

Computational Methods for Analyses of Functional Brain Connectomes

Vikram Ravindra and Ananth Grama

Department of Computer Science
Purdue University

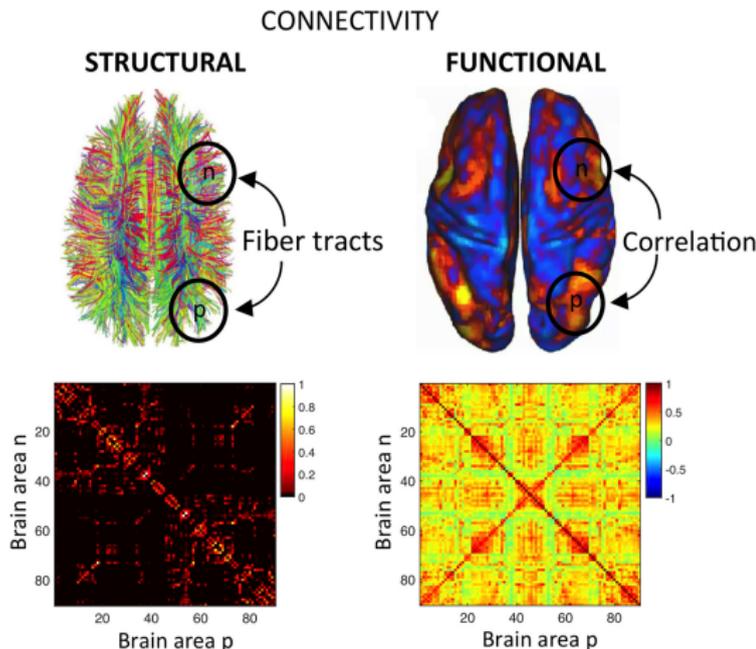
Slides contain results from joint work with Chih-Hao Fang, Prof. David F. Gleich, and Prof. Petros Drineas

- 1 **Introduction to Connectomics**
- 2 **Understanding Visual Stimulus Processing**
 - Finding Signatures in Connectomes
 - Predicting Canonical Response
 - Reconstructing Visual Stimulus
- 3 **References and Related Work**

- 1 Introduction to Connectomics
- 2 Understanding Visual Stimulus Processing
- 3 References and Related Work

Preliminaries: What is a connectome?

- ▶ Connectome is a map of the Brain.
- ▶ Connectomes capture structural or functional connectivity.

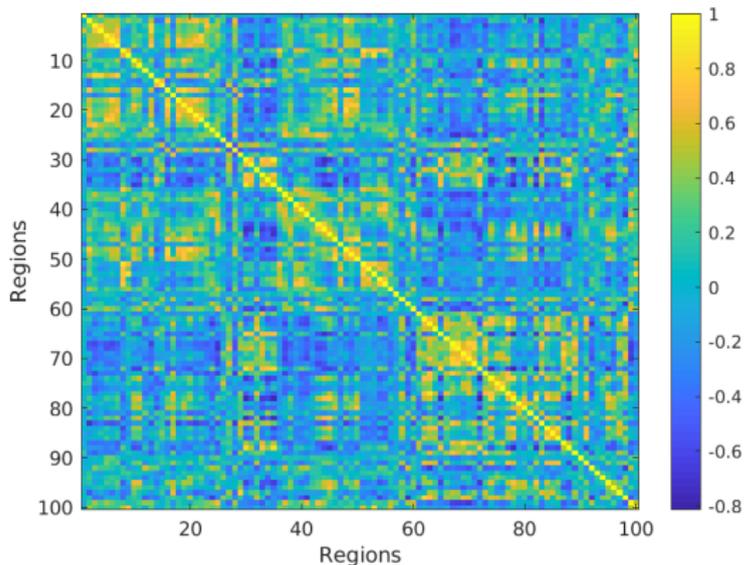


1

¹[Cabral et al] 10.1016/j.neuroimage.2017.03.045

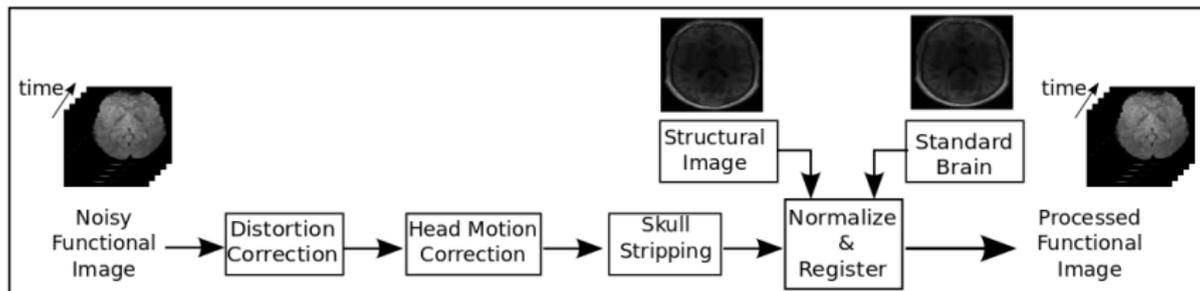
Preliminaries: What is a connectome?

- ▶ Structural connections are often called “wiring diagrams” of the brain
- ▶ Functional connections measure coherence in neuronal firing at rest or while performing a task
- ▶ Our work is focused on functional connectomes



Functional connectome as a *regions* \times *regions* similarity matrix

Preliminaries: A note on pre-processing



An overview of a generic functional MRI preprocessing pipeline

- 1 Introduction to Connectomics
- 2 **Understanding Visual Stimulus Processing**
 - Finding Signatures in Connectomes
 - Predicting Canonical Response
 - Reconstructing Visual Stimulus
- 3 References and Related Work

- ▶ Brain signatures are signals in the connectome that code for individual-level uniqueness
- ▶ Brain signatures can be thought of as **fingerprints**
- ▶ Do brain signatures exist?

OPEN ACCESS Freely available online

PLOS ONE

Connectotyping: Model Based Fingerprinting of the Functional Connectome

Oscar Miranda-Dominguez¹, Brian D. Mills¹, Samuel D. Carpenter¹, Kathleen A. Grant^{1,2}, Christopher D. Kroenke^{1,2,3}, Joel T. Nigg^{4,1}, Damien A. Fair^{1,2,4*}

1Department of Behavioral Neuroscience, Oregon Health & Science University, Portland, Oregon, United States of America, 2Advanced Imaging Research Center, Oregon Health and Science University, Portland, Oregon, United States of America, 3Division of Neuroscience, Oregon National Primate Research Center, Beaverton, Oregon, United States of America, 4Department of Psychiatry, Oregon Health & Science University, Portland, Oregon, United States of America

Abstract

A better characterization of how an individual's brain is functionally organized will likely bring dramatic advances to many fields of study. Here we show a model-based approach toward characterizing resting state functional connectivity MRI (rsfMRI) that is capable of identifying a so-called "connectotype", or functional fingerprint in individual participants. The approach rests on a simple linear model that proposes the activity of a given brain region can be described by the weighted sum of its functional neighboring regions. The resulting coefficients correspond to a personalized model based connectivity matrix that is capable of predicting the time-series of each subject. Importantly, the model itself is subject specific, and has the ability to predict an individual at a later date using a limited number of non-sequential frames. While we show that there is a significant amount of shared variance between models across subjects, the model's ability to discriminate an individual is driven by unique connectives in higher order control regions in frontal and parietal cortices. Furthermore, we show that the connectotype is present in non-human primates as well, highlighting the translational potential of the approach.

