^2 , which is known to be prone to inconsistent feature selection[55]. These connectomes may have been more reliable and stable if instead I had used the Bayesian Information Criterion (BIC), as was done in the cited paper by Sachdeva et. al[55] to improve consistency. I will likely generate new multivariate connectomes with alternate hyper-parameters to assess the effect this hyper-parameter choice.

Multivariate connectome features are generally better predictors of Mullen t-scores than their pairwise counterparts

Table 4.4 reports the classification accuracy trained to predict MSEL t-score groupings from PCA feature sets. The only feature set that predicted t-scores with significantly better than chance accuracy ($p < 0.03$, $n = 10$) was sparse PCA (20 components) for multivariate connectomes. Figure 4.12 shows the confusion matrices for sparse PCs for multivariate (4.12a) and pairwise (4.12b). Further work is required to determine whether this difference in predictive targets between multivariate and pairwise connectomes is consistently replicated in different prediction tasks. In particular, consulting domain experts in early learning assessments may aid in the interpretation of the Mullen scores available for this set of subjects. Closer examination of the relationship between connectomes and Mullen scores, such as Figure 3.17 and Appendix Figures A.3.17-5 is warranted.

4.4 Conclusion

Mullen scores, an instrument for evaluating cognitive development in children under 5 for language, motor, and visual skills, were used to group subjects into 3 achievement groups using unsupervised hierarchical clustering. These groupings were based on the raw (i.e., not age-corrected), then verified against a superset of the data used in this chapter, which included children between 3 months and 5.5 years. It was verified that children younger than 2.5 can be reliably grouped into three groups that appeared uncorrelated with age within this bracket, while the raw scores of older children are positively correlated with age and would thus require use of age-corrected measures. Since the subjects for who I have neural data all fell between 3 months and 2.5 years, I used the raw Mullen scores for prediction.

I reduced the dimensionality of the 32×32 connectomes using PCA and sparse PCA. I used the first 10 and 20 components of each fit to visualize the elements of the connectomes that account for the most variance. I showed that sparse PCA selects for groupings within network pairs, and replicates the lateralized structure that I observed in Chapter 3 multivariate functional connectomes. Furthermore, I created neural feature sets consisting of the contribution of the ROIs of one functional network to the ROIs of another functional network, for example Language-Saliency, Sensorimotor-Dorsal, etc. These were selected based on established functional links in the literature or observed correlations with Mullen scores in Chapter 3.

All connectome feature sets were then used as observation vectors in random

forest classifiers (RFC) to predict the target variables of the labels of the Mullen score groupings. I found that for PCA-based dimensionality reduction of multivariate connectomes, only 20-component sPCA performed significantly better than chance. Of the network-network feature sets, only language-saliency performed better than chance. Pairwise networks fared marginally better, with 20-component PCA, sparse PCA, and sparse PCA with vision excluded, as well as saliency-frontoparietal, sensorimotor-dorsal attention, and language-frontoparietal network pairs predicting Mullen raw scores with better than chance accuracy. However, none of these feature sets predicted Mullen scores with sufficient accuracy to be viable for school readiness prediction for any practical application. Consequently I consider this a null result, but with many avenues for future analysis. In particular, selecting feature sets more finely tailored for their relevance to cognitive development either by more advanced dimensionality reduction techniques, or by specifying ROI-ROI pairs individually rather than on a functional network scale. Further analysis may also investigate the age-corrected Mullen t-scores within narrower age ranges (e.g., 6 months rather than over 2 years). These additional permutations of the prediction analysis in this chapter may elucidate further interesting structure in the functional connectomes. Taken altogether, these observations form a basis for more extensive future work to identify predictive network edges in the functional connectome.

Chapter 5

Acoustic features of complex sounds

5.1 Introduction and Motivation

The auditory system evolved to enable organismal survival by extracting actionable cues from complex natural sounds. It is thus necessary for neural populations in the auditory system to respond to acoustic features that enhance representations of key semantic information. The rodent auditory cortex has been extensively studied using simple parametric stimuli such as pure tones, as well as complex stimuli such as human speech. However, how neural populations in auditory cortex represent the acoustic space spanned by natural sounds remains poorly understood.

To represent ethologically relevant sounds, the auditory system must implement some form of feature selection that extracts a subset of meaningful acoustic properties. In order to survive, an organism must be able to discriminate sounds into behaviorally salient categories. Previous work has demonstrated the presence of categorical representations—that is, neural responses are invariant between different stimuli—of sound in the auditory cortex both in humans [115] and in animal models [116, 46, 52]. In particular, research in the rat auditory cortex has shown that distance between neural responses to human speech phonemes in rat A1 predicts behavioral ability to discriminate phonemes of human speech [117]. This suggests that high-level auditory representations enhance the discriminability of categories of complex sounds.

Research in the auditory cortex has largely emphasized single-neuron responses to simple, non-naturalistic stimuli. However, the rat primary auditory cortex has been shown to have a strongly tonotopic organization (see Figure 1 in [118] and Fig-




Figure 5.1: Figure courtesy of Baratham et al.[119]. Tonotopic map generated from μ ECoG recording from rat A1.

ure 5.1[119]), and encodes temporal and spectral acoustic features using a spatially distributed code [120, 121], as in other species [122, 123] including humans [124]. Furthermore, in other species, such as the ferret, the primary auditory cortex has been shown to have a wide range of spectro-temporal tuning, which permits A1 subpopulations to robustly encode acoustic features that predict, for example, phoneme categories [46]. It is thus likely that the rat primary auditory cortex encodes a variety of spectro-temporal features that enable and enhance categorical discrimination of sounds. Complex naturalistic sounds comprise combinations of spectral and temporal features, and can contain diverse ethologically relevant information. Thus, for natural auditory objects, it is critical to measure population responses for a full neural representation of the sound's acoustic properties in rat primary auditory cortex (A1).

It has been shown that linear models do not fully explain neural responses to complex auditory objects in both humans [47] and animal models [125]. Moreover, neural receptive fields in ferret auditory cortex are stimulus dependent, and STRFs

derived from synthetic sounds are less predictive of A1 response to natural sounds than those derived from natural sound statistics such as human speech [51]. In order to characterize neural representations of sound in the auditory cortex, it is thus necessary to extend current research to include neural responses to acoustically naturalistic, complex stimuli. Recent work in the human auditory system has shown that task-optimized deep neural network classifiers can produce receptive fields that effectively predict fMRI responses to natural auditory objects [126]. This finding indicates an opportunity to construct nonlinear models of neural representations that can be used to learn acoustic receptive fields that more accurately predict neural representations of auditory stimuli.

Despite decades of active research, the central auditory system remains poorly understood. Whereas contemporary efforts to model the primary visual cortex (V1) have pushed forward knowledge in vision science and developed strong theoretical frameworks of visual encoding principles [43], efforts to model the primary auditory cortex have met with less success [125, 127]. Characterization of the response properties of auditory cortex that leverages modern advances in computational modeling and insights from the visual system will serve to advance several frontiers of scientific and technological inquiry.

The prevalence of electrocorticography in human medicine and neuroscience suggests that using ECoG in animal models may facilitate translation from basic neuroscience to application. Mechanisms of auditory categorical representation [115] and semantic encoding [46, 50, 128], topics of interest in current auditory speech processing research, can be investigated in animal models using a semantically and acoustically rich database. The response properties and coding mechanisms studied using μ ECoG in the rat primary auditory cortex will be used to interpret ECoG data found in human subjects, and to develop a cohesive theory of auditory encoding.

The project presented in this chapter aimed to develop data-driven models of feature selectivity and structured variability by introducing key innovations in both experimental and computational methodologies. In particular, the database and acoustic analysis described in the chapter are intended to form the basis of an investigation of how rat primary auditory cortex represents complex natural sounds. To this end, I collected natural sounds with an emphasis on categories that are of ethological importance to a rat, with the intention of spanning the acoustic diversity of these categories. I supplemented these ethologically relevant categories with a wide variety of other sounds, including human speech, music, birdsong, and mechanical sounds. I then worked to identify acoustic features that discriminate semantic categories.

5.2 Methods

Building a complex sound dataset

My goal in assembling the sound database was to span the full range of semantic categories that present ethologically relevant information for rats. Furthermore, the sounds within each category must be acoustically naturalistic, and thoroughly span the range of acoustic features that typify that category. In this section, I will describe the assembly and analysis of the sounds I collected.

Auditory object collection

While collecting sound files, I considered four broad categories: predatory animal vocalizations (danger), non-predatory animal vocalizations, environmental sounds, and manmade sounds. Each of these categories comprises many subcategories that are semantically and acoustically diverse. Figure 5.2 shows the component subcategories.

I collected sound files from a wide array of sources. For categories such as cats, dogs, and vehicle noises, which are commonplace and used widely in various projects (e.g., films and video games), I was able to find a wealth of files on free sound databases. In particular, Zapsplat (with documented permission from the owner of the website) and Project Freesound (creative commons license) were the source of the majority of sound files for: all animal vocalizations with the exception of fox, bat, rat, and songbird vocalizations; all manmade sounds such as music, vehicle and machine sounds, and speech with the exception of included TIMIT files[129]; environmental sounds such as ocean sounds, wind, and snapping twigs. These sources have the benefit of immense volume and diversity, with many potential examples for each desired category.

Quality control was an ongoing concern in selection and exclusion of potential sounds. When scouring the internet for high-quality sound files, it became clear that there is a preponderance of engineered, artificial sound effects. This problem was more prevalent for particularly popular categories, such as cat and dog sounds. Naturalistic acoustic features were a particular priority for animal vocalizations, where auditory perception is likely to be sharpest and most discriminatory. For the animal vocalization category, I excluded any file that I suspected to be artificial based on the metadata labeling or as determined by ear. While this filtering process may have introduced researcher bias to the selection, it was necessary if only to spare us from a staggering number of humans meowing in pale imitation of cats.

I did not include any rodent vocalizations from the online sound databases. Due to their particular ethological relevance, a higher level of quality control was de-




Figure 5.2: Histogram of full sound database by number of discrete number of files (top) and total number of available seconds (bottom) in a given category.

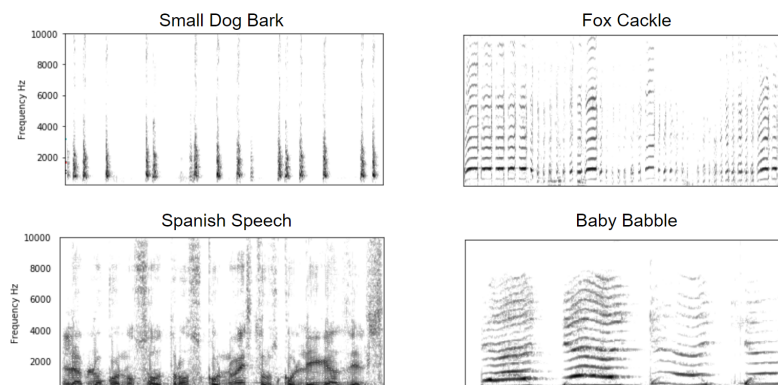


Figure 5.3: Spectrograms of different types of natural vocalizations drawn from the sound database. Note complex harmonic structure and temporal modulation.

sirable. In addition, the acoustic structure of rat ultrasonic vocalizations (USVs) have been shown to have ethological relevance that may contain salient semantic information for other rats[130]. For this reason, I reached out to authors Nobuaki Takahashi and Naoyuki Hironaka for access to their rat vocalization dataset. I received USVs collected from paired male Sprague-Dawley rats, recorded at 192 kHz by a microphone suspended from the cage lid. The files were subdivided based on recorded behavioral cues into three categories: feeding ($n = 12$), fighting ($n = 30$), and movement ($n = 19$).

In order to more thoroughly explore the acoustic features of predator vocalizations, I reached out to Dr. Svetlana Gogleva at Lomonosov Moscow State University for access to her extensive dataset of aggressive and domestic *Vulpes vulpes*[131][132], specifically Dmitry Belyaev’s silver foxes[133]. These sounds were recorded in response to a researcher approaching the foxes in their cages in five stages. Vocalizations were recorded at 22.5 kHz and classified by a researcher into eight types: whine ($n = 9$), moo ($n = 7$), growl ($n = 3$), pant ($n = 2$), cackle ($n = 6$), snort ($n = 5$), cough ($n = 4$), and bark ($n = 2$). While these vocalizations were collected from domesticated foxes, they present a wide behavioral and acoustic features.

I expanded the database further with songbird vocalizations from Dr. David Mets in Professor Michael Brainard’s lab at UCSF. Dr. Mets provided vocalizations from 4 different songbird species, selecting samples that capture the variation of the species. I received over 300 samples each of: Timor zebra finches, spice finches, Indian silverbills, and African silverbills. Finally, in order to include a greater variety




Figure 5.4: Distribution of semantic labels for primary analysis dataset of hand-selected sounds, Set A.

of high-pitched animal vocalizations, I received ($n = 16$) vocalizations by Egyptian fruit bats from Dr. Julie Elie at the Yartsev lab at UC Berkeley.

