

Extracting Reproducible Time-Resolved Resting State Networks using Dynamic Mode Decomposition

James M. Kunert-Graf^{a,*}, Kristian M. Eschenburg^b, David J. Galas^a, J. Nathan Kutz^c, Swati D. Rane^d, Bingni W. Brunton^e

^a*Pacific Northwest Research Institute*

^b*Department of Bioengineering, University of Washington*

^c*Department of Applied Math, University of Washington*

^d*Department of Radiology, University of Washington*

^e*Department of Biology, University of Washington*

Abstract

Resting state networks (RSNs) extracted from functional magnetic resonance imaging (fMRI) scans are believed to reflect the intrinsic organization and network structure of brain regions. Most traditional methods for computing RSNs typically assume these functional networks are static throughout the duration of a scan lasting 5–15 minutes. However, they are known to vary on timescales ranging from seconds to years; in addition, the dynamic properties of RSNs are affected in a wide variety of neurological disorders. Recently, there has been a proliferation of methods for characterizing RSN dynamics, yet it remains a challenge to extract reproducible time-resolved networks. In this paper, we develop a novel method based on dynamic mode decomposition (DMD) to extract networks from short windows of noisy, high-dimensional fMRI data, allowing RSNs from single scans to be resolved robustly at a temporal resolution of seconds. We demonstrate this method on data from 120 individuals from the Human Connectome Project and show that unsupervised clustering of DMD modes discovers RSNs at both the group (gDMD) and the single subject (sDMD) levels. The gDMD modes closely resemble canonical RSNs. Compared to established methods, sDMD modes capture individualized RSN structure that both better resembles the population RSN and better captures subject-level variation. We further leverage this time-resolved sDMD analysis to infer occupancy and transitions among RSNs with high reproducibility. This automated DMD-based method is a powerful tool to characterize spatial and temporal structures of RSNs in individual subjects.

1. Introduction

Resting state networks (RSNs) comprise distinct regions of the brain that exhibit synchronous low-frequency (<0.1 Hz) temporal fluctuations in the absence of explicit tasks [1, 2, 3]. They are most commonly detected using blood-oxygen level-dependent (BOLD) resting state functional magnetic resonance imaging (rs-fMRI), where a series of T2 or T2* weighted images of the brain are acquired repeatedly over the duration of the scan (5–15 minutes) [4, 5]. The data-driven extraction of RSNs from these noisy and high-dimensional datasets is a difficult analytic task, made possible through the use of well-suited methods. Independent Component Analysis (ICA) [66] has become a standard approach in this field, due to its power at detecting and disentangling overlapping, but statistically distinct, signals from noisy, high-dimensional data. The traditional application of ICA and other RSN analyses have focused on static analysis of the images, implicitly assuming that the networks are unaltered throughout the duration of each resting-state scan [5]. The assumption that brain states are static for many minutes has the effect of averaging over the spatial and temporal variability of networks [6].

Although the spatial structure of these RSN patterns is similar throughout the population, their exact structure and dynamics in time vary considerably between individuals [7, 8, 9, 10]. Further, even within individuals, the modes reconfigure dynamically on timescales ranging from seconds to years [11, 12, 13,

*Corresponding author: jkunert@pnri.org

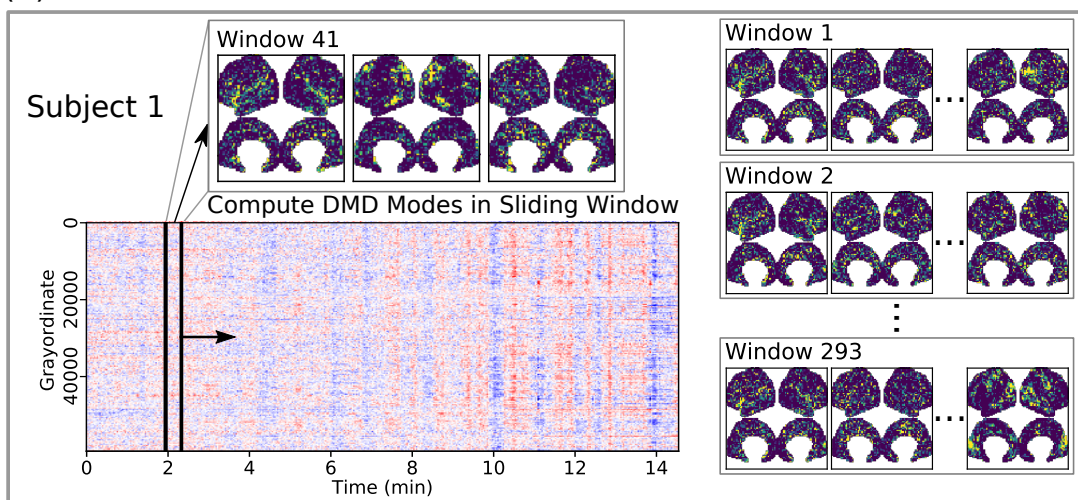
14, 15, 16]. Differences in an individual’s RSN dynamics are important in part because they may serve as useful biomarkers for a variety of neurological dysfunctions, including schizophrenia [17, 18, 19, 20, 21], bipolar disorder [20, 21], autism [22, 23], depression [24], post-traumatic stress disorder [25], attention deficit/hyperactivity disorder [26], and mild cognitive impairment [27]. For example, multiple studies have shown that in schizophrenia, regions of default mode, auditory, motor, and visual resting state networks show differences in their correlations, when compared to a control population [28, 29, 30]. Patients with schizophrenia also show abnormalities in dynamic dwell times in large-scale RSNs [31].