ARTICLES

nature
neuroscience

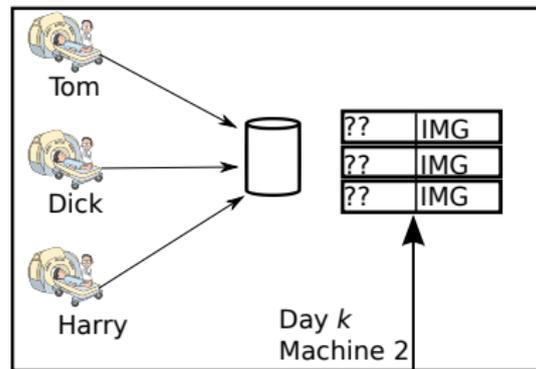
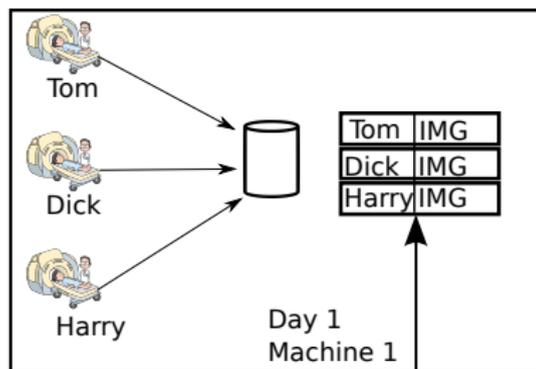
Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity

Emily S Finn^{1,2}, Xilin Shen^{2,3}, Dustin Scheinost², Monica D Rosenberg², Jessica Huang², Marvin M Chun^{1,3,4}, Xenophon Papademetris^{2,5} & R Todd Constable^{1,2,6}

Functional magnetic resonance imaging (fMRI) studies typically collapse data from many subjects, but brain functional organization varies between individuals. Here we establish that this individual variability is both robust and reliable, using data from the Human Connectome Project to demonstrate that functional connectivity profiles act as a "fingerprint" that can accurately identify subjects from a large group. Identification was successful across scan sessions and even between task and rest conditions, indicating that an individual's connectivity profile is intrinsic, and can be used to distinguish that individual regardless of how the brain is engaged during imaging. Characteristic connectivity patterns were distributed throughout the brain, but the frontoparietal network emerged as most distinctive. Furthermore, we show that connectivity profiles predict levels of fluid intelligence: the same networks that were most discriminating of individuals were also most predictive of cognitive behavior. Results indicate the potential to draw inferences about single subjects on the basis of functional connectivity MRI.

PURDUE
UNIVERSITY

Problem Setup



Can we find
corresponding
entries?

- ▶ Can we find connectomic markers for identifying individuals using their connectomes?
- ▶ Can we find markers that identify tasks being performed by subjects?
- ▶ Can we characterize how well people perform these tasks?
- ▶ Can we find brain signatures that are encoded **exclusively** in the functional connectome? (as opposed to the structural connectome)

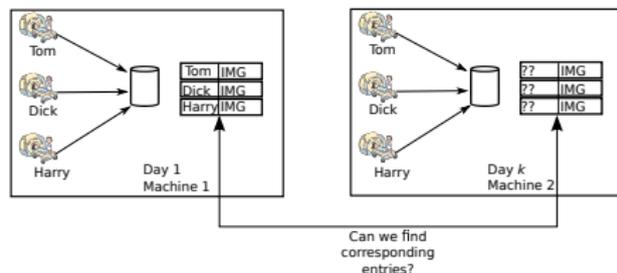
Finding Markers for Identifiability

- ▶ Existing methods do not prescribe spatially localized markers and lack anatomical basis
- ▶ We aim to find localized regions in the brain that code for identity

Our Solution

We use leverage-scores to sample a small-set of discriminating features.

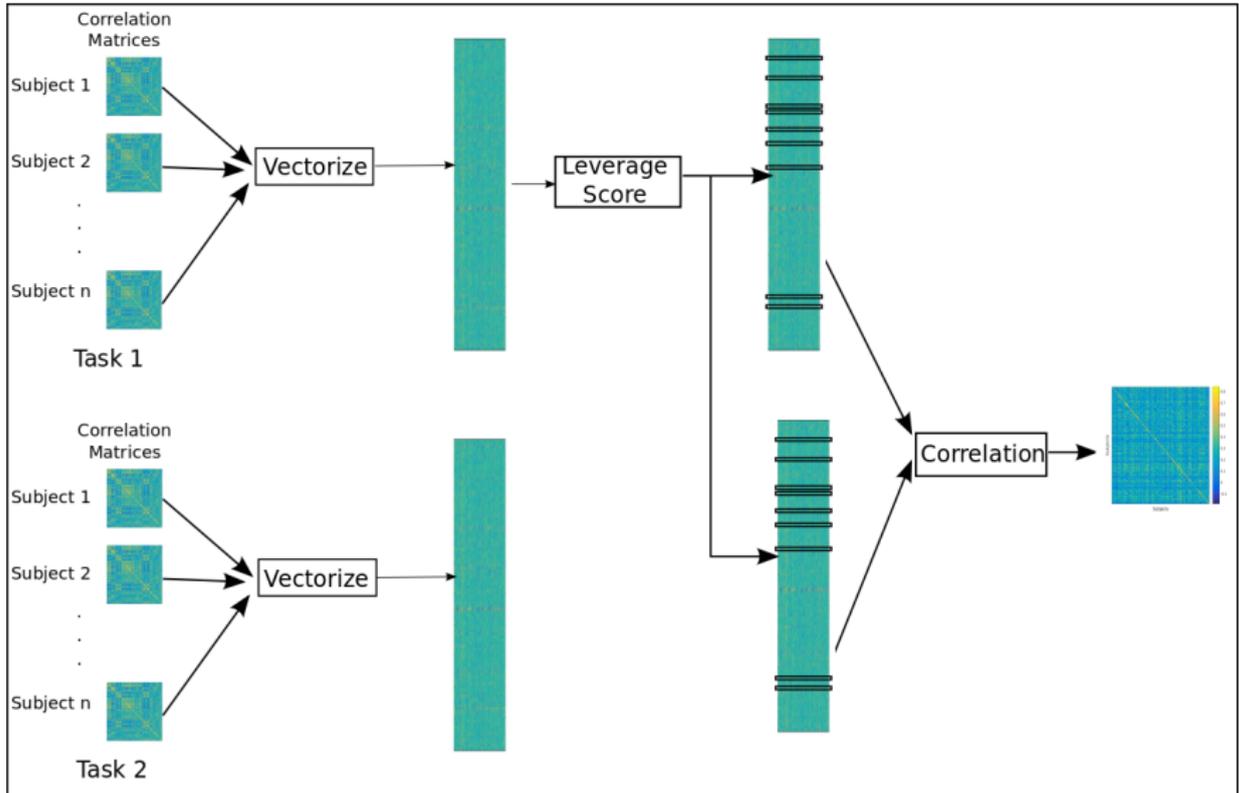
We use leverage-scores of connectomes on day 1 to select features that are strongly correlated in corresponding image on day k



Ravindra et al. SIGMOD '21

Ravindra et al. – Front. in Neuroscience '21

Our Pipeline to find individual-specific markers



Leverage Scores

Given an arbitrary $n \times d$ matrix A , with $n > d$, let U denote the $n \times d$ matrix consisting of the d left singular vectors of A , and let $U(i)$ denote the i -th row of the matrix U as a row vector. Then, the statistical leverage scores of the rows of A are given by

$$l_i = \|U_i\|_2^2 \quad (1)$$

Row Sampling using Leverage Score

```
1: function ROW_SAMPLE(A,s)
2:   Let  $\tilde{A}$  be an empty matrix
3:   for  $t = 1$  to  $s$  do
4:     Let  $A_{i_t, \star}$  be the sampled row, with corresponding probability
        $p_i \propto \|U_{i, \star}\|_2^2$ 
5:     Set  $\tilde{A}_{t, \star} = \frac{1}{\sqrt{sp_i}} A_{i_t, \star}$ 
6:   end for
7: return  $\tilde{A}$ 
8: end function
```

Sampling s rows of A using this algorithm guarantees that ²

$$\mathbb{E}[\|A^T A - \tilde{A}^T \tilde{A}\|_F] \leq \frac{1}{\sqrt{s}} \|A\|_F^2. \quad (2)$$

²[Drineas et al.]doi:10.1137/S0097539704442684

- ▶ HCP has structural, functional MRI (rest and 7 tasks), EEGs, MEGs from 1100 subjects
- ▶ We use functional MRIs from the “unrelated” subset of subjects
- ▶ The data acquisition was done over the course of 2 days
- ▶ On day 1, one resting state fMRI session followed by 4 tasks – language, emotion, gambling and motor
- ▶ On day 15, one resting state fMRI session followed by 3 tasks – relational, social and working memory

Small number of features code individual identity

- ▶ We pick increasingly large set of features with high leverage scores from the Day 1 matrix (of resting state fMRI)
- ▶ We find that a small number of features (around 60 out of 64k) from Day 1 are enough to accurately predict the corresponding identities of connectomes in day k .