I then selected a set of 99 sound files for Set A, a set of auditory objects intended to span the space of semantic categories, with increased representation and acoustic diversity for ethologically relevant categories such as predator vocalizations. Figure 5.4 shows the total number of seconds included for each semantic category. As seen in Figure 5.4, some categories—most notably, songbirds—are strongly over-represented in the set. This is because recording duration varied significantly depending on semantic

category. Songbird call recordings in particular were relatively long (≈ 20 s), while ocean recordings were often well over a minute.

Sound file standardization

All sounds in Set A were standardized for subsequent analysis. In order to match the high sampling rate for the rat USV files, a sampling rate of 192kHz was chosen, and all other files were up-sampled using a Fourier transform. Bit depth was normalized to 32KB. The waveform amplitude was normalized to $|w(t)|_{max} = 1$ with clipping of outliers above the 99.99th percentile. All standardized sound waves were saved to a library of .h5 files, named based on the semantic category and subcategory (e.g., fox3_cackle) in order to facilitate programmatic parsing of semantic labels. Finally, in order to be compatible with the experimental speaker used to present sound stimuli to rats, all stereo sounds were converted to a mono-channel by averaging across left and right audio channels.

Auditory feature analysis

Sounds in Set A were used as a basis for the prediction of semantic labels from acoustic features. I partitioned each soundwave into 200 ms frames and treated each frame as an individual sound, inheriting the semantic labels of its parent sound. Since natural sounds have natural variability in power over time, this frame partitioning ran the risk of producing silent frames with very little spectral power. Inclusion of these silent frames runs the risk of biasing the characterization of semantic categories by their signature acoustic features. To mitigate this concern, I excluded all frames in the bottom 5th percentile of total power, as calculated by the L_1 norm of sound wave $w(t)$. Consult Figure 5.6 for an example of a BioSound-visualized object. This visualisation provides a standard, intuitive visualization of acoustic features, allowing for quick comparison and contrast of sounds.

I used the Theunissen lab python library, BioSound, to calculate the acoustic features of all sounds in Set A[134]. The features calculated are listed and briefly explained below, but a for an extensive description of these features, see Elie&Theunissen,2015[48]:

- **Spectrogram:** The spectrogram of a waveform $w(t)$ is calculated by performing a short-time Fourier transform (STFT) with a sliding window of length ω .
- **Temporal envelope and moments of temporal distribution:** The temporal envelope is a low-pass filtered version absolute value of the sound waveform.




Figure 5.5: Log distribution of the total power contained in each 200ms sound frame, as calculated by the \mathcal{L}_1 -norm of the waveform. 25th, 50th, and 75th percentiles are designated.

This envelope is treated as a distribution of the power of the sound over time, and the mean, standard deviation, skew, and kurtosis of the distribution is calculated.

- **Spectral density and moments of power spectrum distribution:** Power spectral density as calculated by a Fourier transform of the sound waveform. The mean, standard deviation, skew, and kurtosis of the spectral distribution is calculated.
- **Fundamental frequency:** The lowest frequency present in the sound, varying over time. The fundamental frequency was calculated by fitting the harmonic stacks of the sound using the cepstrum[135] $C_p = |\mathcal{F}^{-1} \{ \log (|\mathcal{F}\{f(t)\}|^2) \}|^2$
- **Pitch Saliency:** Pitch saliency is a calculation intended to approximate the percept of the strength of the fundamental frequency of a sound. For example, a flute has high pitch saliency, while the roar of a motorcycle has low pitch saliency. In this case, pitch saliency is calculated as the maximum peak of the sliding autocorrelation of the waveform, which creates peaks at time-

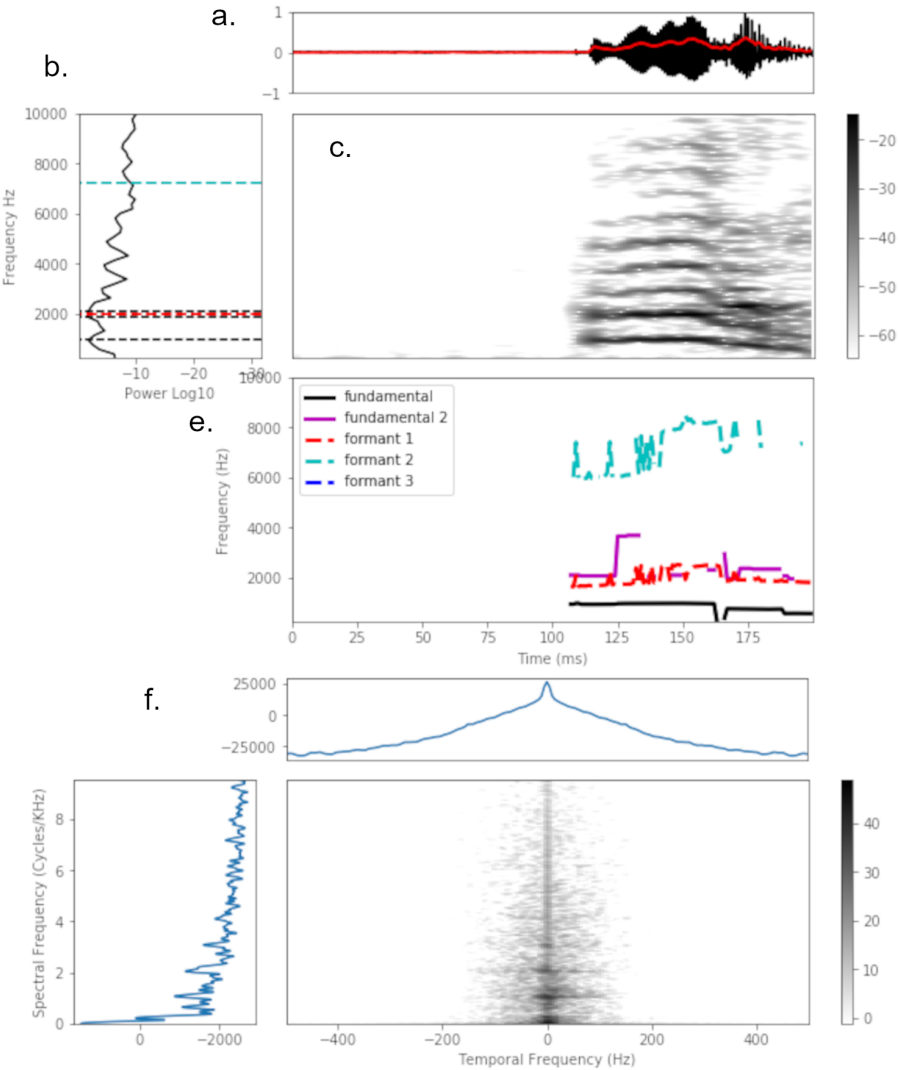


Figure 5.6: Example biosound of a 200ms frame of a fox cackling. **a.** Sound pressure waveform (black) with temporal envelope (red) **b.** Power spectral density (PSD). Quadrants 1, 2, 3 are superimposed in black, mean formants 1 (red) and 2 (cyan). **c.** STFT spectrogram. **d.** Frequency of fundamental frequencies and formants over time.

offsets where periodicity in the waveform creates high coherence. This peak is then normalized by the peak at 0 time offset (essentially normalizing by signal variance).

- **Formants:** Formants are calculated as peaks in the spectral envelope. Some sounds may have only one formant, while others have more.
- **Modulation power spectrum:** The modulation power spectrum is a 2 dimensional Fourier transform of the spectrogram of the sound, which yields a representation of the joint spectral and temporal modulation of the sound. For more information, see Singh&Theunissen, 2003[49].

I used the resulting biosounds to produce an acoustic feature vector for each sample. 29 features were calculated from BioSound features (Table 5.1).

Semantic prediction from auditory features

I investigated separability of semantic labels from acoustic features for all sound frames from Set A. To better understand the acoustic properties of different semantic groups, I visualized acoustic features for different subcategories. I began by visualizing fox, speech, and music frames, because those categories in Set A were of roughly similar sizes, and all three categories include acoustically complex sounds with diverse harmonic activity. See Figure 5.7.

Dimensionality reduction

Visualization of acoustic feature distributions, as in Figure 5.7 showed nonlinear underlying relationships between acoustic features. While no individual acoustic feature alone is strongly semantically discriminatory, dimensionality reduction can be used to reveal underlying structure that may improve separability. I utilized two methods of dimensionality reduction. First, I did principal components analysis to find dimensions of maximum variance in the acoustic space. Second, I used a nonlinear dimensionality reduction, UMAP (uniform manifold approximation and projection).

Principal components analysis: As described briefly in Chapter 4, principal components analysis (or PCA) decomposes the data into orthogonal dimensions that maximize the variance explained for each successive PC. I used PCA to generate an alternative feature set and visualize acoustic features that capture maximum variance.

UMAP: In order to capture structure in the acoustic data that may not be captured by linear methods such as PCA, I also computed feature vectors using

index	Feature	Calculation
0	Temporal mean	$\mu_T = \text{E} [\text{ENV}], \text{ENV} = \text{lowpass}(w(t))$
1	Temporal stdev	$\sigma_T = \sqrt{\text{E}[(w(t) - \mu_T)^2]}$
2	Temporal skew	$\gamma_T = \text{E} \left[\left(\frac{w(t) - \mu_T}{\sigma_T} \right)^3 \right]$
3	Temporal kurtosis	$\text{kurtosis} = \text{E} \left[\left(\frac{w(t) - \mu}{\sigma} \right)^4 \right]$
4	Temporal entropy	$H_t = -\frac{\sum_i^T w(t_i) \log_2 w(t)}{\log T}$
5	Pitch saliency	mean pitch saliency over time
6	Fundamental frequency	mean fundamental frequency
7	Formant 1	mean formant 1
8	Formant 2	mean formant 2
9	Formant 3	mean formant 3
10	Spectral mean	$\mu_S = \text{E}[P(\omega)], P(\omega) = \mathcal{F}\{w(t)\}$
11	Spectral stdev	$\sigma_S = \sqrt{\text{E}[(P(\omega) - \mu_S)^2]}$
12	Spectral skew	$\gamma_S = \text{E} \left[\left(\frac{P(\omega) - \mu_S}{\sigma_S} \right)^3 \right]$
13	spectral kurtosis	$\text{kurtosis} = \text{E} \left[\left(\frac{P(\omega) - \mu_S}{\sigma_S} \right)^4 \right]$
14	Spectral entropy	$H_s = -\frac{\sum_i^M P(\omega_i) \log_2 P(\omega_i)}{\log M}$
15	Q1 PSD	1st quartile of PSD, 25th perc
16	Q2 PSD	2nd quartile of PSD, 50th perc
17	Q3 PSD	3rd quartile of PSD, 75th perc
18	Spectral mod. mean	spectral mod. $= M_s = \sum_t^T 10 \log_{10} \text{MPS}$
19	Spectral mod. stdev	$\sigma(M_s)$
20	Spectral mod. skew	$\gamma(M_s)$
21	Spectral mod. kurtosis	kurtosis of M_s
22	Spectral mod. entropy	$H(M_s)$
23	Temporal mod. mean	temporal mod. $= M_t = \sum_\omega^M 10 \log_{10} \text{MPS}$
24	Temporal mod. stdev	$\sigma(M_t)$
25	Temporal mod. skew	$\gamma(M_t)$
26	Temporal mod. kurtosis	kurtosis of M_t
27	Temporal mod. entropy	$H(M_t)$
28	MPS separability α	SVD $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, $\alpha_{sep} = \frac{\sigma_1}{\sum \sigma}$

Table 5.1: Components of acoustic feature vectors calculated for all 200 ms sound frames.

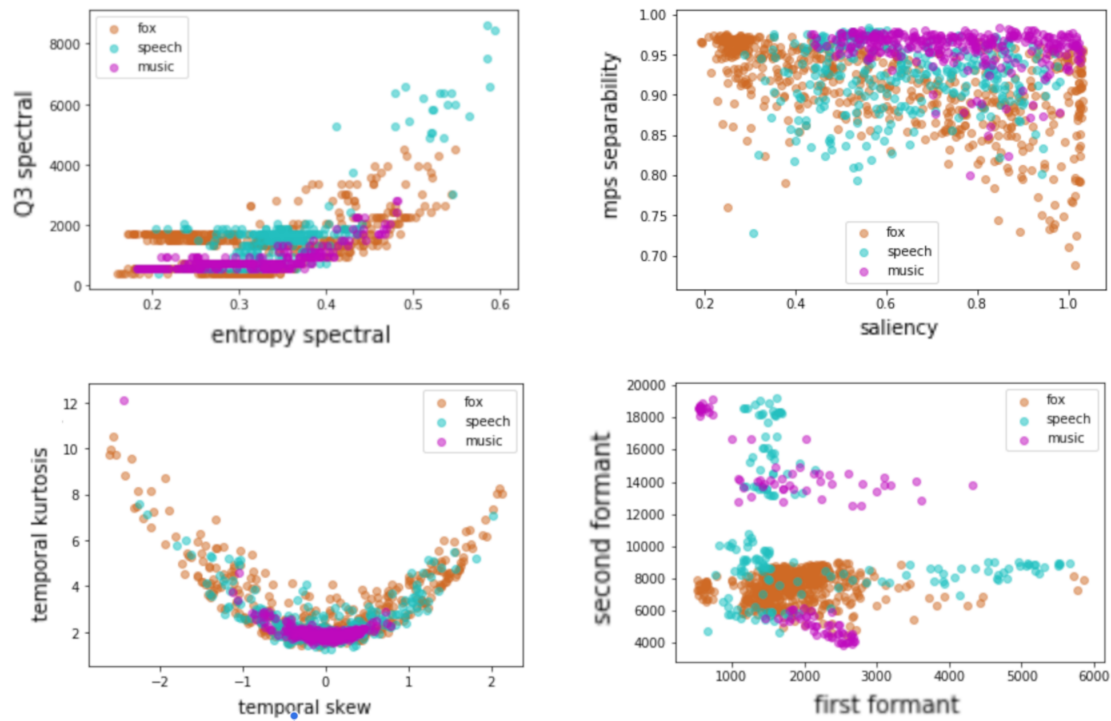


Figure 5.7: **Visualization of selected acoustic features shows some separability between example semantic categories.** In particular, foxes show high variability in their first formant, but tight distribution for their second. Vocalizations (fox and human) have high variability in modulation separability compared to music, which shows consistently high separability.