Technological innovations in brain imaging such as SENSitivity Encoding (SENSE) and simultaneous multi-slice (SMS) acquisitions now allow for improved temporal resolution, enabling investigation of the true temporal dynamics of brain function [32, 33, 34]. One emerging challenge with the advent of the fast acquisition of brain dynamics is the appropriate choice of methods to analyze and interpret data. It follows that much intense work has gone beyond static analyses, calculating RSN structure and activation as they vary in time to investigate a variety of measures of RSN dynamics. Many of these time-resolved RSN analyses use a windowed version of a traditional method, such as correlation [35, 36, 15, 37, 38, 6, 39] or ICA [40, 41]. Windows are then shifted in time and the measures recalculated, leading to methods for characterizing dynamic connectivity [42, 43, 35, 44]. These most common RSN extraction methods make no explicit assumptions about the intrinsic temporal structure of the fMRI data [45, 46]. We note that there are alternatives to the windowed methods such as the wavelet coherence transform approach, which performs a time-frequency decomposition of the resting state signals from a pair of voxels or regions of interest [47]. Another non-windowed approach is to fit a hidden Markov model, which uses Bayesian inference to infer states and their dynamics with the assumption that the system is in precisely one state at any given time [48].

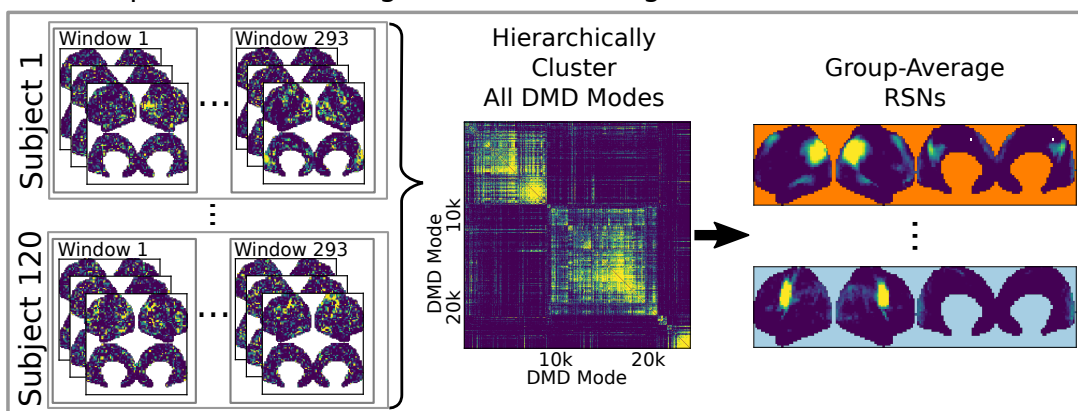
In this work, we present a new framework based on dynamic mode decomposition (DMD) for analyzing resting state BOLD fMRI data. DMD is a spatiotemporal modal decomposition technique [49, 50] ideally suited to extract coherent modes from rs-fMRI data. Similar to ICA, DMD decomposes a signal into constituent coupled spatiotemporal modes [51, 52, 53]. Unlike ICA, DMD constrains the modes to be temporally coherent; specifically, each mode oscillates at a fixed frequency while exponentially growing or decaying. This temporal coherence constraint produces more robust estimates of spatial modes by leveraging the assumption that RSNs have coherent dynamics within short windows of time, while automatically allowing for any phase differences between regions of a network. Thus, in addition to a spatial map of activation, DMD also estimates the temporal frequencies associated with each spatial mode.