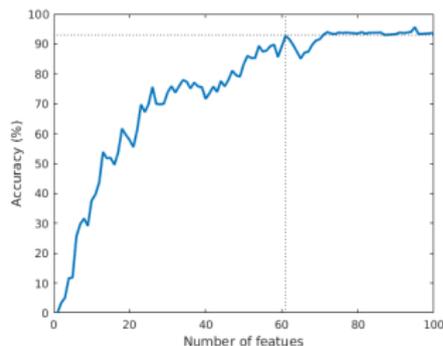


Figure: Only 60 out of 64k features are required to accurately identify connectomes in day k

Anatomical Mapping of individual-specific signatures

- ▶ Recall that each feature is an entry of the correlation matrix, i.e., it represents the similarity between two regions
- ▶ To find high confidence regions, we find high leverage score features for different randomly selected subsets of subjects.
- ▶ We retain the features which occur in a significantly large number of these subsets (p-value $< 10^{-10}$)
- ▶ Regions with high leverage scores are clustered in the parietal and pre-frontal cortex, which agrees with previous studies³ (validation)

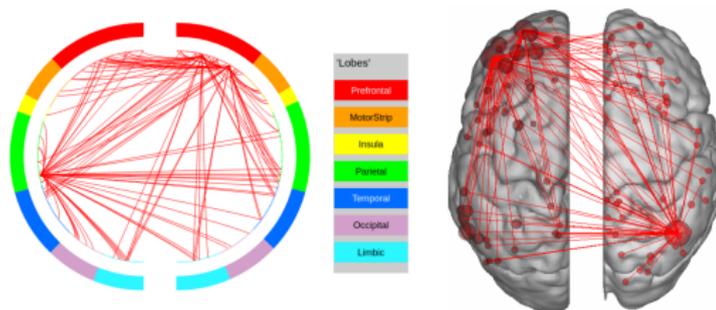


Figure: High Leverage Score Regions are Spatially Localized

Leverage-score features code for individual specific signature in task fMRI

- ▶ We repeat the same procedure with task data
- ▶ For each task, we use first half of session to find features and predict on second part of each session. This gives us a task specific signature
- ▶ We find that the prediction accuracy is generally high (apart from Motor and Working Memory)

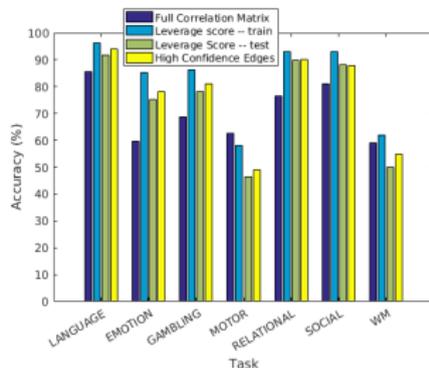


Figure: Prediction accuracies for different tasks

Our Results – Prediction Across Tasks

Markers of a task code for identity even when subjects perform a different task

- ▶ We find the markers for a task and use it to predict the identities of subjects performing a different task
- ▶ We observe that the resting-state markers are the best at predicting identities of subjects performing different tasks.

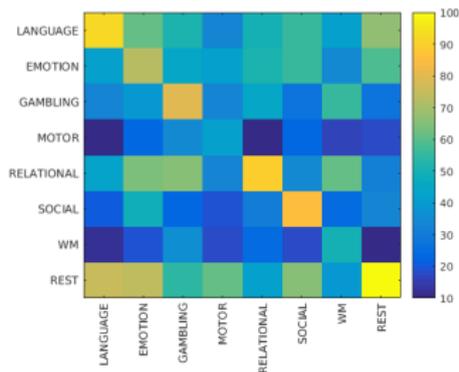


Figure: Prediction accuracies when features are chosen for task i and the test dataset has subjects performing task k

Task Performance can be predicted for fMRIs

- ▶ In the final result with leverage score, we predict the efficacy with which subjects perform tasks
- ▶ Four tasks had performance metrics associated with it
- ▶ We find the leverage-score features for the train set of each task
- ▶ We build a regression model using support vectors and predict the performance metrics of test set

Task	Train nRMSE (in %)	Test nRMSE (in %)
Language	0.33 ± 0.11	1.52 ± 0.20
Emotion	0.28 ± 0.07	0.60 ± 0.37
Relational	0.44 ± 0.04	2.74 ± 0.34
Working Memory	0.57 ± 0.12	1.93 ± 0.41

Table: Task-wise prediction error expressed as normalized root-mean-squared error.

- 1 Introduction to Connectomics
- 2 **Understanding Visual Stimulus Processing**
 - Finding Signatures in Connectomes
 - Predicting Canonical Response
 - Reconstructing Visual Stimulus
- 3 References and Related Work

- ▶ What are naturalistic stimuli? – Viewing a movie, listening to an audio book
- ▶ Inferences drawn from neuronal response to naturalistic stimuli inform us about cognitive processes
- ▶ How similar is our neuronal response to the same input?
- ▶ How similar is our neuronal activity when we recall from memory?
- ▶ How similar are our responses to the same movie, but in different languages?

We answer some of these important questions in our study.

Our Contributions: Predicting Canonical Response

- ▶ We aim to find a representation for a given fMRI frame that is consistent across subjects (we assume that all subjects are viewing the exact same stimulus at each point in time)
- ▶ We want these representations to also capture frame similarity (we want two similar frames to have similar representations)
- ▶ Our approach uses archetypal analysis.

- ▶ For a given matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, archetypal analysis finds the best minimal convex hull to describe the data-points

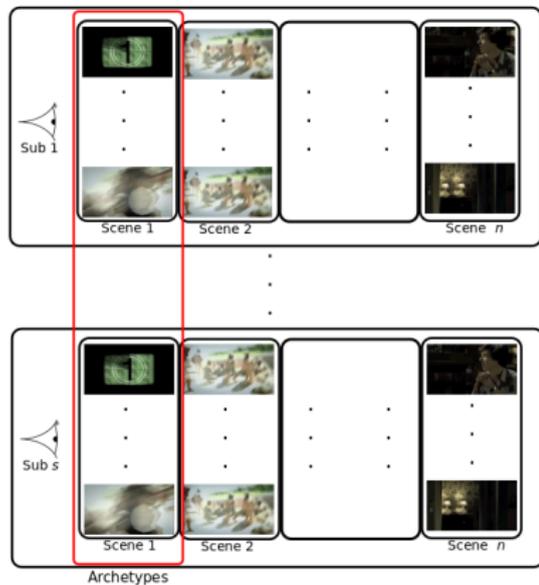
$$\begin{aligned} \min_{\mathbf{C}, \mathbf{S}} \quad & \|\mathbf{X} - \mathbf{XCS}\|_F \\ \text{s.t.} \quad & |C_{j,*}| = 1 \quad \forall j, \\ & |S_{*,i}| = 1 \quad \forall i, \\ & \mathbf{C} \geq 0 \quad \mathbf{S} \geq 0 \end{aligned} \tag{3}$$

- ▶ \mathbf{XC} represents the hull, or the *archetypes*
- ▶ Column j of \mathbf{S} represents the fractions of all archetypes which explains the data-point j .

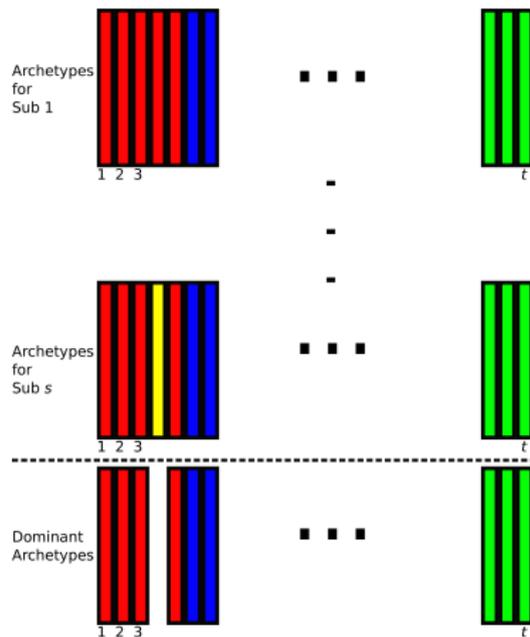
Our Approach

- ▶ Divide the episode into scenes (this is usually done by experts). Each scene corresponds to a few tens of fMRI frames
- ▶ Stack all fMRI frames for a scene from all subjects into a matrix, and do archetyal analysis.
- ▶ Archetypes are representative snap-shots of brains corresponding to the scene.
- ▶ For each fMRI frame of each subject, find the closest archetype
- ▶ Find the consensus archetype for each time-frame. We call this the *Dominant Archetype* for the frame.
- ▶ The set of dominant archetypes corresponding to each time-frame forms our canonical representation of the scene

Our Approach



(a)



(b)