Uniform Manifold Approximation and Projection (UMAP). UMAP is a non-linear method of dimensionality reduction that assumes a Riemannian manifold on which the data is uniformly distributed[136]. The nonlinear nature of the algorithm makes mapping UMAP features back to acoustic space difficult. I used UMAP to create a low-dimensional feature set and tested their power to predict semantic labels.

Classification of natural sounds from acoustic feature vectors

I used a support vector machine (SVM) to classify all semantic categories based on different sets of acoustic features. Described very much in brief, SVM classifies data by mapping feature vectors to a high-dimensional space using a given kernel function $K(\vec{x}_1, \vec{x}_2)$, such that the distance between data points \vec{x}_1 , and \vec{x}_2 is calculated with that function instead of a linear dot product. This function can be a polynomial of degree d or an exponential function such as the sigmoid (hyperbolic tangent) or Gaussian radial basis function (RBF), where $\mathbf{K}(\vec{x}_1, \vec{x}_2) = \exp\left(-\frac{\|\vec{x}_1 - \vec{x}_2\|^2}{2\sigma^2}\right)$. The SVM then learns a hyperplane in the high-dimensional feature space that separates the target labels of the data with the maximized margin[137].

I used a SVM with an RBF kernel to classify all semantic groups in set A based on z-scored acoustic features, PCs, and UMAP features. Since semantic group sizes were not even between labels, I calculated chance-level scores by shuffling labels between data points and using the score of an SVM classifying the shuffled dataset. I repeated this process with $n_p = 50$ partitions of the data and set chance classification levels to the mean score of all partitions. I also excluded songbird vocalizations from initial classification efforts, since the category was so large as to drown out all other labels.

In Figure 5.8 I visualize confusion matrices, which compare the true label to the predicted label, to identify categories with high overlap in acoustic feature spaces.

5.3 Results and Discussion

Classification of semantics of natural sounds performs better than chance, but is subject to confusion in categories with high acoustic variability

Chance-level accuracy was ascertained from a 50-fold partition of shuffled classification to be 10.3%. Accuracy for z-scored features for 200 ms frames of Set A frames was 73.9%. This classification performance is replicated using PCA-test accuracy for an SVM using the first 15 principal components (PCs) was 72.9%. I noted from




Figure 5.8: **Prediction of semantic labels from acoustics performs well above chance, but semantic categories with high levels of acoustic variability are incorrectly assigned to other sound categories.** Confusion matrices with predicted semantic labels on the x-axis and true labels on the y axis. **Left:** first 15 PCA components. **Right:** Five UMAP dimensions.

visual inspection of the confusion matrix that certain semantic categories, specifically animal vocalizations and human speech, showed higher levels of confusion than others. For example, animal vocalizations with low sample size representation in the dataset, such as sheep and geese, were never predicted, but instead binned under labels with high acoustic variability such as foxes, dogs, and speech (see Figure 5.8). Given that PCs are simply linear combinations of the original features, it appears that these preferred bins such as foxes and speech span a large acoustic space, therefore lying close to other categories.

By contrast, classification accuracy for 10-component UMAP was relatively low: 63.1%, and varying significantly each time UMAP components are re-generated. Examination of the confusion matrix in Figure 5.8 shows similar dominance by a few categories: foxes, cats, corvids, and speech. Interestingly, while acoustic feature PCs bin many sounds under “dog,” UMAP seems to prefer cats and corvids. This

difference in preferred error categories may suggest that the features selected by this instance of UMAP change the distances of semantic categories in acoustic space, for example by moving the geese, horses, and pigs closer to cats.

It is necessary to note that these initial classification attempts are highly flawed. While the size of the dataset used provides a large number of samples, the examples within categories are far from statistically independent, since datapoints are frames drawn from a smaller number of files. For example, all fox datapoints are drawn from 8 files, all cat samples are drawn from 9 files, and all duck sounds are drawn from 2 files. This places a strong caveat on the nature of the classification task—are I measuring prediction of the semantic category, or simply similarity to the rest of the vocalization the sample is from? This analysis should be repeated using a greater range of sound files for each category, more balanced category sizes, and more careful quality control of 200 ms data frames included in the classification set.

5.4 Conclusion

Natural sounds contain a wealth of semantic information encoded by complex spectrotemporal structure. In this chapter, I described the collection and standardization of a large, rich dataset of high-quality natural sounds developed for use in experimental neuroscience. I aimed to collect both a wide array of sounds spanning as much of semantic space as possible, while also going “deep” on semantic categories with ethological salience for common animal models such as rats. For example, predator vocalizations such as fox and cat vocalizations, as well as ultrasonic vocalizations for rats and bats. I analysed the acoustic features of a selected subset of the natural sounds and tested their predictive power of semantic categories using simple machine learning techniques such as support vector machines.

The purpose of this work was to form the initial dataset for an extensive project aiming to characterize how neural representations of sound in the auditory cortex encode acoustic features that enhance categorical boundaries. This project ran aground due to my failure to collect usable neural data in acute experiments in rats, and subsequent move into the project detailed in chapters 3 and 4. Fortunately, this research is now to be continued by undergraduate researcher Vitto Resnick, who has sequences of natural sounds for use in an experimental context. He has graciously provided examples of this stimulus set, displayed in Figure 5.9, showing the concatenation of different natural sound clips. These stimulus sets will be presented to anesthetized rats, and μ ECoG (electrocorticography) arrays such as the one used to generate Figure 5.1 will be used to record the neural response of auditory cortex. This data will be used by members of the Bouchard lab to characterize distributed representations

of complex natural sounds in rat auditory cortex.

Chapter 6

Conclusion

Modern neuroimaging and data science provide a dizzying array of methods by which the mechanisms of neural computation can be investigated. Advances in statistical machine learning allow us to extract meaning from high-dimensional datasets, thus opening new frontiers in the analysis of neuroimaging techniques that This dissertation has discussed two of many open questions in neuroscience: first, how to estimate the connectivity underlying non-invasive whole-brain data recordings; second, how to identify salient structure of complex stimuli in order to elucidate neural representations of sensory information.

This thesis proposes a novel method of estimating functional connectivity that uses the stable, selective ensemble statistical machine learning framework UoI LASSO to generate parsimonious multivariate connectomes from EEG and fMRI data. I showed that the resulting connectomes were sparse and small-world, and spatially distributed in both recording modalities. In the EEG dataset, the contrast between multivariate functional networks and the more correlation-based pairwise connectomes was stark. Graph analysis of the functional connectivity networks showed that pairwise connectomes showed a tendency to lattice structure, suggesting heavy connection between proximate nodes, with daisy chaining connections substituting for long-distance connections. This lattice structure is highly inefficient from a graph theoretic perspective—for signal to propagate long distances across a lattice network, it must navigate many short pathways. In a system such as the brain, this type of network would be particularly suboptimal, since the propagation of signals through neural pathways involves complex cascades of highly synchronized activity at the molecular and cellular level, as well as the heavy metabolic cost of maintaining synaptic connections. It has been shown repeatedly in a variety of contexts that the brain ruthlessly prioritizes efficiency when performing its myriad complex and overlapping operations, from sensory encoding[43][44] to cortical connectivity[98]. A

lattice network functional connectivity estimate in EEG is thus likely the result of failing to explain away shared signal caused by volume conduction[138]. Further investigation confirmed that pairwise coupling is highly correlated with inter-electrode distance, further supporting the interpretation that pairwise estimates of functional connectivity for EEG are likely irreparably vulnerable to strong spurious correlations caused by signal leakage. The multivariate connectomes estimated from this data showed a mixture of short and long-range connections, and while connectomes for different subjects showed a range of sparsity levels (20-60%), all connectomes were decisively small-world. These properties persisted across age groups and frequency bands. Future studies should continue the analysis of the EEG connectomes by further quantifying their graph theoretic properties such as community detection, modularity, and the eigenspectrum of the Laplacian. These measures should be compared across frequency bands and age groups to quantify potential predictors of developmental maturity.

In fMRI, the difference between multivariate and pairwise connectomes was less immediately apparent. Thresholded pairwise networks were not as spatially concentrated or lattice-like. However, closer inspection showed core differences in the structure of multivariate connectomes. Many coupling weights were set consistently to zero, while coactivation of functional networks tended to be concentrated in a small subset of ROI pairs. Furthermore, multivariate connectomes selected for lateralized structure with strong positive connections for ipsilateral pathways and strong negative connections for contralateral pathways. Lateral structure and relationships between positive and negative connections were “washed out” in pairwise networks due to the inherent constraints of the method—without explaining away variance while generating coupling coefficients, many correlations are faintly positive due to shared connections. Further studies should dig deeper into the variance of multivariate connectomes, particularly network edges (ROI pairs) that have high variability between subjects. The visual network was one such locus. The lateral regions of the visual network were often selected as a regressor while fitting the activity of other ROIs, but the coupling coefficients vary widely between connectomes. It would be enlightening to examine the stability of these couplings across multiple folds fitting the same subject.

I trained random forest classifiers to predict cognitive developmental measures (the Mullen scales) from multivariate fMRI functional connectomes. I found that standard principal components performed no better than chance as predictors. This is likely because high-variance couplings in the networks are not necessarily encoding high-level cognitive tasks. Select network pairs performed better, albeit not by much (scores about 1 standard deviation better than chance at best). While this work yielded largely negative results, there are many options yet to explore. The selection

of neural features in particular offers many possibilities, such as the use of sparse prediction algorithms to learn combinations of ROI pairs that are statistically tailored to predict the Mullen scores.

Chapter 5 describes the collection and preliminary analysis of a large database of complex natural sounds. This database was created to be used as a source of naturalistic auditory stimuli, in particular for rats. There is ongoing work in the Bouchard lab to create stimulus sets from this database and collect neural data using μ ECoG arrays to record from cortical micro-columns in rat primary auditory cortex and adjoining areas. Beyond the collection of neural data, further studies may use machine learning techniques to predict semantic information from acoustic structure, thus uncovering spectrotemporal receptive fields that optimally discriminate semantic categorical boundaries. These normative models of acoustic discrimination can then be compared to neural responses to the stimuli in order to better understand cortical representations of acoustic features.

Chapter

Appendix

A Correlations of Mullen Scores to Connectomes

Included here are the Spearman correlation coefficients and p-values for pairwise and multivariate connectomes and Mullen Fine Motor, Gross Motor, Expressive Language, and Visual Reception raw and t-scores, as well as the Early Learning Composite. These figures are referenced in Chapter 3. The equivalent figure for Expressive Language is included as Figure 3.17.

B Principal components for multivariate and pairwise connectomes

Included here are the first 20 principal components (PCs) and sparse PCs for multivariate and pairwise functional connectomes used for prediction in Chapter 4.