Here we develop and validate a novel method to extract reproducible, time-resolved RSNs based on DMD. To demonstrate our method, we apply it to rs-fMRI data from 120 individuals in the Human Connectome Project (HCP) dataset [54, 55, 56]. DMD is computed for rs-fMRI data in windows, and the spatial components of the DMD, known as DMD modes, represent coherent spatial correlations across brain areas (Fig. 1A). We first consider group-DMD (gDMD), where hierarchical clustering discovers stereotypical modes present within the full population (Fig. 1B). These modes are strongly clustered, with clusters and sub-clusters corresponding directly to canonical RSNs. We next applied our DMD analysis independently to single scans (sDMD), as illustrated in Fig. 1C. Our technique extracts and characterizes individualized RSN modes and their dynamics at a resolution of ~ 3 s. To assess our time-resolved results, we compute the subject-level default mode networks (DMNs) through DMD and find that the individualized variations we observe agree with those derived through group-ICA (gICA) and dual regression. Importantly, this agreement is in spite of the fact that our DMD approach uses only data from a single subject, whereas the gICA approach regresses against a reference RSN computed from the entire population. We show that the sDMD modes resemble the gICA modes more closely and more robustly than subject-level DMNs computed via single-scan spatial ICA. Our subject-level analysis lends well to characterizing the RSN dynamics, including how frequently each RSN is active and the probabilities of transitions between different RSNs. We show that the sDMD derived RSN occupancies are reproducible within each individual across independent scans, and these dynamical properties may be used to build subject-level network representations of RSN dynamics. These results lay the groundwork for the application of the DMD family of methods to fMRI data; DMD is a modular approach that extends easily to multi-resolution analyses and control theoretic frameworks, a key advantage over ICA-based methods that have no directly analogous extensions. We suggest that DMD is an improved method for robust, reproducible extraction of functional connectivity modes from short windows of rs-fMRI data, and these time-resolved modes are ideally suited to characterize RSN dynamics on the single-subject level.

(A) Windowed DMD



(B) Group-Level Clustering of DMD Modes (gDMD)



(C) Subject-Level Clustering of DMD Modes (sDMD)




Figure 1: A schematic of our method using DMD to analyze resting state BOLD fMRI data. **(A)** Spatiotemporal DMD modes are extracted from short sliding windows of rs-fMRI data from 120 subjects from the Human Connectome Project. We use 23-sec windows and slide in 3-sec steps over each 15-min scan. **(B)** Group-DMD (gDMD) clusters all modes from the full population to extract group-averaged RSNs. **(C)** Subject-level DMD (sDMD) clusters modes from a single scan, which yields both individualized RSNs and their corresponding dynamics.

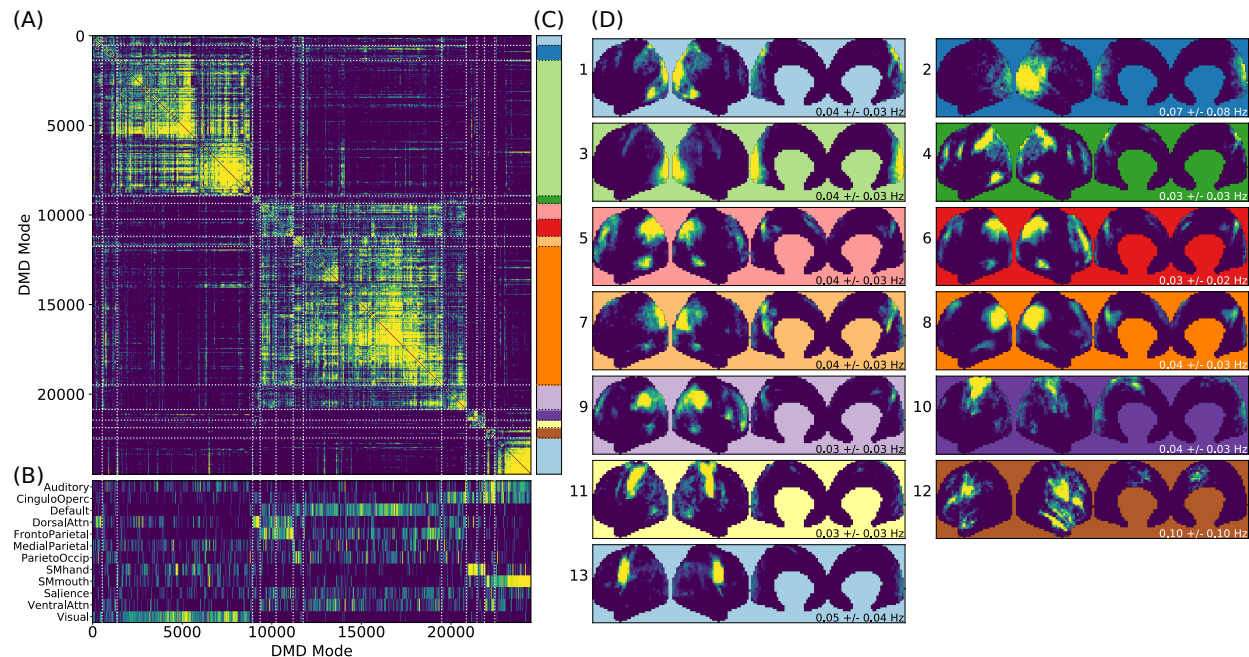


Figure 2: Results from gDMD using modes extracted from scans of 120 individuals, revealing clusters closely resembling known RSNs. (A) The hierarchically clustered correlation matrix between gDMD modes. (B) Overlap of each mode with RSNs from [57]. (C) Clusters automatically extracted from the hierarchical clustering results. A different threshold can be chosen to derive finer or coarser sub-clusters. (D) Averaged modes within each of the 13 clusters, with the corresponding DMD temporal frequencies. Several of the automatically identified DMD mode clusters visually correspond to canonical RSNs.