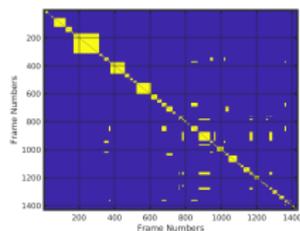
- ▶ We use a dataset⁴ with subjects viewing an episode of BBC's Sherlock
- ▶ 17 subjects, approximately 50 minutes viewing, 2000 fMRI frames and 50 scenes (expert annotation)
- ▶ We use ROIs from Hippocampus and the dorsal Default Mode Network (dDMN)
- ▶ Hippocampus has been implicated in episodic memory and dDMN is active for viewing stimuli



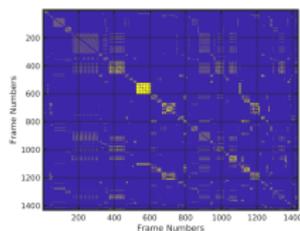
⁴[Chen et al.] doi:10.1038/n.4450

Our Results – Frame Similarity

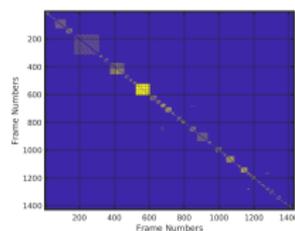
- ▶ We aim to validate that frames with similar/same archetypes are actually similar. (e.g., they share annotations)
- ▶ We find the pairwise archetypal similarity of frames
- ▶ We calculate pairwise annotation similarity (rated by experts, on the basis of excitement, emotion, indoor/outdoor, etc)
- ▶ Note the similarity in block diagonal structure. We find that the correlation between the values in this block diagonal is high (> 0.9) whereas correlation elsewhere is low (< 0.2).



(c) Framewise similarity of dominant archetypes (i.e., from fMRI space)



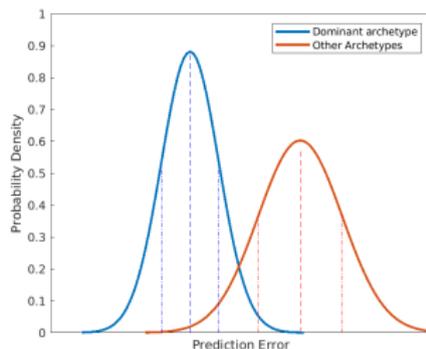
(d) Framewise similarity of expert annotations



(e) Framewise similarity of archetypes and annotations

Our Results – Prediction of response

- ▶ We hold out 30% subjects and train on the remaining. We use the framewise dominant archetypes of the train-set to measure how well we can predict the frame-wise response on test set
- ▶ Our results show that 77% of the test-set frames are closer to their corresponding dominant archetypes, as opposed to any other archetypes.
- ▶ Moreover, the nRMSE is 20% with dominant archetypes, and on an average is 45% with other archetypes



Our Results – Stable representation of neuronal response

- ▶ For each scene, we stack the fMRI matrices of all subjects into a population matrix.
- ▶ AA on all scene-wise population matrices gives us stable representation of each frame for a population – 93.32% of all time-frames in every epoch were consistently closer to the same archetype across subjects.

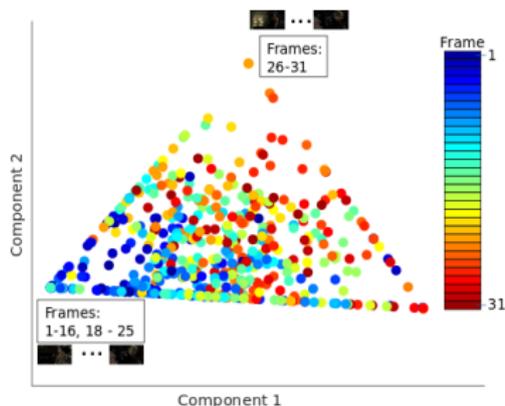


Figure: Scene-wise archetypes across subjects for scene 50.

Conclusion

- ▶ We show a new approach to deconvolving a population of functional images, which could lead to novel insights
- ▶ In particular, we show that AA gives a stable, frame-wise representation of brain activity
- ▶ We show that these representations correlate strongly with frame-wise annotation. They can also predict response in new subjects who are shown the same stimulus.

- 1 Introduction to Connectomics
- 2 **Understanding Visual Stimulus Processing**
 - Finding Signatures in Connectomes
 - Predicting Canonical Response
 - **Reconstructing Visual Stimulus**
- 3 References and Related Work

- ▶ Previously, we observed that complex, dynamic neuronal response to continuous visual inputs can be expressed by low-rank representations.
- ▶ Now, we show that we can predict visual objects (such as face, car, etc.).
- ▶ Further, we reconstruct individual video frames from fMRI.
- ▶ Our approach uses deep-learning architectures

Technical Approach

- ▶ Divide video (and corresponding fMRI) inputs into train and test.
- ▶ We train an encoder-decoder neural network on the training video clips (using cross-validation).
- ▶ We use the trained encoder to convert video frames into latent representations.
- ▶ We train a map from the fMRI response to latent vectors corresponding to the same visual input frames
- ▶ We reconstruct video frames from fMRI using previously trained map and decoder.

Technical Approach

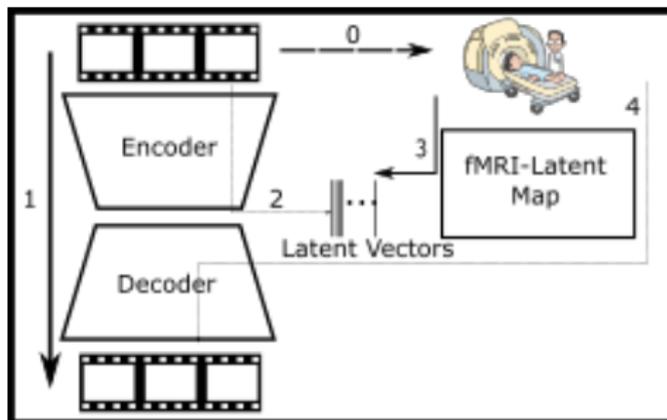


Figure: Schematic representation of our setup. The solid arrows correspond to training steps.

Encoder-Decoder Architecture

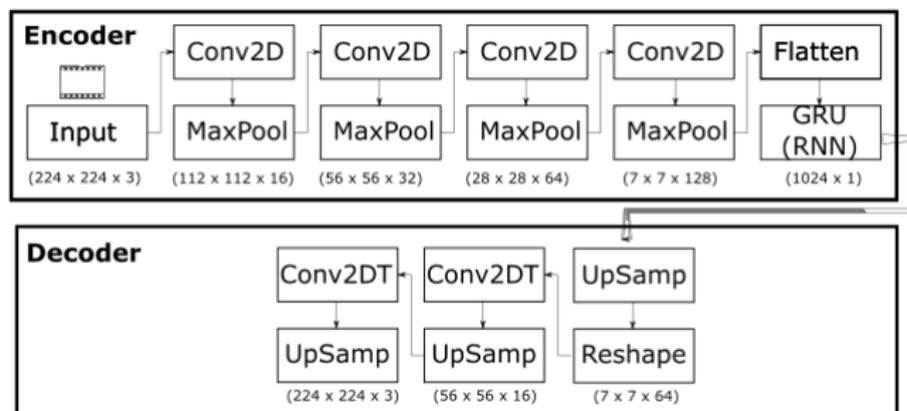


Figure: Each input frame is of dimensions $224 \times 224 \times 3$. The encoder is a sequence of convoluted neural networks (CNNs) and pooling layers, with increasing number of channels and reduction in spatial dimensions. Finally, we flatten out the tensor and input into a Gated Recurrent Unit (GRU). The decoder is a sequence of upsampling and CNN transpose. The output of the final layer of decoder is a video frame with dimensions that match the input.



Figure: We have 2 layers of GRU with 1024 dimensions and tanh activation. Finally, this is followed by a linear dense layer.

Similarity in Cluster Structure

- ▶ We show that the cluster structure of predicted latent vectors computed using fMRIs is similar to latent vectors computed from video stimulus.
- ▶ This shows that the output of the map (that inputs fMRI) is a good approximation to output of the encoder (that inputs video frames), thereby allowing us to make predictions about video stimuli

Similarity in Cluster Structure

- ▶ We train the encoder-decoder and fMRI-Latent Map as described earlier. This gives us the following:
 - ▶ Latent vectors corresponding to train set within each iteration of cross-validation (which both the encoder and non-linear map have previously seen).
 - ▶ Latent vectors in the hold-out set in the training procedure of the non-linear map (which the encoder has seen, but the map has not seen in the current iteration).
 - ▶ Latent vectors corresponding to the test video and corresponding test fMRI (which neither the encoder nor the map have previously encountered).
- ▶ In each case, we cluster the outputs of encoder and non-linear map separately using k-means for $k = \{3, 4, \dots, 20\}$.
- ▶ Then, we compute the adjusted rand index (ARI) between the clustering obtained.