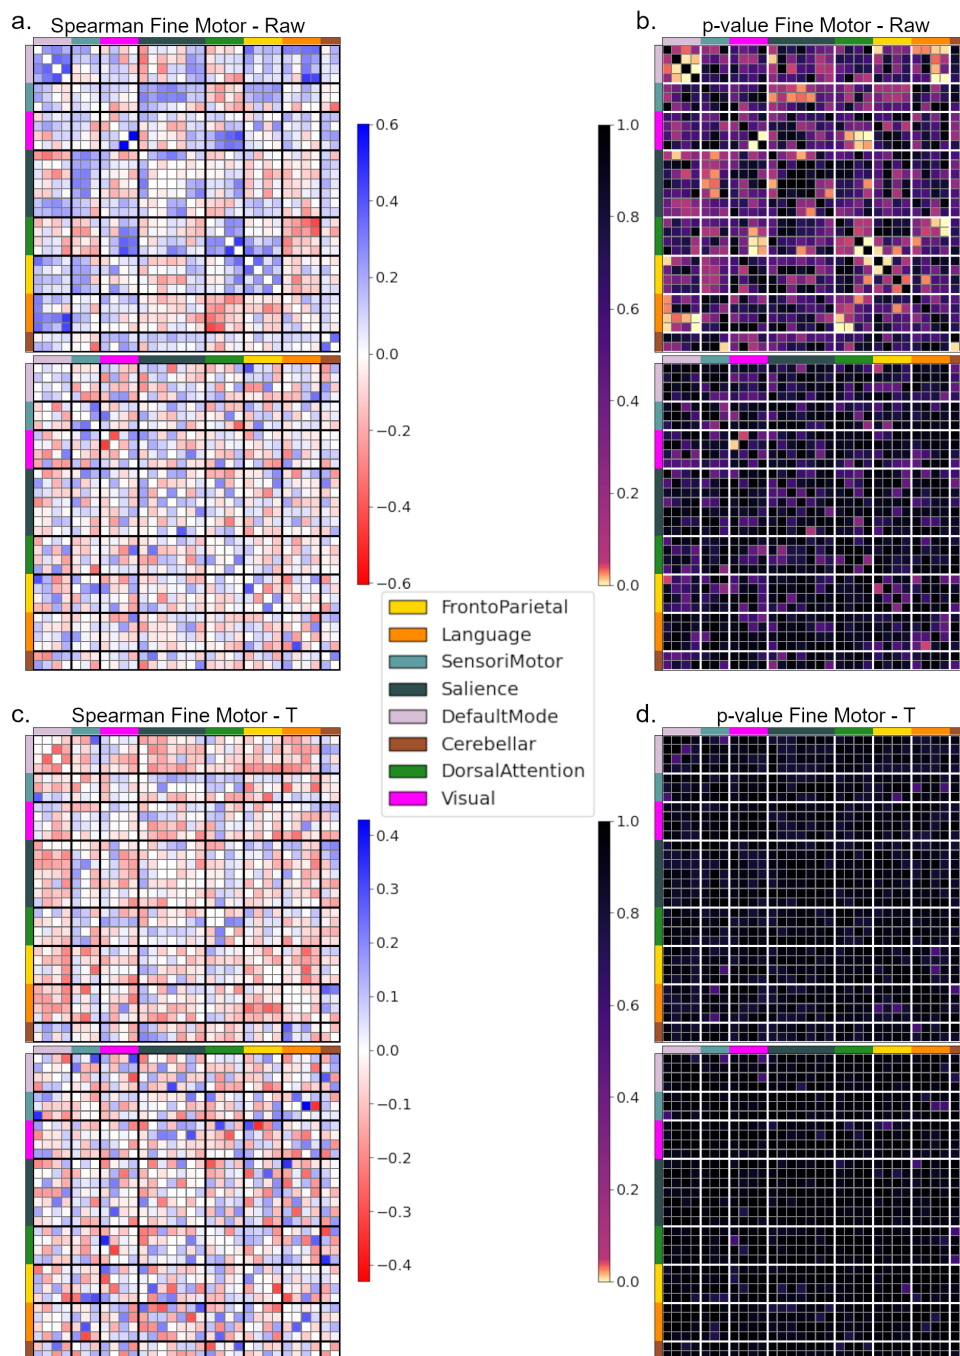


Figure 1: **Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations.** **a.** Spearman correlation coefficients and **b.** p-values between MSEL Fine Motor raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. **c.** Spearman correlation coefficients and **d.** p-values between MSEL Fine Motor t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.

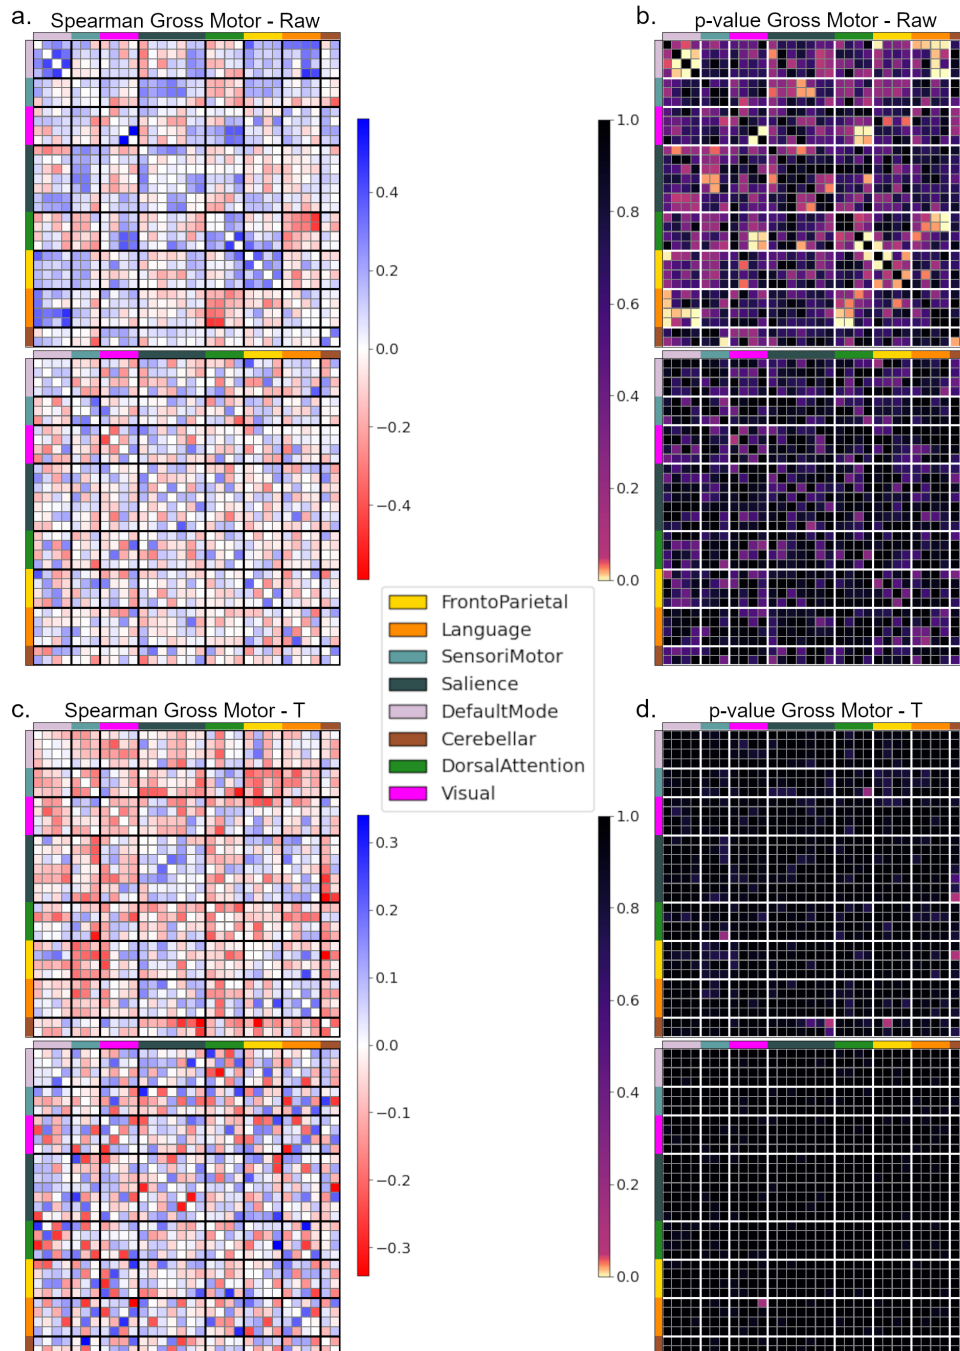


Figure 2: **Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations.** **a.** Spearman correlation coefficients and **b.** p-values between MSEL Gross Motor raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. **c.** Spearman correlation coefficients and **d.** p-values between MSEL Gross Motor t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.

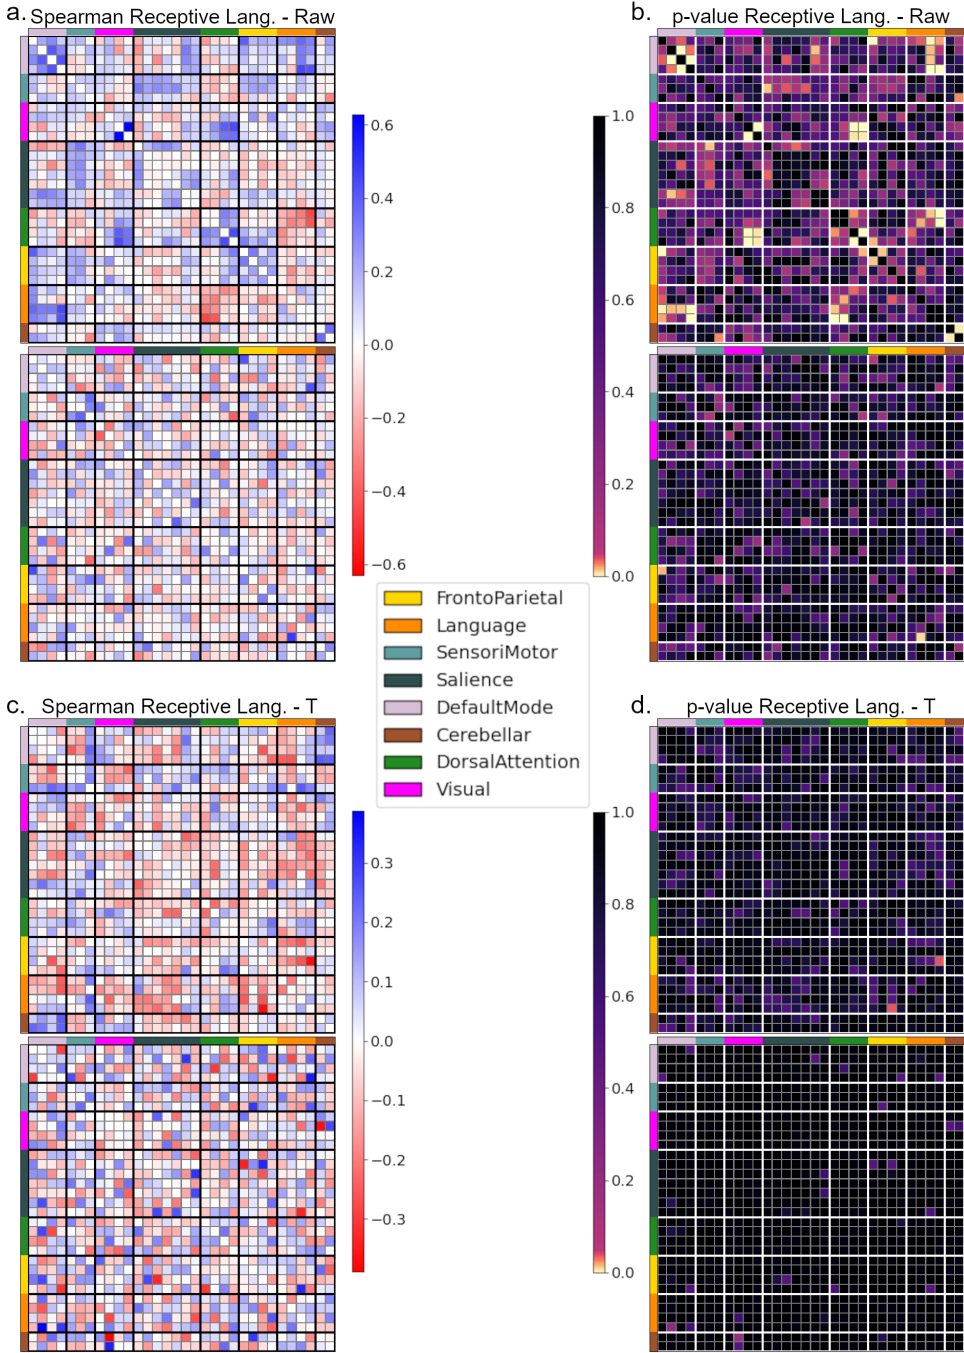


Figure .3: **Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations.** **a.** Spearman correlation coefficients and **b.** p-values between MSEL Receptive Language raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. **c.** Spearman correlation coefficients and **d.** p-values between MSEL Receptive Language t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.