2. Results

Here we present results using our method combining dynamic mode decomposition (DMD) and unsupervised clustering to extract resting state networks (RSNs) from Human Connectome Project (HCP) data. We show results at both the group level (gDMD) and the subject level (sDMD), as summarized in Fig. 1. The sDMD results extract reproducible, time-resolved RSNs and their transition dynamics for individual scans. Details of our approach are described in the Methods. Code to assist in downloading the correct data, run all analyses described in this paper, and generate the corresponding figures are openly available at https://github.com/kunert/DMD_RSN.

2.1. Extracted group-level modes resemble Resting State Networks (RSNs)

The BOLD data we analyze is a collection of 15-min resting-state scans from 120 unrelated individuals from the HCP. Each scan has 1200 frames acquired at a temporal resolution of ~ 0.72 seconds. The data had been pre-processed with the HCP minimal pre-processing pipeline [56] and is freely available to download at <https://www.humanconnectome.org/>. As part of this pre-processing, the data are mapped to the “grayordinate” space, where each point maps onto the cortical surface.

To resolve modes in time, we apply DMD to short windows of data within each individual scan. In the results shown here (Fig. 2), we use ~ 23 -sec windows sliding in ~ 3 -sec increments. This windowing scheme corresponds to windows of 32 frames sliding by 4 frames, producing a total over 293 overlapping windows per each 1200-frame scan. DMD decomposes the data in each window into a set of spatiotemporal modes, and we truncate to the first 8 temporally coherent modes; this truncation produces a low-rank representation of the dominant spatiotemporal patterns in the data. Our results presented here are relatively robust to the number of modes calculated within each window, as described further in Appendix C.

To show that the DMD modes thus computed resemble known RSNs, we use an unsupervised clustering approach to automatically identify and label clusters of modes, then compare them to canonical RSNs. Specifically, we perform agglomerative clustering to hierarchically cluster our set of DMD modes based

on the average spatial correlation between clusters [59, 60] (as implemented in the `cluster.hierarchy` package of SciPy 1.0.0 [58]). In gDMD, we consider as a group all the DMD modes from all windows of 120 individuals in the dataset. The hierarchically clustered correlation matrix is shown in Fig. 2A, where the modes have been reordered so that this matrix is strongly block diagonal. In other words, groups of modes cluster naturally. Fig. 2B shows the correlation of each DMD mode to each of a set of reference RSNs [57]. Importantly, the automatically-grouped clusters appear to correspond to distinct RSNs. For instance, the strongly clustered block in the lower-right of the correlation matrix shows very strong correlations to the mouth sensorimotor RSN.

The clustering establishes boundaries that separate modes into distinct clusters, shown in different colors in Fig. 2C. This procedure uses a threshold that determines the resultant size and granularity of the clusters. Here we choose a threshold on the cophenetic distance (i.e. the “distance” threshold option of the `scipy.cluster.hierarchy` function `fclust` [58]) of 0.955, a relatively coarse grouping of the modes into 13 clusters whose averages are shown in Fig. 2D. These cluster averages, extracted at the population level, appear to correctly resemble known resting state networks. For instance, cluster 3 (light green) resembles the visual network, cluster 8 (dark orange) resembles the default network, and cluster 10 (dark purple) resembles the sensorimotor hand network.

Every DMD spatial mode has a corresponding temporal eigenvalue, which gives information about the oscillation frequency of that mode in time. Although these eigenvalues have not been used in the clustering process, they describe the characteristic oscillation frequency of each cluster. The mean and standard deviation for each cluster’s oscillation frequency is shown in Fig. 2D. As a validation that RSNs have low frequencies of oscillation, these frequencies are generally below or around 0.10 Hz. Further analysis of the frequency content of extracted gDMD clusters is discussed in Appendix B.

2.2. Subject-level modes and dynamics

One key advantage of our approach using DMD followed by unsupervised clustering is that it can be performed equally well on any subset of modes. We are particularly interested in performing the analysis on single scans, which we refer to as sDMD. Fig. 3 shows the results from hierarchically clustering modes from a single 15-minute scan. Similarly to the group results, this similarity matrix is strongly block diagonal, and the blocks correspond to different canonical RSNs (Fig. 3A, bottom). Note that we use a slightly lower threshold on the cophenetic distance than in Fig. 2 (0.95 instead of 0.955), yielding a larger number of finer-grained subclusters.

Importantly, in this single-scan analysis, the clusters are easily interpreted temporally: each cluster shows temporal dynamics defined by the time windows in which its constituent modes are found, as shown in Fig. 3B. Largely, clusters are active over periods spanning many consecutive windows and extended periods of time. Notably, many different clusters are observed to co-occur in time. This overlap of modes in the same window poses no problems for the sDMD approach, but it violates the assumptions made by other time-resolved methods (such as hidden Markov models) that require the system is in a single state at any particular time.