Result – Similarity in Cluster Structure

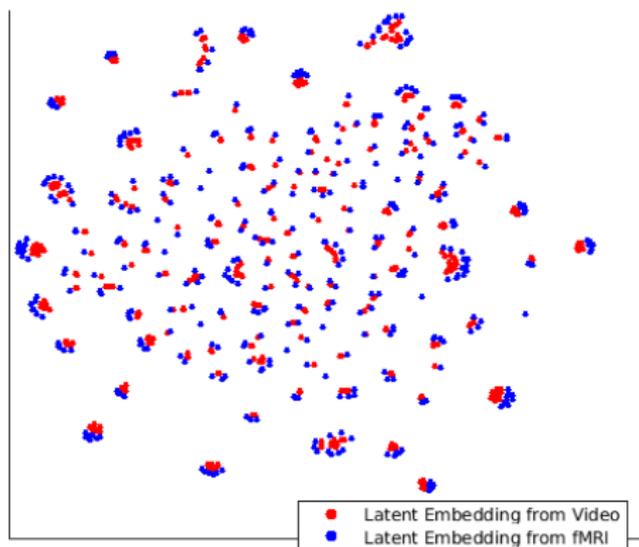
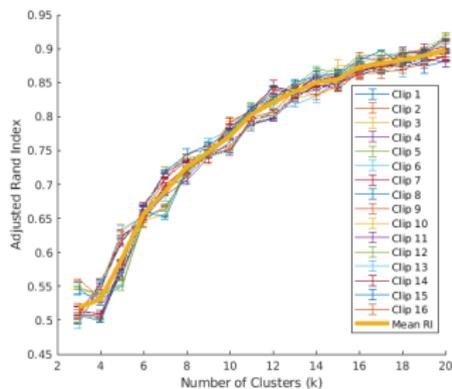
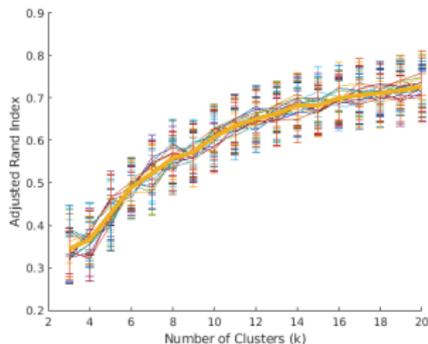


Figure: Visualization of latent representations computed from video frames processed using a trained encoder (blue) and from corresponding fMRI frames processed using a trained non-linear map (red) for a sample subject, after reducing to two dimensions using t-SNE.

Result – Similarity in Cluster Structure



(a)



(b)

Figure: Plots show the relationship between cluster groupings and adjusted rand-index (ARI) when (a) the encoder has been previously trained on the video frames; (b) neither neural network has previously seen the video or corresponding fMRI frame.

Result – Similarity in Cluster Structure

- ▶ The high values of both measures show that embeddings obtained from the two routes are similar.
- ▶ Therefore, fMRI processed using the trained map can be used to approximate the video frames and encoder.
- ▶ We note that the measures plateau for larger values of k .

Prediction of Visual Features

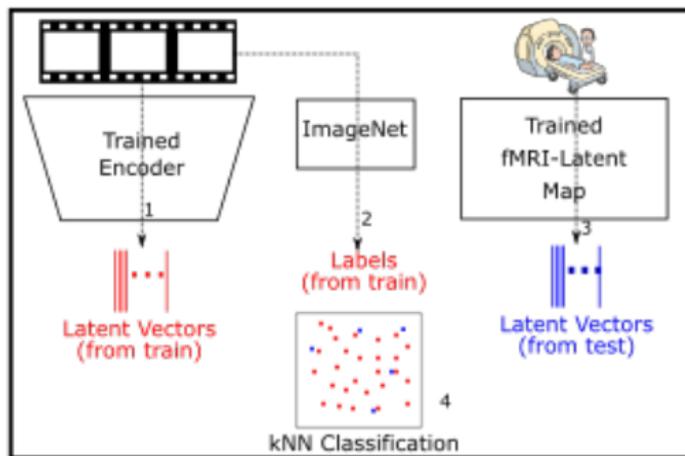


Figure: Prediction of visual features from latent space. We assume that the encoder-decoder and the fMRI-latent map are previously trained. First, using the train video frames, we obtain the (actual) latent vectors. Then, we input the same video frames into pre-trained ImageNet to classify each of the frames. Following this, we predict latent vectors by inputting test fMRI through a pretrained map. We then do k-NN classification between actual and predicted latent vectors to predict visual features in fMRI frames.

Result – Prediction of Visual Features

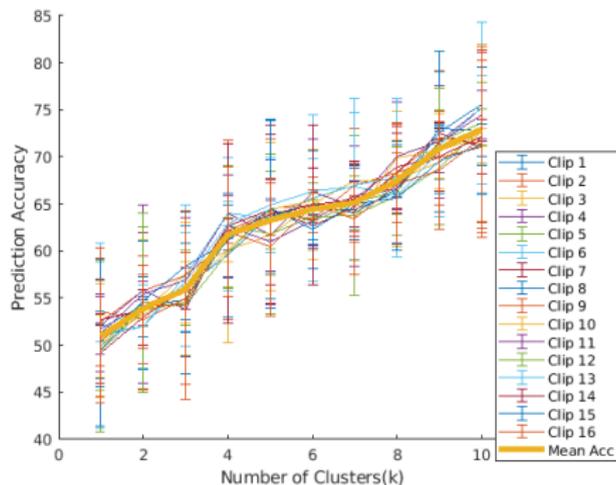


Figure: Prediction accuracy of objects in visual stimulus. This figure shows the prediction accuracy when matching objects/ classes recognized in visual stimulus (in latent space) and its corresponding k -closest neighbors among all train video frames (in latent space). We see that the prediction accuracy is $> 70\%$ when we allow the fMRI frame be matched to one of its ten closest euclidean neighbours in latent space. We use ImageNet to recognize entities in images.

Result – Prediction of Visual Features

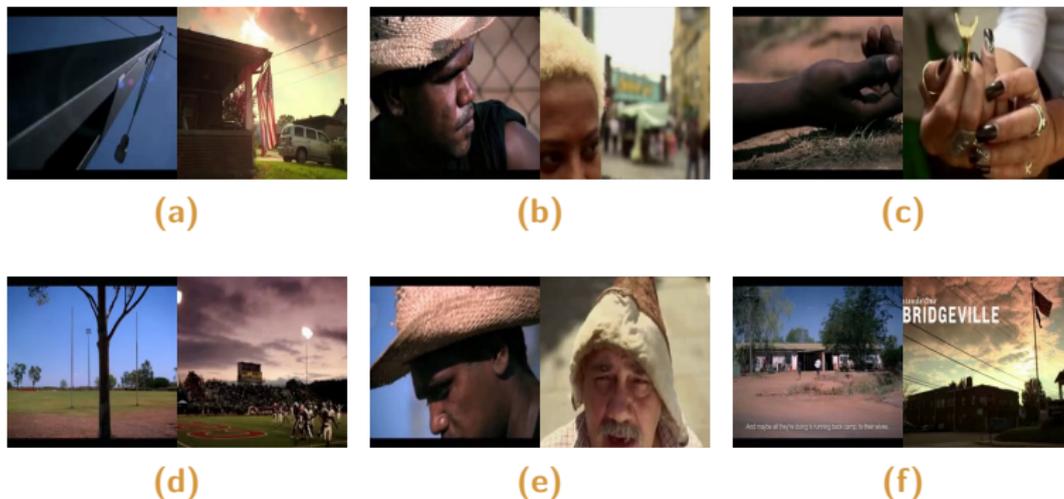


Figure: Examples of visual stimulus reconstructed from latent representation. The left frame in each of these examples is computed using latent vectors from visual stimulus, whereas the right frame is the its closest euclidean neighbour among latent vectors computed from fMRI frames in the training set. In each of the images, we can see similarity in objects/ faces in the background or foreground.