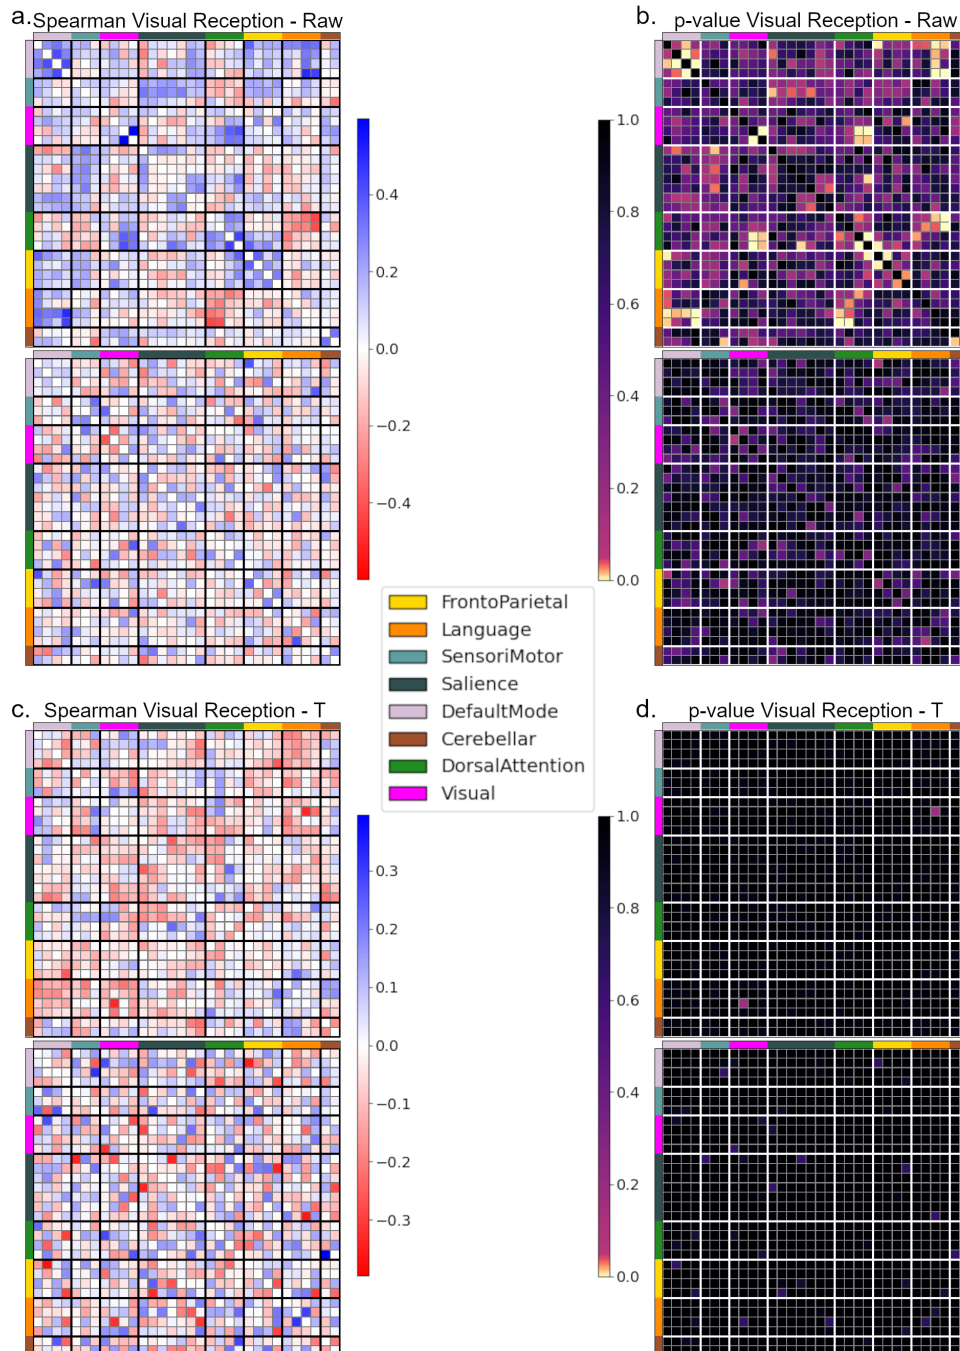


Figure 4: **Pairwise connectomes show widespread significant correlations with raw scores, while multivariate networks have no significant correlations.** **a.** Spearman correlation coefficients and **b.** p-values between MSEL Visual Reception raw scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices. **c.** Spearman correlation coefficients and **d.** p-values between MSEL Visual Reception t-scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.

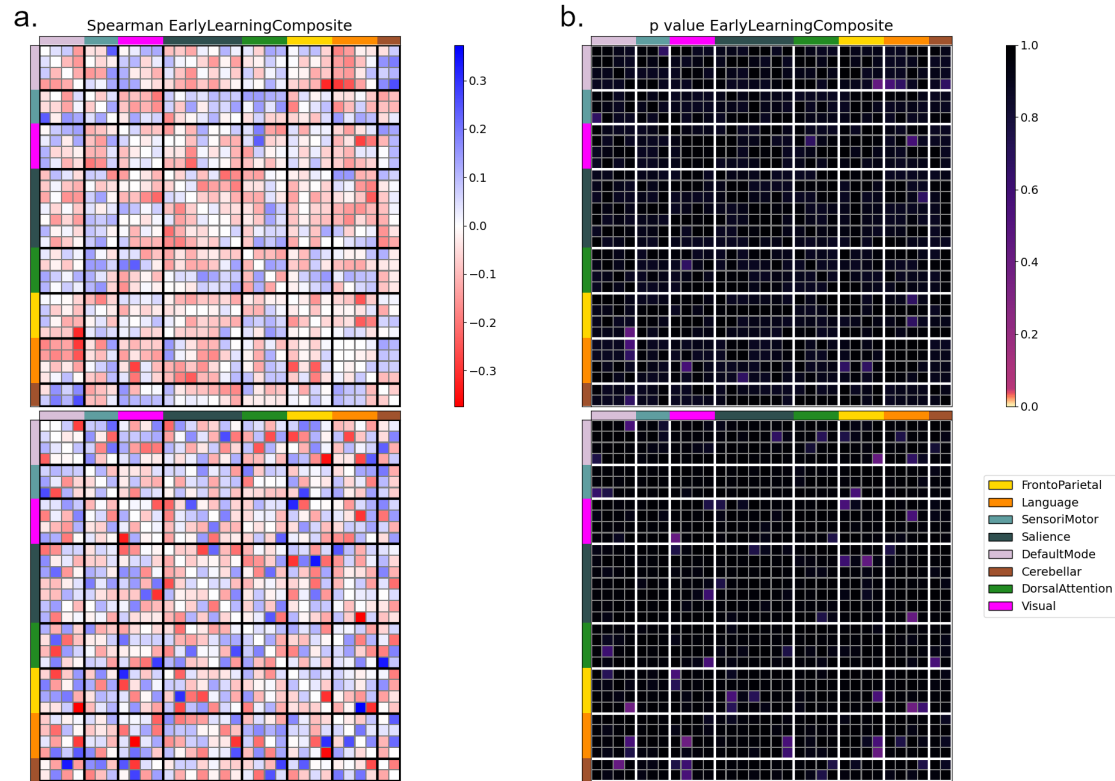


Figure .5: **Neither multivariate or pairwise connectomes show statistically significant correlates to Early Learning Composite score.** a. Spearman correlation coefficients and b. p-values between MSEL Early Learning Composite scores and pairwise (top) and multivariate (bottom) connectomes plotted as weight matrices.




Figure .6: Cumulative variance captured for multivariate connectome PCA. Blue lines show total variance captured by 20 PCs, while red lines show 80% variance captured.




Figure .7: Cumulative variance captured for pairwise connectome PCA. Blue lines show total variance captured by 20 PCs, while red lines show 80% variance captured.




Figure .8: Principal components 1 (top) and 11 (bottom) of multivariate connectomes.




Figure .9: Principal components 2 (top) and 12 (bottom) of multivariate connectomes.

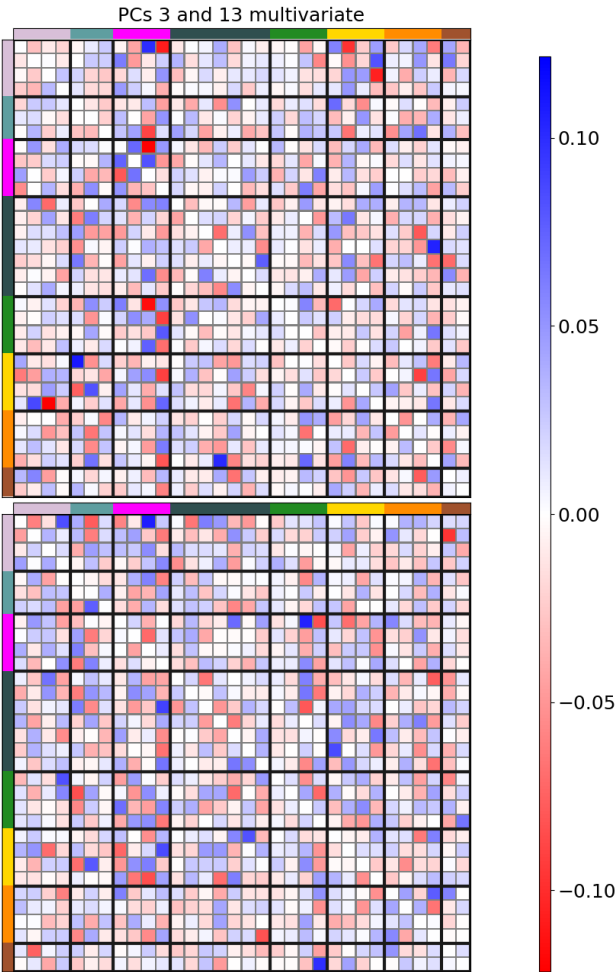


Figure .10: Principal components 3 (top) and 13 (bottom) of multivariate connectomes.




Figure .11: Principal components 4 (top) and 14 (bottom) of multivariate connectomes.




Figure .12: Principal components 5 (top) and 15 (bottom) of multivariate connectomes.




Figure .13: Principal components 6 (top) and 16 (bottom) of multivariate connectomes.




Figure .14: Principal components 7 (top) and 17 (bottom) of multivariate connectomes.




Figure .15: Principal components 8 (top) and 18 (bottom) of multivariate connectomes.

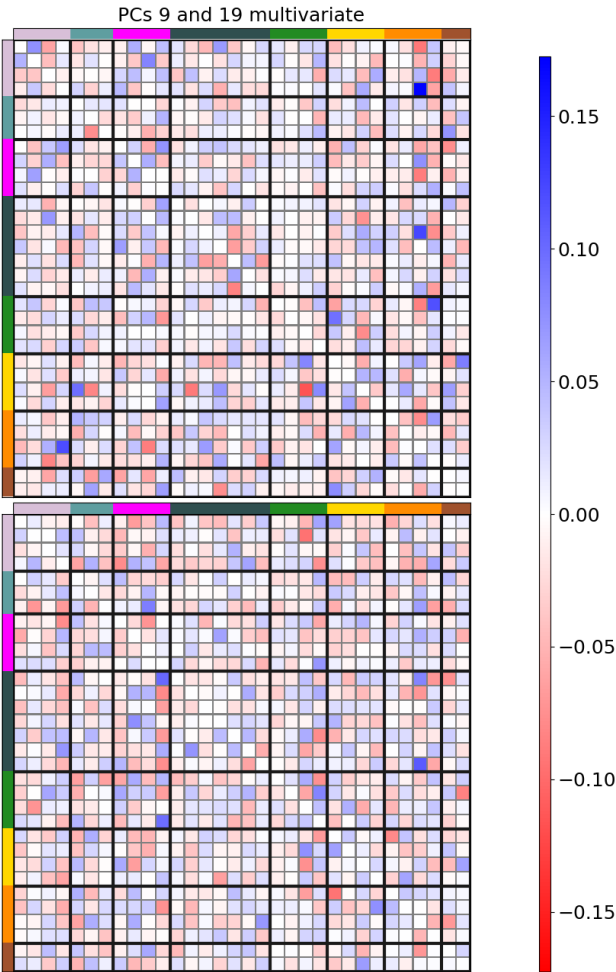


Figure .16: Principal components 9 (top) and 19 (bottom) of multivariate connectomes.

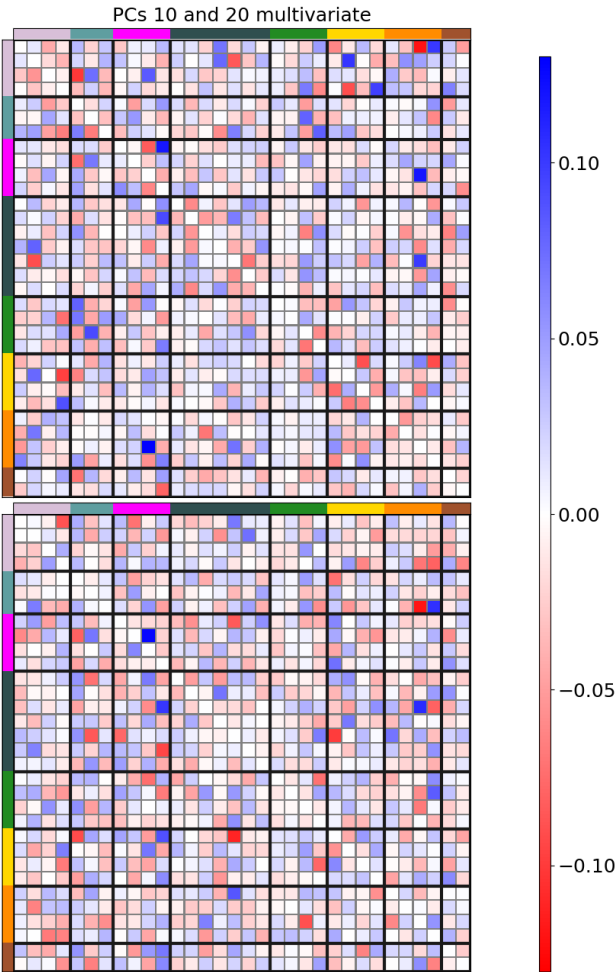


Figure .17: Principal components 10 (top) and 20 (bottom) of multivariate connectomes.