2.3. Subject-level modes capture spatial heterogeneity

Clusters derived by sDMD are able to reliably capture individual variability of RSNs. To quantify this feature, we compare one particular resting state network, the default mode network (DMN), as extracted by three different methods: sDMD, ICA, and group-ICA with dual regression (gICA). The gICA networks were computed by the HCP and are available on ConnectomeDB as part of the *High-level rfMRI Connectivity Analyses* data release. In short, this approach generates high-quality group ICA modes using the dataset of 1200 individuals then uses dual regression to adapt each group mode to the heterogenous structure seen within a particular scan.

Fig. 4 compares the DMNs extracted by ICA, sDMD, and gICA. Specifically, we use spatial-ICA on the entire window of scan data, as opposed to the sliding-window approach of sDMD. ICA and sDMD were both run multiple times with varying hyper-parameters, and for each method we choose the output that most strongly correlates with the canonical DMN. Specifically, ICA was performed multiple times as the fitting process is inherently stochastic, and will sometimes fail to extract a quality mode; thus we performed ICA

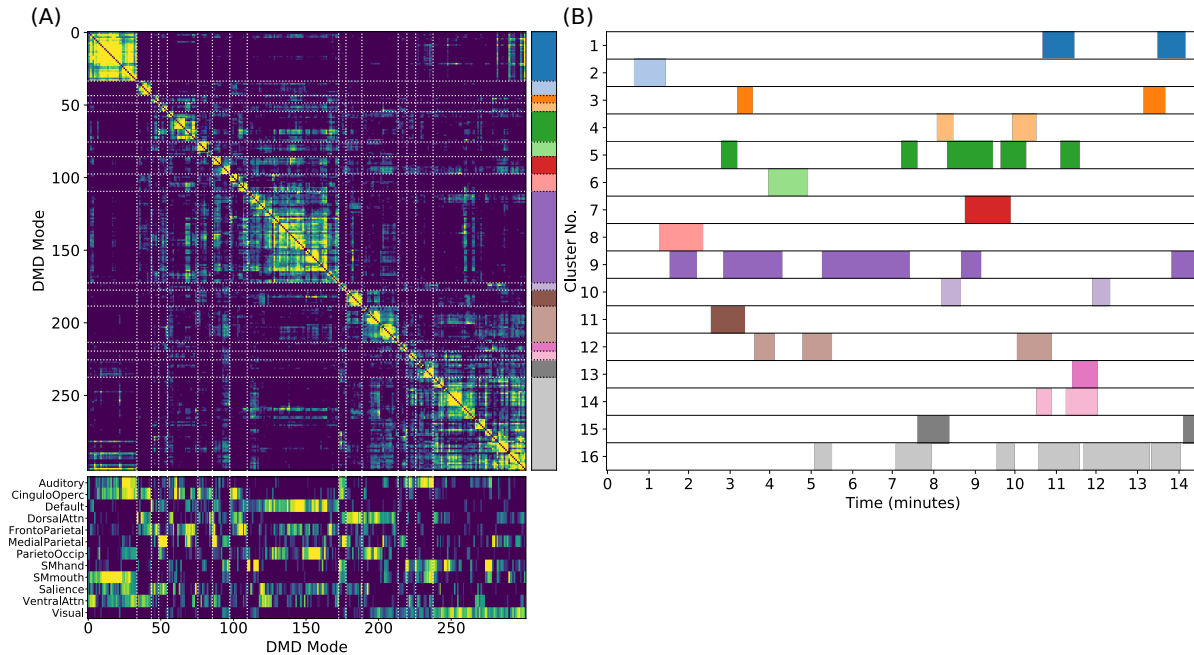


Figure 3: Results from sDMD using modes extracted from a single scan. **(A)** The hierarchically clustered correlation matrix between the sDMD modes, and their overlap with the RSNs of [57]. **(B)** Temporal activation of each cluster within the scan. Each row corresponds to a distinct cluster, and columns are colored in if the corresponding cluster is active within that time window.

on each scan ten times, each time randomly varying the number of modes in the range 8-16, and kept the mode which most strongly resembled the gICA DMN. sDMD, on the other hand, is inherently deterministic and was only performed once to produce a single set of modes; clustering of these modes was performed over a range of clustering parameters (the pre-clustering mode mask threshold, and the cophenetic distance threshold for forming clusters) and we kept the result most closely resembling the gICA mode. Importantly, the ICA and sDMD analyses were computed on data from single scans, whereas gICA produces personalized DMNs for each scan based on a population-derived DMN using dual regression. As shown in the examples from 3 different subjects in Fig. 4A, qualitatively speaking, ICA and sDMD are both able to extract DMNs that resemble the personalized gICA DMNs. In particular, note that some of the spatial heterogeneity of DMNs for individual subjects is recapitulated in the ICA and sDMD networks.