Face Detection

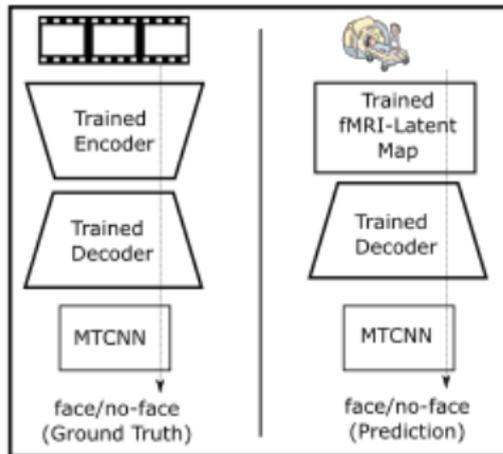


Figure: Schematic Representation of Face Detection. After training the encoder-decoder and the fMRI-Latent map, we detect faces as follows. First, test videos are passed through the encoder-decoder, and the output of the decoder is processed using MTCNN, a pretrained network to detect faces. This output serves as the ground truth. Then, we pass the corresponding fMRI frames into the map and the resulting latent vector through the decoder. The output of this is fed into MTCNN to give our prediction of face/ no-face. Comparing the labels of ground truth and our prediction, we assess the performance of our framework.

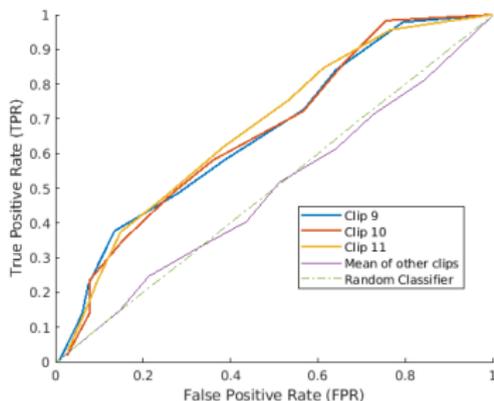


Figure: Receiver Operating Characteristic (ROC) obtained by varying confidence thresholds for the three stages of MTCNN. We note excellent performance of our framework for scenes with significant number of faces.

Result – Face Reconstruction



Figure: Examples of reconstructed faces. In each of these cases, the image of the left is the original image, and the image on the right is reconstructed from corresponding fMRI frames. We input test fMRI frames into the non-linear map to obtain predicted latent vectors. We input these vectors into the previously trained decoder.

References

- ▶ V Ravindra, H Nassar, DF Gleich, and A Grama, Rigid Graph Alignment, ComplexNets, 2019.
- ▶ V Ravindra and A Grama, Characterizing Similarity of Visual Stimulus from Associated Neuronal Response, IJCAI 2020
- ▶ V Ravindra and A Grama, De-anonymization attacks in Neuroimaging Datasets, SIGMOD 2021.
- ▶ V Ravindra, P Drineas, and A Grama, Constructing Compact Brain Connectomes for Individual Fingerprinting, Frontiers in Neuroscience 2021.
- ▶ V Ravindra, G Sanders, and A Grama, Identifying Coherent Subgraphs in Dynamic Brain Networks, ICIP 2021.
- ▶ V Ravindra, H Nassar, DF Gleich, and A Grama, Aligning Spatially Constrained Graphs, IEEE Tran. on Knowledge and Data Engg., 2022
- ▶ (In review) V Ravindra, CH Fang, and A Grama, May I see what you see? Predicting Visual Features from Neuronal Activity, 2022

Rigid Graphs and Alignment

- ▶ Graphs are commonly used to model real-world systems
- ▶ Many of these systems have fixed relative positions for nodes – we call such graphs “rigid graphs”

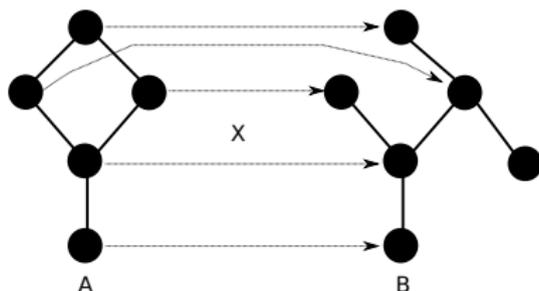


Figure: Examples of Rigid Graphs

Ravindra et al. – ComplexNets '19
Ravindra et al. – In Review, '21

Topological Network Alignment

- ▶ Let $A(V_A, E_A)$ and $B(V_B, E_B)$ be two graphs
- ▶ The problem of network alignment is to find correspondence between nodes of A and B
- ▶ Example: $\mathbf{X} \in \{0, 1\}^{|V_A| \times |V_B|}$ encodes the correspondence between A and B



$$\mathbf{x}_{ij} = \begin{cases} 1 & \text{if } i \in V_A \text{ is matched with } j \in V_B \\ 0 & \text{otherwise} \end{cases}$$

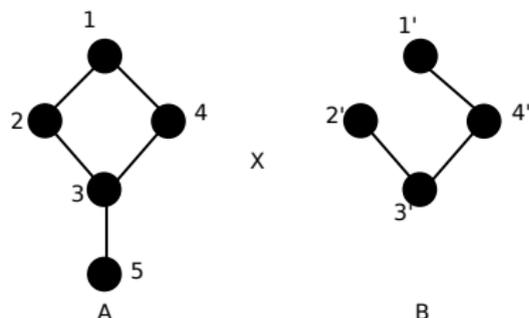
Topological Network Alignment: Formulation

- ▶ Recall: A and B are the adjacency matrices and \mathbf{X} encodes the correspondence
- ▶ Let $\mathbf{L} \in \mathbb{R}^{|V_A| \times |V_B|}$ denote the prior matrix that provides external information about the similarity between pairs of nodes in A and B
- ▶ Let α and β denote the relative importance of prior knowledge and edge-overlap

$$\begin{aligned} \max_{\mathbf{X}} \quad & \alpha \mathbf{L} \bullet \mathbf{X} + \beta A \bullet \mathbf{X} B \mathbf{X}^T \\ \text{s.t.} \quad & \sum_i \mathbf{X}_{ij} \leq 1 \quad \forall j = 1 \dots |V_B|, \\ & \sum_j \mathbf{X}_{ij} \leq 1 \quad \forall i = 1 \dots |V_A|, \quad \mathbf{X}_{ij} \in \{0, 1\} \end{aligned} \tag{4}$$

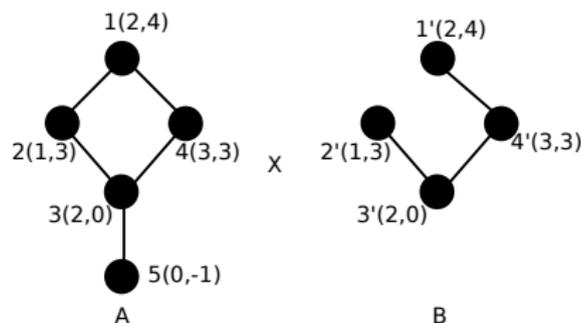
Topological Network Alignment

▶ Example



- ▶ Nodes 1,2,3,4 can be assigned to any cyclic ordering of 1',2',3',4'
- ▶ In applications where nodes represent physical entities (say, nodes represent brain regions), we need to take spatial considerations into account

Rigid Graph Alignment (RGA)



- ▶ We have $A(V_A, E_A, C_A)$ and $B(V_B, E_B, C_B)$
- ▶ Also, edges of rigid graphs have preferred (fixed) lengths. Therefore, we can perform orthogonal transformations from C_A to C_B

$$\begin{aligned} \min_{\Omega} \quad & \|C_A - C_B \Omega\|_F^2 \\ \text{s.t.} \quad & \Omega^T \Omega = I \end{aligned} \tag{5}$$

Rigid Graph Alignment – Formulation

- ▶ We can combine the objectives of network alignment and rigid body constraints
- ▶ Note that rows of C_B are permuted in accordance with \mathbf{X} (which is given by the topological aligner)

$$\mathbf{F} = \max_{\mathbf{X}, \Omega} \underbrace{\alpha \mathbf{L} \bullet \mathbf{X}}_{\text{Update Prior}} + \underbrace{\beta \mathbf{A} \bullet \mathbf{X} \mathbf{B} \mathbf{X}^T}_{\text{Network Alignment}} - \underbrace{\gamma \|\mathbf{C}_A - \mathbf{X} \mathbf{C}_B \Omega\|_F^2}_{\text{Structural Alignment}} \quad (6)$$

s.t. $\sum_i \mathbf{X}_{ij} \leq 1 \quad \forall j = 1 \dots |V_B|,$
 $\sum_j \mathbf{X}_{ij} \leq 1 \quad \forall i = 1 \dots |V_A|, \quad \mathbf{X}_{ij} \in \{0, 1\}$