Figure .18: Sparse principal components 1 (top) and 11 (bottom) of multivariate connectomes.




Figure .19: Sparse principal components 2 (top) and 12 (bottom) of multivariate connectomes.




Figure .20: Sparse principal components 3 (top) and 13 (bottom) of multivariate connectomes.




Figure .21: Sparse principal components 4 (top) and 14 (bottom) of multivariate connectomes.




Figure .22: Sparse principal components 5 (top) and 15 (bottom) of multivariate connectomes.




Figure .23: Sparse principal components 6 (top) and 16 (bottom) of multivariate connectomes.




Figure .24: Sparse principal components 7 (top) and 17 (bottom) of multivariate connectomes.




Figure .25: Sparse principal components 8 (top) and 18 (bottom) of multivariate connectomes.




Figure .26: Sparse principal components 9 (top) and 19 (bottom) of multivariate connectomes.




Figure .27: Sparse principal components 10 (top) and 20 (bottom) of multivariate connectomes.




Figure .28: Principal components 1 (top) and 11 (bottom) of pairwise connectomes.




Figure .29: Principal components 2 (top) and 12 (bottom) of pairwise connectomes.




Figure .30: Principal components 3 (top) and 13 (bottom) of pairwise connectomes.




Figure .31: Principal components 4 (top) and 14 (bottom) of pairwise connectomes.




Figure .32: Principal components 5 (top) and 15 (bottom) of pairwise connectomes.




Figure .33: Principal components 6 (top) and 16 (bottom) of pairwise connectomes.




Figure .34: Principal components 7 (top) and 17 (bottom) of pairwise connectomes.




Figure .35: Principal components 8 (top) and 18 (bottom) of pairwise connectomes.




Figure .36: Principal components 9 (top) and 19 (bottom) of pairwise connectomes.




Figure .37: Principal components 10 (top) and 20 (bottom) of pairwise connectomes.




Figure .38: Sparse principal components 1 (top) and 11 (bottom) of pairwise connectomes.




Figure .39: Sparse principal components 2 (top) and 12 (bottom) of pairwise connectomes.

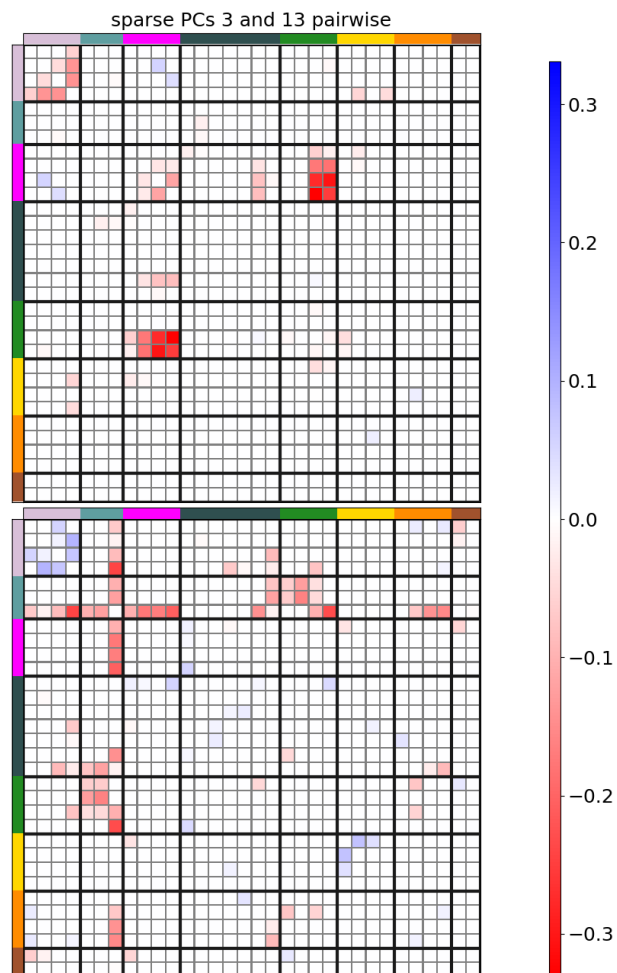


Figure .40: Sparse principal components 3 (top) and 13 (bottom) of pairwise connectomes.




Figure .41: Sparse principal components 4 (top) and 14 (bottom) of pairwise connectomes.




Figure .42: Sparse principal components 5 (top) and 15 (bottom) of pairwise connectomes.




Figure .43: Sparse principal components 6 (top) and 16 (bottom) of pairwise connectomes.




Figure .44: Sparse principal components 7 (top) and 17 (bottom) of pairwise connectomes.




Figure .45: Sparse principal components 8 (top) and 18 (bottom) of pairwise connectomes.




Figure .46: Sparse principal components 9 (top) and 19 (bottom) of pairwise connectomes.




Figure .47: Sparse principal components 10 (top) and 20 (bottom) of pairwise connectomes.

Bibliography

- [1] Jean-Michel Mongeau et al. “Multimodal integration across spatiotemporal scales to guide invertebrate locomotion”. In: *Integrative and comparative biology* 61.3 (2021), pp. 842–853.
- [2] Bin He et al. “Electrophysiological imaging of brain activity and connectivity—challenges and opportunities”. In: *IEEE transactions on biomedical engineering* 58.7 (2011), pp. 1918–1931.
- [3] Sungho Tak and Jong Chul Ye. “Statistical analysis of fNIRS data: a comprehensive review”. In: *Neuroimage* 85 (2014), pp. 72–91.
- [4] Gérard Derosière et al. *NIRS-measured prefrontal cortex activity in neuroergonomics: strengths and weaknesses*. 2013.
- [5] Sylvain Baillet. “Magnetoencephalography for brain electrophysiology and imaging”. In: *Nature neuroscience* 20.3 (2017), pp. 327–339.
- [6] C.J. Stam. “Use of magnetoencephalography (MEG) to study functional brain networks in neurodegenerative disorders”. In: *Journal of the Neurological Sciences* 289.1 (2010). Mental Dysfunction in Parkinson’s Disease, pp. 128–134. ISSN: 0022-510X. DOI: <https://doi.org/10.1016/j.jns.2009.08.028>. URL: <https://www.sciencedirect.com/science/article/pii/S0022510X09007849>.
- [7] Neil Schaul. “The fundamental neural mechanisms of electroencephalography”. In: *Electroencephalography and clinical neurophysiology* 106.2 (1998), pp. 101–107.
- [8] Josef Parvizi and Sabine Kastner. “Promises and limitations of human intracranial electroencephalography”. In: *Nature neuroscience* 21.4 (2018), pp. 474–483.
- [9] Denis Le Bihan. “Looking into the functional architecture of the brain with diffusion MRI”. In: *Nature reviews neuroscience* 4.6 (2003), pp. 469–480.

- [10] Denis Le Bihan. “Diffusion, confusion and functional MRI”. In: *Neuroimage* 62.2 (2012), pp. 1131–1136.
- [11] Derek K Jones. “Challenges and limitations of quantifying brain connectivity in vivo with diffusion MRI”. In: *Imaging in Medicine* 2.3 (2010), p. 341.
- [12] Franck Amyot et al. “A review of the effectiveness of neuroimaging modalities for the detection of traumatic brain injury”. In: *Journal of neurotrauma* 32.22 (2015), pp. 1693–1721.
- [13] Felix Biessmann et al. “Analysis of multimodal neuroimaging data”. In: *IEEE reviews in biomedical engineering* 4 (2011), pp. 26–58.
- [14] Bokkyu Kim and Carolee Winstein. “Can neurological biomarkers of brain impairment be used to predict poststroke motor recovery? A systematic review”. In: *Neurorehabilitation and neural repair* 31.1 (2017), pp. 3–24.
- [15] Elisa Menozzi and Anthony HV Schapira. “Exploring the genotype–phenotype correlation in GBA-Parkinson disease: clinical aspects, biomarkers, and potential modifiers”. In: *Frontiers in Neurology* 12 (2021), p. 694764.
- [16] Joanne Ryan et al. “Phenotypic heterogeneity in dementia: a challenge for epidemiology and biomarker studies”. In: *Frontiers in public health* 6 (2018), p. 181.
- [17] Frank Birklein et al. “Complex regional pain syndrome—phenotypic characteristics and potential biomarkers”. In: *Nature Reviews Neurology* 14.5 (2018), pp. 272–284.
- [18] Izabela Guimarães Barbosa et al. “Executive dysfunction in euthymic bipolar disorder patients and its association with plasma biomarkers”. In: *Journal of affective disorders* 137.1-3 (2012), pp. 151–155.
- [19] Noriaki Yahata, Kiyoto Kasai, and Mitsuo Kawato. “Computational neuroscience approach to biomarkers and treatments for mental disorders”. In: *Psychiatry and clinical neurosciences* 71.4 (2017), pp. 215–237.
- [20] Richard E Frye et al. “Emerging biomarkers in autism spectrum disorder: a systematic review”. In: *Annals of translational medicine* 7.23 (2019).
- [21] Daniel Ruivo Marques et al. “Insomnia disorder and brain’s default-mode network”. In: *Current neurology and neuroscience reports* 18 (2018), pp. 1–4.
- [22] Wilson Truccolo et al. “Single-neuron dynamics in human focal epilepsy”. In: *Nature neuroscience* 14.5 (2011), pp. 635–641.

- [23] Zhiqiang Zhang et al. “Altered spontaneous neuronal activity of the default-mode network in mesial temporal lobe epilepsy”. In: *Brain research* 1323 (2010), pp. 152–160.
- [24] Diane L Damiano. “Activity, activity, activity: rethinking our physical therapy approach to cerebral palsy”. In: *Physical therapy* 86.11 (2006), pp. 1534–1540.
- [25] Timothy J Teyler and Pascal DiScenna. “The hippocampal memory indexing theory.” In: *Behavioral neuroscience* 100.2 (1986), p. 147.
- [26] Timothy J. Teyler and Jerry W. Rudy. “The hippocampal indexing theory and episodic memory: Updating the index”. In: *Hippocampus* 17.12 (2007), pp. 1158–1169. DOI: <https://doi.org/10.1002/hipo.20350>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/hipo.20350>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/hipo.20350>.
- [27] Suzanne Corkin. “Lasting Consequences of Bilateral Medial Temporal Lobectomy: Clinical Course and Experimental Findings in H.M”. eng. In: *Seminars in neurology* 4.2 (1984), pp. 249–259. ISSN: 0271-8235.
- [28] Alexander R Backus et al. “Mnemonic convergence in the human hippocampus”. en. In: *Nat. Commun.* 7.1 (Sept. 2016), p. 11991.
- [29] Daniel Levenstein et al. “On the Role of Theory and Modeling in Neuroscience”. In: *Journal of Neuroscience* 43.7 (2023), pp. 1074–1088. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1179-22.2022. eprint: <https://www.jneurosci.org/content/43/7/1074.full.pdf>. URL: <https://www.jneurosci.org/content/43/7/1074>.
- [30] György Buzsáki, Costas A Anastassiou, and Christof Koch. “The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes”. In: *Nature reviews neuroscience* 13.6 (2012), pp. 407–420.
- [31] Eric Courchesne. “15 Cognitive Components of the Event-Related Brain Potential: Changes Associated with Development”. In: *Tutorials in Event Related Potential Research: Endogenous Components*. Ed. by Anthony W.K. Gaillard and Walter Ritter. Vol. 10. Advances in Psychology. North-Holland, 1983, pp. 329–344. DOI: [https://doi.org/10.1016/S0166-4115\(08\)62046-4](https://doi.org/10.1016/S0166-4115(08)62046-4). URL: <https://www.sciencedirect.com/science/article/pii/S0166411508620464>.
- [32] Eric Courchesne, Leo Ganz, and Anthony M. Norcia. “Event-Related Brain Potentials to Human Faces in Infants”. In: *Child Development* 52.3 (1981), pp. 804–811. ISSN: 00093920, 14678624. URL: <http://www.jstor.org/stable/1129080> (visited on 03/14/2024).