We show that sDMD reliably outperforms ICA in quantitative comparisons with gICA on a subject-by-subject basis. Fig. 4B shows that that DMNs from both ICA and sDMD correlate significantly more strongly with gICA results from the same individual (bottom panel) than with gICA results from the rest of the population (top panel). Specifically, the distribution of correlations for ICA has a lower mean and longer tail of low correlation values than the distribution of correlations for sDMD. This result indicates that sDMD extracts similar individualized DMN structures as gICA using data from only a single subject. Further, the sDMD analysis provides unambiguous time dynamics of the subject-level modes. The gICA or ICA modes can be correlated against the scan data to yield a measure of activation in time, but this is a continuous measure which requires thresholding and will be significantly nonzero in the case of any overlapping signal or noise. Our clustering approach, on the other hand, unambiguously identifies a cluster as active/inactive within a particular window of time.

2.4. Dynamic properties of modes are reproducible

Our proposed gDMD and sDMD analyses produce reproducible dynamics in addition to reliable spatial structures. This reproducibility criterion is critical for interpreting dynamic properties as meaningful reflections of the underlying functional networks. Useful dynamic properties to compute include how often

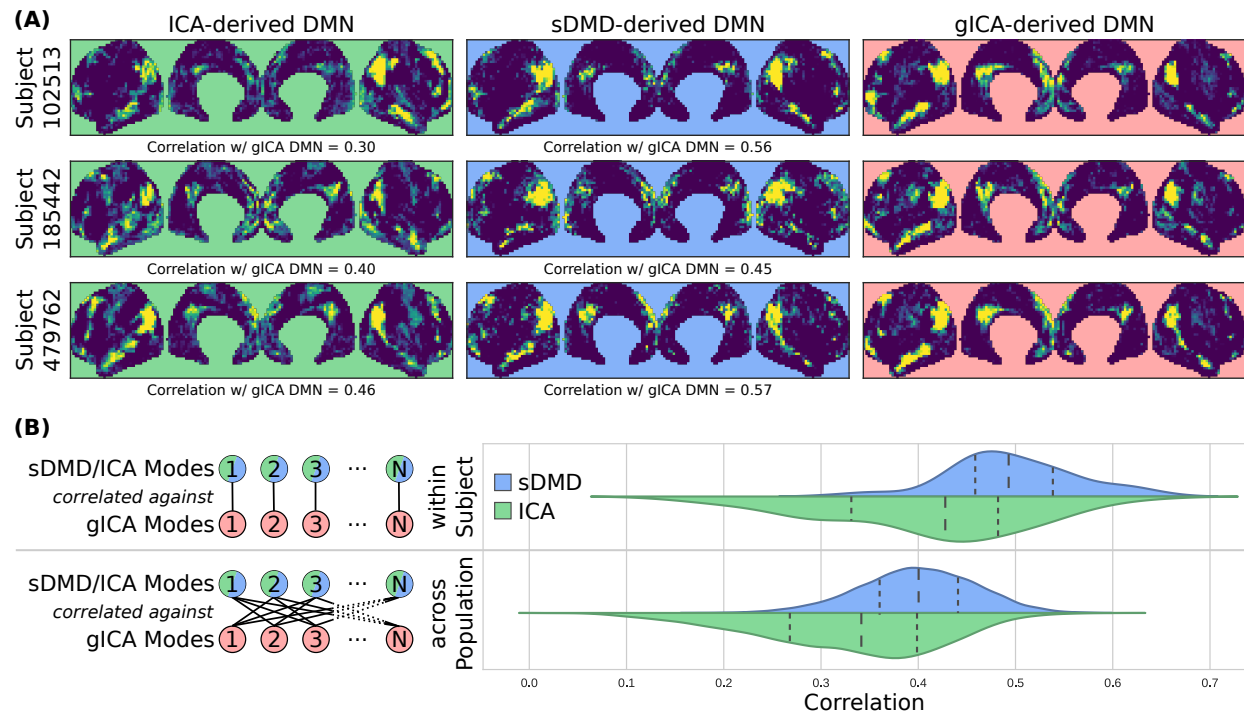


Figure 4: sDMD is able to capture the individual heterogeneity in spatial DMN structure. **(A)** Each column shows examples of subject-level DMNs as calculated by: (i) traditional ICA, (ii) sDMD, and (iii) group-ICA with dual regression (gICA). gICA yields the cleanest modes but requires a high-quality group-level mask, whereas ICA and sDMD use only data from a single scan. **(B)** The spatial correlation of ICA/sDMD modes with gICA modes. We compare how strongly each ICA/sDMD mode correlates with the gICA mode from the same individual, as well as how strongly it correlates with the full population's gICA modes. Both methods find modes which correlate more strongly with DMNs from the same individual, indicating that they all capture similar subject-level variations in DMN structure. sDMD outperforms ICA in this regard, however, and correlates more strongly with the gICA modes. At the same time, sDMD provides more unambiguous temporal information than either approach.

each mode is active, and how often pairs of modes are active together (either simultaneously, or with one mode transitioning into another). To assess reproducibility, we take advantage of the HCP resting state dataset, where each of the 120 subjects underwent four (4) different 15-min resting state scans. Specifically, we repeat the same analysis for each of these four sets of scans separately and quantify to what extent our DMD analysis extracts similar modes and dynamic properties. Here we examine dynamic networks extracted by gDMD, looking at two simple measure of dynamics: the occupancy of and the transfer matrices between networks. Next, we quantify the reproducibility of these dynamics at a single-subject level by using sDMD on each of the 4 scans for each subject.