- ▶ We can rewrite the objective function as follows

$$\begin{aligned}
 \mathbf{F} = \max_{\mathbf{X}} \quad & \underbrace{\alpha \mathbf{L} \bullet \mathbf{X}}_{\text{Update Prior}} + \underbrace{\beta \mathbf{A} \bullet \mathbf{X} \mathbf{B} \mathbf{X}^T}_{\text{Network Alignment}} + \underbrace{\gamma \mathbf{C}_A \boldsymbol{\Omega}^T \mathbf{C}_B^T \bullet \mathbf{X}}_{\text{Structural Alignment}} \quad (7) \\
 \text{s.t.} \quad & \sum_i \mathbf{X}_{ij} = 1 \quad \forall j = 1 \dots |V_B|, \\
 & \sum_j \mathbf{X}_{ij} = 1 \quad \forall i = 1 \dots |V_A|, \quad \mathbf{X}_{ij} \in \{0, 1\}
 \end{aligned}$$

- ▶ This form of the objective function suggests that the optimal \mathbf{X} is one that maximizes the network alignment and structural alignment, while respecting the prior.

- ▶ The RGA objective suggests that the prior \mathbf{L} should be proportional to the similarity between the coordinates, i.e., $C_A \mathbf{\Omega}^T C_B^T$.
- ▶ The prior for topological alignment must be sparse

$$\mathbf{L}_{i,j} = \begin{cases} \exp(-\|C_{A_i} - C_{B_j}\|_2^2) & \|C_{A_i} - C_{B_j}\|_2^2 \leq d_i^k \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Rigid Graph Alignment – Algorithm

- ▶ Our algorithm uses alternating procedure: we fix Ω and populate the prior, which is input to the aligner
- ▶ The aligner returns the match \mathbf{X} , which we use to refine the structural alignment

Rigid Graph Alignment

- 1: **Input:** Graphs $A(V_A, E_A, C_A)$ and $B(V_B, E_B, C_B)$, α, β, γ
- 2: **Output:** Aligned graphs A and B
- 3: **repeat**
- 4: $\mathbf{L} = \text{get_prior}(C_A, C_B)$
- 5: $\mathbf{X} = \text{align}(A, B, \mathbf{L})$
- 6: $B = \mathbf{X}B\mathbf{X}^T$
- 7: $\Omega = \text{transform_coordinates}(C_A, C_B, \mathbf{X})$
- 8: $C_B = \mathbf{X}C_B\Omega$
- 9: **until** converged

- ▶ Rigid Graph Alignment can be used in conjunction with network aligners that require priors
- ▶ The structural alignment problem, called the Orthogonal Procrustes Problem can be solved using Kabsch's algorithm

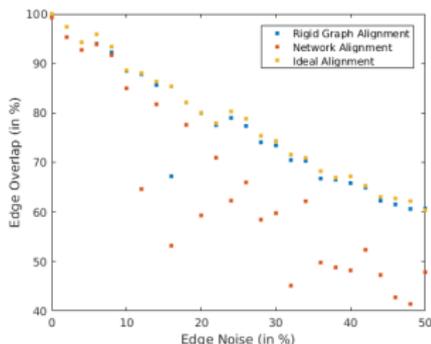
- ▶ We use both real and synthetic data
- ▶ Synthetic data is used to explore performance and robustness
- ▶ Real data is used to characterize accuracy in application settings

- ▶ We create graphs with 100,500 and 1000 nodes (results shown here are for 1000 nodes)
- ▶ We create graphs using Preferential Attachment, and Erdos-Renyi models (results shown here are for Preferential Attachment)
- ▶ **Step 1:** We create a uniform 3D grid. Each point on the grid is given a node with a biased coin-toss. Then, we add edges according to Preferential Attachment Model.
- ▶ **Step 2:** We create graph B by perturbing graph A as follows
 - ▶ We add and delete edges at random (edge noise)
 - ▶ We move the physical position of nodes (node noise)
- ▶ **Step 3:** We run Rigid Graph Alignment with our choice of aligner (results for `netalignmbp`⁵ shown here).

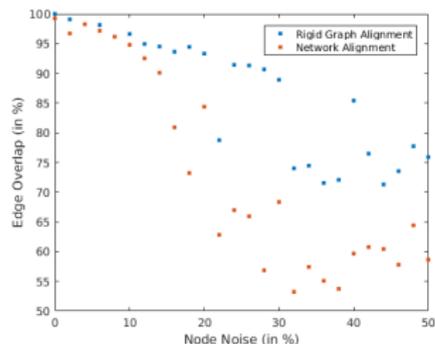
⁵[Bayati et al.] doi:10.1145/2435209.2435212

Our Results – Synthetic Experiments

- ▶ **Case 1:** We introduce noise in edges only. This decreases the number of *true* edges. RGA finds these true edges with high accuracy.
- ▶ **Case 2:** we introduce noise to nodes only. RGA corrects for these errors, whereas a wrong prior adversely affects the performance of regular topological network alignment algorithms



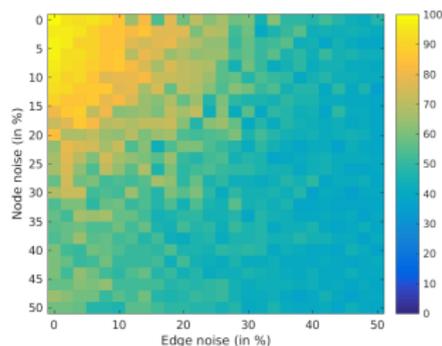
(a) Rigid Graph Alignment is robust to added/deleted edges



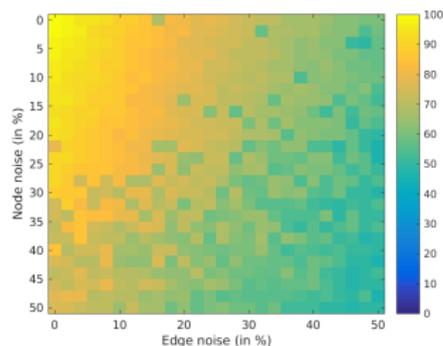
(b) Rigid Graph Alignment is robust to shift in physical positions of nodes

Our Results – Synthetic Experiments

- ▶ We introduce noise to both nodes and edges
- ▶ Rigid graph alignment is robust to noise in terms of edge overlap ($A \bullet \mathbf{X}B\mathbf{X}^T$)



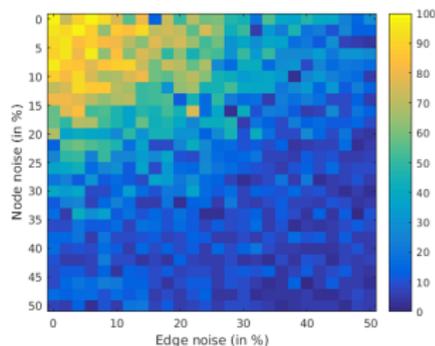
(c) Edge overlap drops with noise



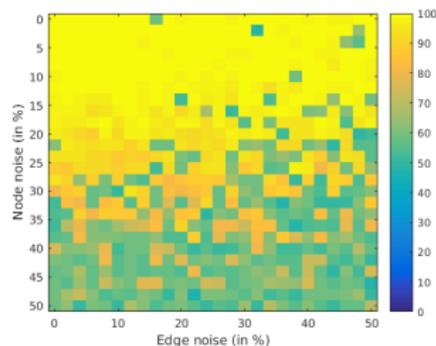
(d) Rigid Graph Alignment is robust to noise

Our Results – Synthetic Experiments

- ▶ Our graph generation procedure gives us correct node alignments, which can be used to compute node overlap
- ▶ Rigid Graph Alignment matches nodes well, even in presence of noise



(e) Node Overlap drops with noise



(f) Rigid Graph Alignment is robust to noise

- ▶ We use 2 sessions of 20 resting state fMRI from the HCP data
- ▶ For each session, we register to first image of the session (head-motion correction) and register to first image
- ▶ **We don't register to subject's own structural MRI or to a standard MNI coordinate system. This makes our analysis entirely functional.**
- ▶ Then, we generate a *voxel* \times *voxel* correlation matrix and threshold (95th percentile) to create an adjacency matrix
- ▶ Each node has a physical position (which is the 3D coordinate of the voxel)