- [33] Taeko Sasai-Sakuma and Yuichi Inoue. “Differences in electroencephalographic findings among categories of narcolepsy-spectrum disorders”. en. In: *Sleep Med.* 16.8 (Aug. 2015), pp. 999–1005.
- [34] M E Raichle et al. “A default mode of brain function”. en. In: *Proc. Natl. Acad. Sci. U. S. A.* 98.2 (Jan. 2001), pp. 676–682.
- [35] Daniel P Kennedy, Elizabeth Redcay, and Eric Courchesne. “Failing to deactivate: resting functional abnormalities in autism”. In: *Proceedings of the National Academy of Sciences* 103.21 (2006), pp. 8275–8280.
- [36] Vladimir L Cherkassky et al. “Functional connectivity in a baseline resting-state network in autism”. In: *Neuroreport* 17.16 (2006), pp. 1687–1690.
- [37] Vinod Menon. “Large-scale brain networks and psychopathology: a unifying triple network model”. In: *Trends in Cognitive Sciences* 15.10 (2011), pp. 483–506. ISSN: 1364-6613. DOI: <https://doi.org/10.1016/j.tics.2011.08.003>. URL: <https://www.sciencedirect.com/science/article/pii/S1364661311001719>.
- [38] Scott Marek and Nico U F Dosenbach. “The frontoparietal network: function, electrophysiology, and importance of individual precision mapping”. en. In: *Dialogues Clin. Neurosci.* 20.2 (June 2018), pp. 133–140.
- [39] Muriel M K Bruchhage et al. “Functional connectivity correlates of infant and early childhood cognitive development”. en. In: *Brain Struct. Funct.* 225.2 (Mar. 2020), pp. 669–681. URL: <https://creativecommons.org/licenses/by/4.0>.
- [40] Alexander G Huth et al. “Decoding the semantic content of natural movies from human brain activity”. en. In: *Front. Syst. Neurosci.* 10 (Oct. 2016), p. 81.
- [41] Shinji Nishimoto et al. “Reconstructing visual experiences from brain activity evoked by natural movies”. en. In: *Curr. Biol.* 21.19 (Oct. 2011), pp. 1641–1646.
- [42] Nico U. F. Dosenbach et al. “Prediction of Individual Brain Maturity Using fMRI”. In: *Science* 329.5997 (2010), pp. 1358–1361. DOI: 10.1126/science.1194144. eprint: <https://www.science.org/doi/pdf/10.1126/science.1194144>. URL: <https://www.science.org/doi/abs/10.1126/science.1194144>.
- [43] Bruno A Olshausen and David J Field. “Wavelet-like receptive fields emerge from a network that learns sparse codes for natural images”. In: *Nature* 381 (1996), pp. 607–609.

- [44] Evan C Smith and Michael S Lewicki. “Efficient auditory coding”. en. In: *Nature* 439.7079 (Feb. 2006), pp. 978–982.
- [45] Shihab Shamma. “The acoustic features of speech sounds in a model of auditory processing: vowels and voiceless fricatives”. In: *Journal of Phonetics* 16.1 (1988). Representation of Speech in the Auditory Periphery, pp. 77–91. ISSN: 0095-4470. DOI: [https://doi.org/10.1016/S0095-4470\(19\)30467-X](https://doi.org/10.1016/S0095-4470(19)30467-X). URL: <https://www.sciencedirect.com/science/article/pii/S009544701930467X>.
- [46] Nima Mesgarani et al. “Phoneme representation and classification in primary auditory cortex”. In: *The Journal of the Acoustical Society of America* 123.2 (2008), pp. 899–909.
- [47] Brian N Pasley et al. “Reconstructing speech from human auditory cortex”. In: *PLoS biology* 10.1 (2012), e1001251.
- [48] Julie E Elie and Frédéric E Theunissen. “The vocal repertoire of the domesticated zebra finch: a data-driven approach to decipher the information-bearing acoustic features of communication signals”. en. In: *Anim. Cogn.* 19.2 (Mar. 2016), pp. 285–315.
- [49] Nandini C. Singh and Frédéric E. Theunissen. “Modulation spectra of natural sounds and ethological theories of auditory processing”. In: *The Journal of the Acoustical Society of America* 114.6 (Dec. 2003), pp. 3394–3411. ISSN: 0001-4966. DOI: 10.1121/1.1624067. eprint: https://pubs.aip.org/asa/jasa/article-pdf/114/6/3394/8093145/3394_1_online.pdf. URL: <https://doi.org/10.1121/1.1624067>.
- [50] Nima Mesgarani et al. “Mechanisms of noise robust representation of speech in primary auditory cortex”. In: *Proceedings of the National Academy of Sciences* 111.18 (2014), pp. 6792–6797.
- [51] Stephen V David et al. “Rapid synaptic depression explains nonlinear modulation of spectro-temporal tuning in primary auditory cortex by natural stimuli”. In: *Journal of Neuroscience* 29.11 (2009), pp. 3374–3386.
- [52] Joji Tsunada, Jung Hoon Lee, and Yale E Cohen. “Representation of speech categories in the primate auditory cortex”. In: *Journal of neurophysiology* 105.6 (2011), pp. 2634–2646.
- [53] Bing Du et al. “fMRI brain decoding and its applications in brain-computer interface: A survey”. en. In: *Brain Sci.* 12.2 (Feb. 2022), p. 228.

- [54] Quentin Huys, Joshua Vogelstein, and Peter Dayan. “Psychiatry: Insights into depression through normative decision-making models”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Koller et al. Vol. 21. Curran Associates, Inc., 2008. URL: https://proceedings.neurips.cc/paper_files/paper/2008/file/d04d42cdf14579cd294e5079e0745411-Paper.pdf.
- [55] Pratik S. Sachdeva et al. “Improved inference in coupling, encoding, and decoding models and its consequence for neuroscientific interpretation”. In: *Journal of Neuroscience Methods* 358 (2021), p. 109195. ISSN: 0165-0270. DOI: <https://doi.org/10.1016/j.jneumeth.2021.109195>. URL: <https://www.sciencedirect.com/science/article/pii/S0165027021001308>.
- [56] Naomi Altman and Martin Krzywinski. “The curse(s) of dimensionality”. en. In: *Nat. Methods* 15.6 (June 2018), pp. 399–400.
- [57] Surya Ganguli and Haim Sompolinsky. “Compressed Sensing, Sparsity, and Dimensionality in Neuronal Information Processing and Data Analysis”. In: *Annual Review of Neuroscience* 35.1 (2012). PMID: 22483042, pp. 485–508. DOI: [10.1146/annurev-neuro-062111-150410](https://doi.org/10.1146/annurev-neuro-062111-150410). eprint: <https://doi.org/10.1146/annurev-neuro-062111-150410>. URL: <https://doi.org/10.1146/annurev-neuro-062111-150410>.
- [58] Kristofer Bouchard et al. “Union of Intersections (UoI) for Interpretable Data Driven Discovery and Prediction”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/788d986905533aba051261497ecffcbb-Paper.pdf.
- [59] Mark Bear, Barry Connors, and Michael A Paradiso. *Neuroscience: exploring the brain, enhanced edition: exploring the brain*. Jones & Bartlett Learning, 2020.
- [60] Vinod Menon. “20 years of the default mode network: A review and synthesis”. In: *Neuron* 111.16 (2023), pp. 2469–2487. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2023.04.023>. URL: <https://www.sciencedirect.com/science/article/pii/S0896627323003082>.
- [61] Vinod Menon and Lucina Q Uddin. “Saliency, switching, attention and control: a network model of insula function”. In: *Brain structure and function* 214 (2010), pp. 655–667.

- [62] Miranda Scolari, Katharina N Seidl-Rathkopf, and Sabine Kastner. “Functions of the human frontoparietal attention network: Evidence from neuroimaging”. In: *Current Opinion in Behavioral Sciences* 1 (2015). Cognitive control, pp. 32–39. ISSN: 2352-1546. DOI: <https://doi.org/10.1016/j.cobeha.2014.08.003>. URL: <https://www.sciencedirect.com/science/article/pii/S2352154614000138>.
- [63] Paola Odriozola et al. “Insula response and connectivity during social and non-social attention in children with autism”. In: *Social cognitive and affective neuroscience* 11.3 (2016), pp. 433–444.
- [64] Julien Bastin et al. “Direct Recordings from Human Anterior Insula Reveal its Leading Role within the Error-Monitoring Network”. In: *Cerebral Cortex* 27.2 (Jan. 2016), pp. 1545–1557. ISSN: 1047-3211. DOI: [10.1093/cercor/bhv352](https://doi.org/10.1093/cercor/bhv352). eprint: <https://academic.oup.com/cercor/article-pdf/27/2/1545/10907583/bhv352.pdf>. URL: <https://doi.org/10.1093/cercor/bhv352>.
- [65] Tor D. Wager et al. “An fMRI-Based Neurologic Signature of Physical Pain”. In: *New England Journal of Medicine* 368.15 (2013). PMID: 23574118, pp. 1388–1397. DOI: [10.1056/NEJMoa1204471](https://doi.org/10.1056/NEJMoa1204471). eprint: <https://doi.org/10.1056/NEJMoa1204471>. URL: <https://doi.org/10.1056/NEJMoa1204471>.
- [66] Emma Twait, Rola Farah, and Tzipi Horowitz-Kraus. “Decreased functional connectivity of the salience network during narrative comprehension in children with reading difficulties: An fMRI study”. In: *NeuroImage: Clinical* 20 (2018), pp. 987–992. ISSN: 2213-1582. DOI: <https://doi.org/10.1016/j.nicl.2018.10.006>. URL: <https://www.sciencedirect.com/science/article/pii/S2213158218303164>.
- [67] Lucina Q. Uddin. “Chapter 3 - Functions of the Salience Network”. In: *Salience Network of the Human Brain*. Ed. by Lucina Q. Uddin. San Diego: Academic Press, 2017, pp. 11–16. ISBN: 978-0-12-804593-0. DOI: <https://doi.org/10.1016/B978-0-12-804593-0.00003-5>. URL: <https://www.sciencedirect.com/science/article/pii/B9780128045930000035>.
- [68] Devarajan Sridharan, Daniel J Levitin, and Vinod Menon. “A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks”. In: *Proceedings of the National Academy of Sciences* 105.34 (2008), pp. 12569–12574.
- [69] Steven L Bressler et al. “Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention”. en. In: *J. Neurosci.* 28.40 (Oct. 2008), pp. 10056–10061.

- [70] Maurizio Corbetta, Gaurav Patel, and Gordon L Shulman. “The reorienting system of the human brain: from environment to theory of mind”. en. In: *Neuron* 58.3 (May 2008), pp. 306–324.
- [71] Kristafor Farrant and Lucina Q Uddin. “Asymmetric development of dorsal and ventral attention networks in the human brain”. en. In: *Dev. Cogn. Neurosci.* 12 (Apr. 2015), pp. 165–174.
- [72] Luigi Rolando. *Saggio sopra la vera struttura del cervello dell’uomo e degl’animali e sopra le funzioni del sistema nervoso di Luigi Rolando.*. 2. nella stamperia da SSRM privilegiata, 1809.
- [73] Jia-Hong Gao et al. “Cerebellum implicated in sensory acquisition and discrimination rather than motor control”. In: *Science* 272.5261 (1996), pp. 545–547.
- [74] Richard B Ivry and Juliana V Baldo. “Is the cerebellum involved in learning and cognition?” In: *Current opinion in neurobiology* 2.2 (1992), pp. 212–216.
- [75] Mario Manto et al. “Consensus paper: roles of the cerebellum in motor control—the diversity of ideas on cerebellar involvement in movement”. en. In: *Cerebellum* 11.2 (June 2012), pp. 457–487.
- [76] Catherine J Stoodley and Jeremy D Schmahmann. “Functional topography in the human cerebellum: a meta-analysis of neuroimaging studies”. In: *Neuroimage* 44.2 (2009), pp. 489–501.
- [77] Christophe Habas et al. “Distinct Cerebellar Contributions to Intrinsic Connectivity Networks”. In: *Journal of Neuroscience* 29.26 (2009), pp. 8586–8594. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1868-09.2009. eprint: <https://www.jneurosci.org/content/29/26/8586.full.pdf>. URL: <https://www.jneurosci.org/content/29/26/8586>.
- [78] Yen-Ling Chen et al. “Resting-state fMRI mapping of cerebellar functional dysconnections involving multiple large-scale networks in patients with schizophrenia”. In: *Schizophrenia Research* 149.1 (2013), pp. 26–34. ISSN: 0920-9964. DOI: <https://doi.org/10.1016/j.schres.2013.05.029>. URL: <https://www.sciencedirect.com/science/article/pii/S0920996413002843>.
- [79] D. Tomasi and N. D. Volkow. “Resting functional connectivity of language networks: characterization and reproducibility”. In: *Molecular Psychiatry* 17.8 (Aug. 2012), pp. 841–854. ISSN: 1476-5578. DOI: 10.1038/mp.2011.177. URL: <https://doi.org/10.1038/mp.2011.177>.