Modes were paired between the four independent scans based upon their spatial similarity; this matching procedure, as well as visualizations of all extracted modes, is detailed in Appendix A. Fig. 5A shows the six most highly correlated modes. Fig. 5B shows the occupancy matrix O_{ij} of these 6 modes, where the value in the i^{th} row and j^{th} column corresponds to how often modes i and j are simultaneously active. The diagonal O_{ii} shows how often each mode i is active, and the matrix is symmetric. Next, Fig. 5C shows a transfer matrix T_{ij} , which contains the probability that, if mode i is active, then mode j will be active 30 seconds later. This matrix not symmetric, and it encodes rich dynamical information that can be interpreted as a dynamic network model, as shown schematically in Fig. 5D.

We next use these same 6 modes from the gDMD and repeat the analysis for individual single scans, analyzing the extent to which the dynamic properties are unique to each individual. The occupancy of each mode was broken up into how often it is present in individual scans, and Fig. 6 shows the results of the subject-level occupancy correlation between two different sets of scans of the same individuals. Correlation coefficients and p -values are calculated for all possible pairs of the 4 different scans, and Fig. 6 reports the median correlation across all scan comparisons and the corresponding p -value. For five of the six modes,

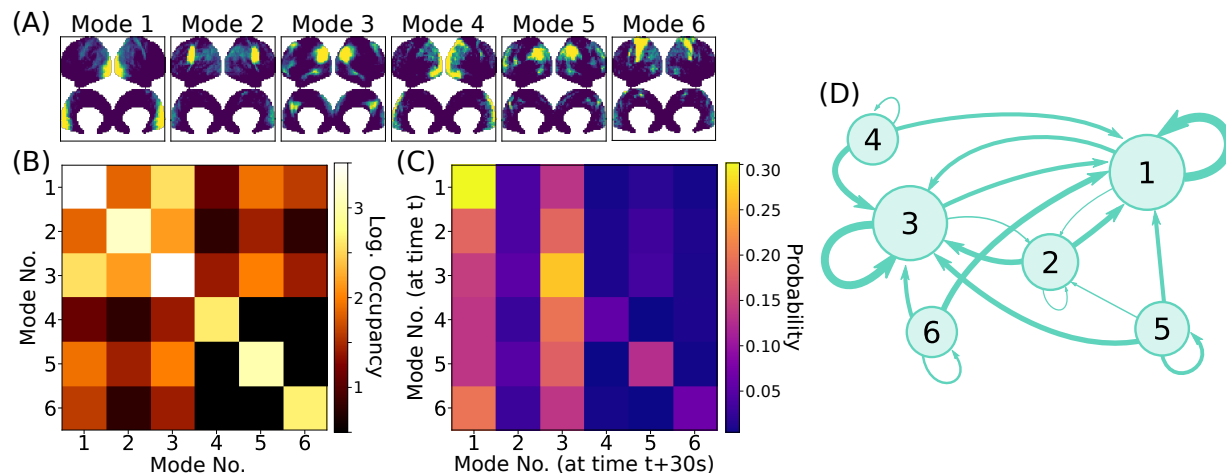


Figure 5: Dynamic properties of networks as extracted by gDMD revealed by the occupancy and transitions of modes. (A) gDMD was performed on four independent sets of scans of the same population of 120 individuals. The six modes plotted here are those extracted most reliably from the four datasets. (B) The occupancy matrix O_{ij} : each entry shows how often mode i and j are active simultaneously (the diagonal simply shows how often mode i is active). (C) The transfer matrix T_{ij} : this shows the probability that if mode i is active, mode j will be active 30 seconds later. (D) The transfer matrix can be interpreted as a network which characterizes RSN dynamics.

occupancy results from different scans of the same individual are positively correlated, with the strongest and most significant correlations observed for the most active modes. Note that the high p -value of Mode 4 does not imply that it is not meaningful, but rather suggest that its activity may not correlate within an individual across different scans (as would be the case, for example, if mode activity was uniformly probable across the population). Thus, our approach extracts reproducible, individualized dynamic properties of resting state networks.