RGA significantly improves the quality of alignment

- ▶ We aligned functional connectomes of the same subject, across sessions
- ▶ We observe a significant improvement in edge overlap in a small number of iterations
- ▶ Regular topological alignment (iteration 1 in plot)

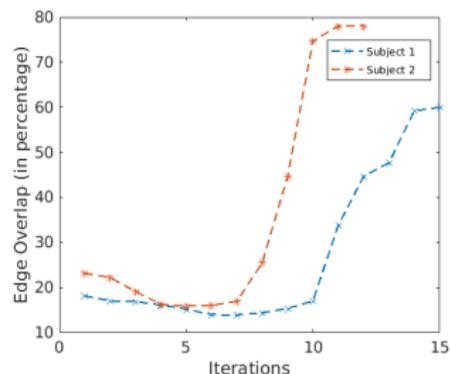
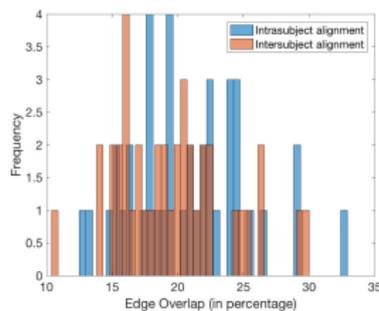


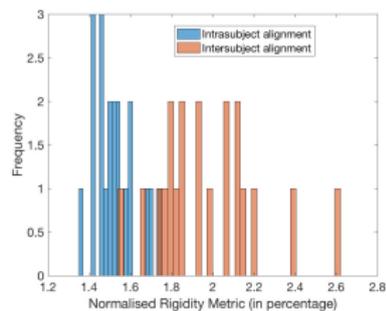
Figure: Rigid Graph Alignment improves alignment between brain networks in a small number of iterations.

Selecting suitable metrics for Identification

- ▶ We align two images of same subject (intra-subject alignment) and we align two images across subjects (inter-subject alignment)
- ▶ In terms of edge-overlap, it is difficult to distinguish between inter-subject and intra-subject alignments
- ▶ However, if we measure the error in terms of rigidity (i.e., $\|C_A - \mathbf{X}C_B\Omega\|_F$), the inter-subject and intra-subject distributions are separable.



(a) Edge Overlap is not a good metric to identify



(b) Rigidity of edge lengths is a good metric to identify

- ▶ Dynamic graphs are natural abstractions for modeling correlations in brain activity.
- ▶ An important problem in the context of these dynamic correlation graphs is the discovery of sets of regions of the brain, whose activity level is temporally coherent.
- ▶ These manifest as temporally persistent sub-graphs that are strongly connected, referred to as *coherent subgraphs*.
- ▶ In this paper, we present a model and method for identifying coherent subgraphs in dynamic correlation graphs.

- ▶ Let $\mathbf{M} \in \mathbb{R}^{r \times t}$ denote a time-series matrix, with r sources (time series) and t time-points.
- ▶ We aim to identify vertices that participate in large coherent subgraphs.
- ▶ In these nodes, internal connections are stronger than external ones.

$$\mathbf{M} = \begin{bmatrix} \mathbf{S} \\ \mathbf{Y} \end{bmatrix} = \begin{bmatrix} \mathbb{1}_s b^T + \mathbf{Z} \\ \mathbf{Y} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_s \\ \mathbf{U}_Y \end{bmatrix} \boldsymbol{\Sigma} \mathbf{V}^T \quad (9)$$

- ▶ Claim 1: The expected value of edge weights that are internal to \mathcal{C} is higher than expected value of edge weights that are external to \mathcal{C}
- ▶ Claim 2: Let $\mathbf{U}\Sigma\mathbf{V}^T$ be the rank k thin SVD of \mathbf{M} with $\sigma_k = \sigma_{min}$. The noise in the SVD coordinates associated with coherent subgraphs $|\mathcal{C}|$ is bounded.
- ▶ Claim 3: If \mathbf{M} is full rank and any spectral coordinate (row j of \mathbf{U}) is close to that of $i \in \mathcal{C}$, then j is also in \mathcal{C} .
- ▶ Proofs are in paper.

- ▶ We use fMRIs from HCP dataset. We use rest, relational memory, gambling tasks, emotional processing , language processing, and social cognition.
- ▶ Each fMRI time-series data is a *regions* (n) \times *time-points* (t) matrix \mathbf{M} .
- ▶ The session is divided into windows of length l ($=5,10,20$) and a slide size s ($=1,2,5$).
- ▶ We denote the windows as $\mathbf{W}_1, \dots, \mathbf{W}_k$, and $\mathbf{W}_i \in \mathbb{R}^{n \times l} \forall i \in \{1, \dots, k\}$.

Spectral coordinates have cluster structure similar to that of the time window matrix

- ▶ Goal: Show that cluster structure of the rows of a time window matrix \mathbf{W}_i and that of the rows of its left singular matrix \mathbf{U}_i , i.e., its *spectral coordinates*, are similar.
- ▶ We divide the time-series into windows \mathbf{W}_i and compute corresponding left-singular vector matrices \mathbf{U}_i
- ▶ We cluster the rows of both matrices separately and compute the Jaccard coefficient by retaining different number of singular vectors and varying number of clusters

		# Singular vectors			
		1	5	10	20
# clusters	3	96.26	94.25	94.31	96.22
	5	93.52	92.79	92.71	93.57
	10	88.12	85.79	86.79	87.89

Table: \mathcal{J} (Agreement) (in %) of clusters obtained by clustering time-windows (of size 20) and those obtained by clustering the left singular vectors for EMOTION task in the Human Connectome Project. This table shows that the agreement does not improve as more singular vectors are used.

Coherent subgraphs are reliable indicators of rest activity

- ▶ For each time-window matrix \mathbf{W}_i , we cluster corresponding vectors $\mathbf{U}_i^{(1)}$. We use the metadata from the HCP dataset to annotate every frame as “task” or “rest”.
- ▶ We perform k-means clustering on the rows of $\mathbf{U}_i^{(1)}$ and predicted rest correctly for $98.7 \pm 0.3\%$ of all time points.
- ▶ We use *conductance* as the metric to assess strength of the coherent subgraphs

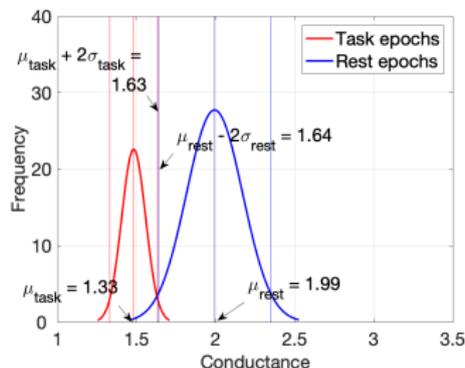


Figure: Histogram of conductance when subjects are at rest (blue) and during activity (red) during EMOTION task fit to Normal Distribution. The separation reveals that synchronous brain activity increases while performing task, which is evidenced by increase in similarity on internal edges and decrease in similarity of border edges. The blue and red vertical lines represent the means and one standard deviations of the respective distributions.

Coherent subgraphs of snapshots predict the task condition

- ▶ We annotate each time-window with its corresponding task condition such as win & loss conditions in the GAMBLING task, or face & shape conditions in EMOTION task using the metadata available with the HCP dataset.
- ▶ Then, we divide the time-points into train and test splits to perform a 10-fold cross validation. For each train/ test time-window \mathbf{W}_i , we find the top left-singular vector $\mathbf{U}_i^{(1)}$ and stack the vectors into \mathbf{L}_{train} and \mathbf{L}_{test} , respectively.
- ▶ Finally, we restrict the feature space to the clique nodes, as obtained previously. We then use linear Support Vector Machines (SVM) to predict the task conditions of test windows using the learning parameters from the train windows.

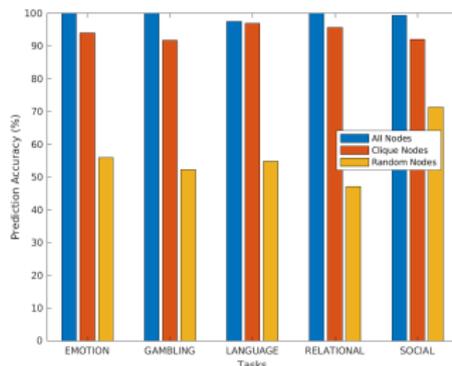


Figure: Accuracy in prediction of task-condition for 5 tasks in the HCP dataset. Our results that near-cliques encode the nature of activity in fMRIs. A random selection of nodes performs poorly in comparison.