- [80] Lucy R. Chai et al. “Functional Network Dynamics of the Language System”. In: *Cerebral Cortex* 26.11 (Oct. 2016), pp. 4148–4159. ISSN: 1047-3211. DOI: 10.1093/cercor/bhw238. eprint: <https://academic.oup.com/cercor/article-pdf/26/11/4148/17308774/bhw238.pdf>. URL: <https://doi.org/10.1093/cercor/bhw238>.
- [81] Yaqiong Xiao et al. “Development of the intrinsic language network in preschool children from ages 3 to 5 years”. en. In: *PLoS One* 11.11 (Nov. 2016), e0165802.
- [82] Wei Gao et al. “Functional connectivity of the infant human brain: Plastic and modifiable”. en. In: *Neuroscientist* 23.2 (Apr. 2017), pp. 169–184.
- [83] Amna Rehman and Yasir Al Khalili. *Neuroanatomy, Occipital Lobe*. StatPearls Publishing, July 2023.
- [84] Takashi Hanakawa et al. “Functional properties of brain areas associated with motor execution and imagery”. en. In: *J. Neurophysiol.* 89.2 (Feb. 2003), pp. 989–1002.
- [85] Angela R Laird et al. “Behavioral interpretations of intrinsic connectivity networks”. en. In: *J. Cogn. Neurosci.* 23.12 (Dec. 2011), pp. 4022–4037.
- [86] K J Friston et al. “Functional connectivity: the principal-component analysis of large (PET) data sets”. en. In: *J. Cereb. Blood Flow Metab.* 13.1 (Jan. 1993), pp. 5–14.
- [87] R Matthew Hutchison et al. “Dynamic functional connectivity: promise, issues, and interpretations”. en. In: *Neuroimage* 80 (Oct. 2013), pp. 360–378.
- [88] Timothy O. Laumann et al. “On the Stability of BOLD fMRI Correlations”. In: *Cerebral Cortex* 27.10 (Sept. 2016), pp. 4719–4732. ISSN: 1047-3211. DOI: 10.1093/cercor/bhw265. eprint: <https://academic.oup.com/cercor/article-pdf/27/10/4719/19832217/bhw265.pdf>. URL: <https://doi.org/10.1093/cercor/bhw265>.
- [89] Yasaman Shahhosseini and Michelle F Miranda. “Functional connectivity methods and their applications in fMRI data”. en. In: *Entropy (Basel)* 24.3 (Mar. 2022), p. 390.
- [90] Seong-Gi Kim, Wolfgang Richter, and Kāmil Uğurbil. “Limitations of temporal resolution in functional MRI”. In: *Magnetic resonance in medicine* 37.4 (1997), pp. 631–636.
- [91] Ravi S Menon and Bradley G Goodyear. “Spatial and temporal resolution in fMRI”. In: *Functional MRI: an introduction to methods* (2001), pp. 145–158.

- [92] April R Levin et al. “BEAPP: The batch electroencephalography Automated Processing Platform”. en. In: *Front. Neurosci.* 12 (Aug. 2018), p. 513.
- [93] Laurel J Gabard-Durnam et al. “The Harvard automated processing pipeline for electroencephalography (HAPPE): Standardized processing software for developmental and high-artifact data”. In: *Front. Neurosci.* 12 (Feb. 2018).
- [94] Susan Whitfield-Gabrieli and Alfonso Nieto-Castanon. “Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks”. en. In: *Brain Connect.* 2.3 (July 2012), pp. 125–141.
- [95] Gesa Hartwigsen et al. “Phonological decisions require both the left and right supramarginal gyri”. en. In: *Proc. Natl. Acad. Sci. U. S. A.* 107.38 (Sept. 2010), pp. 16494–16499.
- [96] Dmitry Smirnov et al. “Fronto-parietal network supports context-dependent speech comprehension”. In: *Neuropsychologia* 63 (2014), pp. 293–303. ISSN: 0028-3932. DOI: <https://doi.org/10.1016/j.neuropsychologia.2014.09.007>. URL: <https://www.sciencedirect.com/science/article/pii/S0028393214003121>.
- [97] Anne Hsu et al. “Modulation Power and Phase Spectrum of Natural Sounds Enhance Neural Encoding Performed by Single Auditory Neurons”. In: *Journal of Neuroscience* 24.41 (2004), pp. 9201–9211. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.2449-04.2004. eprint: <https://www.jneurosci.org/content/24/41/9201.full.pdf>. URL: <https://www.jneurosci.org/content/24/41/9201>.
- [98] Danielle Smith Bassett and Ed Bullmore. “Small-world brain networks”. en. In: *Neuroscientist* 12.6 (Dec. 2006), pp. 512–523.
- [99] Danielle S Bassett and Edward T Bullmore. “Small-world brain networks revisited”. en. In: *Neuroscientist* 23.5 (Oct. 2017), pp. 499–516.
- [100] M. Oberhuber et al. “Four Functionally Distinct Regions in the Left Supramarginal Gyrus Support Word Processing”. In: *Cerebral Cortex* 26.11 (Oct. 2016), pp. 4212–4226. ISSN: 1047-3211. DOI: 10.1093/cercor/bhw251. eprint: <https://academic.oup.com/cercor/article-pdf/26/11/4212/7389755/bhw251.pdf>. URL: <https://doi.org/10.1093/cercor/bhw251>.
- [101] Weidong Cai et al. “Inhibition-related modulation of salience and frontoparietal networks predicts cognitive control ability and inattention symptoms in children with ADHD”. en. In: *Mol. Psychiatry* 26.8 (Aug. 2021), pp. 4016–4025.

- [102] Lisa Emery et al. “Age-related changes in neural activity during performance matched working memory manipulation”. en. In: *Neuroimage* 42.4 (Oct. 2008), pp. 1577–1586.
- [103] Ronnie Krupnik, Yossi Yovel, and Yaniv Assaf. “Inner Hemispheric and Interhemispheric Connectivity Balance in the Human Brain”. In: *Journal of Neuroscience* 41.40 (2021), pp. 8351–8361. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1074-21.2021. eprint: <https://www.jneurosci.org/content/41/40/8351.full.pdf>. URL: <https://www.jneurosci.org/content/41/40/8351>.
- [104] Dale Walker et al. “Prediction of school outcomes based on early language production and socioeconomic factors”. In: *Child development* 65.2 (1994), pp. 606–621.
- [105] Angela Blums et al. “Building links between early socioeconomic status, cognitive ability, and math and science achievement”. In: *Journal of Cognition and Development* 18.1 (2017), pp. 16–40.
- [106] Lucy A Lurie et al. “Mechanisms linking socioeconomic status and academic achievement in early childhood: Cognitive stimulation and language”. In: *Cognitive Development* 58 (2021), p. 101045.
- [107] Marian E Williams, Lara Sando, and Tamara Glen Soles. “Cognitive tests in early childhood: Psychometric and cultural considerations”. In: *Journal of Psychoeducational Assessment* 32.5 (2014), pp. 455–476.
- [108] Dana Charles McCoy et al. “Measuring early childhood development at a global scale: Evidence from the Caregiver-Reported Early Development Instruments”. In: *Early childhood research quarterly* 45 (2018), pp. 58–68.
- [109] James A Berkovec et al. “Mortgage discrimination and FHA loan performance”. In: *Mortgage Lending, Racial Discrimination and Federal Policy*. Routledge, 2018, pp. 289–305.
- [110] Julia Dressel and Hany Farid. “The accuracy, fairness, and limits of predicting recidivism”. In: *Science Advances* 4.1 (2018), eaao5580. DOI: 10.1126/sciadv.aao5580. eprint: <https://www.science.org/doi/pdf/10.1126/sciadv.aao5580>. URL: <https://www.science.org/doi/abs/10.1126/sciadv.aao5580>.
- [111] Joe H. Ward. “Hierarchical Grouping to Optimize an Objective Function”. eng. In: *Cluster Analysis*. London: SAGE Publications Ltd, 2012, pp. II225–. ISBN: 9780857021281.

- [112] G P McCabe, Peter H A Sneath, and Robert R Sokal. “Numerical taxonomy: The principles and practice of numerical classification”. In: *J. Am. Stat. Assoc.* 70.352 (Dec. 1975), p. 962.
- [113] Véronique Daneault et al. “Cerebral functional networks during sleep in young and older individuals”. en. In: *Sci. Rep.* 11.1 (Mar. 2021), p. 4905.
- [114] Dragutin Petkovic et al. “Improving the explainability of Random Forest classifier – user centered approach”. In: *Biocomputing 2018*, pp. 204–215. DOI: 10.1142/9789813235533_0019. eprint: https://www.worldscientific.com/doi/pdf/10.1142/9789813235533_0019. URL: https://www.worldscientific.com/doi/abs/10.1142/9789813235533_0019.
- [115] Edward F Chang et al. “Categorical speech representation in human superior temporal gyrus”. In: *Nature neuroscience* 13.11 (2010), pp. 1428–1432.
- [116] Phil Reed et al. “Speech perception in rats: use of duration and rise time cues in labeling of affricate/fricative sounds”. In: *Journal of the experimental analysis of behavior* 80.2 (2003), pp. 205–215.
- [117] Crystal T Engineer et al. “Cortical activity patterns predict speech discrimination ability”. In: *Nature neuroscience* 11.5 (2008), pp. 603–608.
- [118] Sharon L Sally and Jack B Kelly. “Organization of auditory cortex in the albino rat: sound frequency”. In: *Journal of neurophysiology* 59.5 (1988), pp. 1627–1638.
- [119] Vyassa L. Baratham et al. “Columnar Localization and Laminar Origin of Cortical Surface Electrical Potentials”. In: *Journal of Neuroscience* 42.18 (2022), pp. 3733–3748. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1787-21.2022. eprint: <https://www.jneurosci.org/content/42/18/3733.full.pdf>. URL: <https://www.jneurosci.org/content/42/18/3733>.
- [120] Michael P Kilgard and Michael M Merzenich. “Distributed representation of spectral and temporal information in rat primary auditory cortex”. In: *Hearing research* 134.1-2 (1999), pp. 16–28.
- [121] Daniel B Polley et al. “Multiparametric auditory receptive field organization across five cortical fields in the albino rat”. In: *Journal of neurophysiology* 97.5 (2007), pp. 3621–3638.
- [122] Michael M Merzenich, Paul L Knight, and G Linn Roth. “Representation of cochlea within primary auditory cortex in the cat”. In: *Journal of neurophysiology* 38.2 (1975), pp. 231–249.
- [123] Jennifer K Bizley et al. “Functional organization of ferret auditory cortex”. In: *Cerebral cortex* 15.10 (2005), pp. 1637–1653.

- [124] Melissa Saenz and Dave RM Langers. “Tonotopic mapping of human auditory cortex”. In: *Hearing research* 307 (2014), pp. 42–52.
- [125] Israel Nelken. “Processing of complex stimuli and natural scenes in the auditory cortex”. In: *Current opinion in neurobiology* 14.4 (2004), pp. 474–480.
- [126] Alexander JE Kell et al. “A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy”. In: *Neuron* 98.3 (2018), pp. 630–644.
- [127] Andrew J King and Israel Nelken. “Unraveling the principles of auditory cortical processing: can we learn from the visual system?” In: *Nature neuroscience* 12.6 (2009), pp. 698–701.
- [128] Christopher R Holdgraf et al. “Rapid tuning shifts in human auditory cortex enhance speech intelligibility”. In: *Nature communications* 7.1 (2016), p. 13654.
- [129] J. S. Garofolo et al. *DARPA TIMIT Acoustic Phonetic Continuous Speech Corpus CDROM*. 1993.
- [130] Nobuaki Takahashi, Makio Kashino, and Naoyuki Hironaka. “Structure of rat ultrasonic vocalizations and its relevance to behavior”. en. In: *PLoS One* 5.11 (Nov. 2010), e14115.
- [131] Svetlana Gogoleva et al. “To bark or not to bark? Vocalization in red foxes selected for tameness or aggressiveness toward humans”. In: *Bioacoustics The International Journal of Animal Sound and its Recording* 18 (July 2008). DOI: 10.1080/09524622.2008.9753595.
- [132] Svetlana Gogoleva et al. “Sign and strength of emotional arousal: Vocal correlates of positive and negative attitudes to humans in silver foxes (*Vulpes vulpes*)”. In: *Behaviour* 147 (Dec. 2010). DOI: 10.2307/25799782.
- [133] Lee Alan Dugatkin. “The silver fox domestication experiment”. en. In: *Evolution (N. Y.)* 11.1 (Dec. 2018).
- [134] Theunissen, Frederic. *BioSound*. URL: <https://github.com/theunissenlab/soundsig>.
- [135] B P Bogert, M J R Healy, and J W Tukey. “The Quefrency Alanalysis [sic] of Time Series for Echoes: Cepstrum, Pseudo Autocovariance, Cross-Cepstrum and Saphe Cracking”. In: *Proceedings of the Symposium on Time Series Analysis (M. Rosenblatt*. Vol. 15. New York: Wiley, 1963, pp. 209–243.

- [136] Leland McInnes, John Healy, and James Melville. *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. 2020. arXiv: 1802.03426 [stat.ML].
- [137] Corinna Cortes and Vladimir Vapnik. “Support-vector networks”. en. In: *Mach. Learn.* 20.3 (Sept. 1995), pp. 273–297.
- [138] Katarzyna J. Blinowska and Maciej Kaminski. “Functional Brain Networks: Random, “Small World” or Deterministic?” English. In: *PLoS One* 8.10 (Oct. 2013). Copyright - © 2013 Blinowska, Kaminski. This is an open-access article distributed under the terms of the Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/> (the “License”), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. Notwithstanding the ProQuest Terms and Conditions, you may use this content in accordance with the terms of the License; Last updated - 2023-12-01; SubjectsTermNotLitGenreText - Poland. URL: <https://www.proquest.com/scholarly-journals/functional-brain-networks-random-small-world/docview/1447630384/se-2>.