3. Discussion

In this work we present a novel framework based on the dynamic mode decomposition (DMD) for extracting time-resolved resting state networks at both the group (gDMD) and single-subject (sDMD) levels. DMD decomposes high-dimensional time-series data into a sum of dominant coupled spatiotemporal modes. One way to think about DMD is that it is similar to a combination of applying principal component analysis (PCA) in space and the Fourier transform in time, as it extracts both spatial maps and their coupled frequency content simultaneously. We demonstrate our approach using data from 120 subjects from the Human Connectome Project (HCP). DMD to extract coherent spatiotemporal modes from short, sliding windows of data, and patterns modes are revealed by unsupervised clustering. The gDMD clusters correspond to the average modes within a population and closely resemble canonical RSNs. The sDMD use only data from a single scan, calculating subject-level RSNs and their temporal dynamics simultaneously. When sDMD is used to extract subject-level Default Mode Networks from the data of a single scan, it does so more reliably than ICA, while also having unambiguous time resolution. Finally, we show that the extracted temporal dynamics are highly reproducible within subjects and between trials.

DMD extracts additional information beyond what we have fully considered within the scope of this work. By construction, each DMD has an intrinsic temporal frequency. We have made note of these frequencies (e.g. as labels within Fig. 2) but have drawn no conclusions from them beyond noting that they are plausible compared to the known literature (i.e. generally $< 0.1\text{Hz}$). However, the preliminary analyses in Appendix B suggest that networks may have distinct and reproducible frequency content which could be characterized by future studies. Similarly, it remains to be systematically investigated the phase information in each mode and how they relate to resting state network organization.

Beyond data that has already been pre-processed and mapped onto “grayordinate” cortical surface space, it may be desirable to analyze volumetric data directly. In principle, there is no reason why our DMD

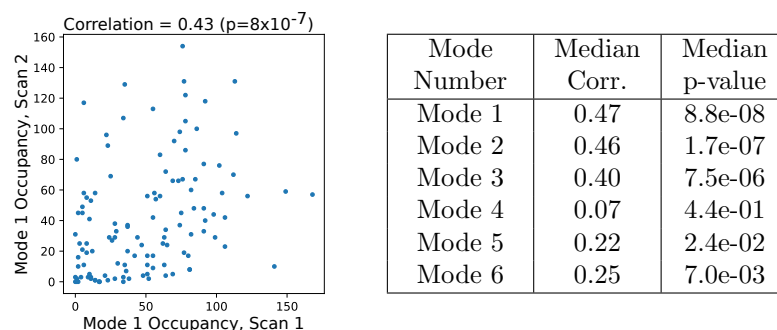


Figure 6: Occupancy of modes computed by sDMD is correlated across different scans of the same subjects. The 6 modes are the same as in those in Fig. 5. Spearman’s rank correlation is computed of mode occupancies between all pairs of the 4 scans. The plot at left shows the occupancy of Mode 1 (O_{11}) for each subject in two different scans (REST1_LR and REST2_LR). There is a significant correlation; subjects with high activity in one scan are more likely to have high activity in another. In the table at right, we have correlated all possible pairs of the four scans and report the median correlation (and corresponding p -value). The correlations have $p < 0.05$ for all but mode 4.

approach cannot be applied equally well to voxel data; nevertheless, the quality of the networks thus extracted would be unknown and must be thoroughly characterized. Similarly, care should be taken when applying our approach to data acquired at different spatial and temporal resolutions. We have made all of our code openly available in the hope that these analyses can be performed readily and reproducibly.

This demonstration that DMD is capable of extracting modes reliably from short windows of data, with similar or better performance than ICA, lays the groundwork for future methodological developments. Indeed, a key motivation for the use of DMD is that a number of powerful extensions to DMD have been developed that could be readily applied to BOLD fMRI data. DMD with Control (DMDc [61]) uses a control theoretic framework, allowing intrinsic dynamics to be estimated from those driven by external stimuli. DMDc is a natural framework to analyze task-based fMRI scans. Multi-resolution DMD (mrDMD [62]) improves the quality of the temporal information extracted over multiple timescales. These straightforward extensions of DMD to multi-resolution and control frameworks have no direct counterpart in ICA, giving DMD a critical advantage over ICA in future analyses of multi-scale signal extraction and control. Similarly, Optimized DMD [63] is an implementation of the DMD algorithm which is more computationally expensive, but significantly improves the precision of the extracted frequency information, which would enable the study of how characteristic frequencies may vary between networks, as suggested by the preliminary analyses of Appendix B.

It is increasingly clear that RSN dynamics are significantly altered by neurological disorders, and their impact on the spatial and temporal structures of RSNs may be diverse and highly individualized. Rather than searching only for patterns in RSNs strongly resembling those seen in the general population, it may be fruitful to explore methods that are agnostic to the average. We propose that DMD is an approach well suited for data-driven extraction of individualized structure and dynamics of networks from single scans, and we believe the development of the DMD family of methods opens new doors for exploring and characterizing whole-brain dynamics as captured by fMRI.

4. Methods

4.1. Human Connectome Project rs-fMRI Data

We used rs-fMRI data from 120 unrelated individuals in the Human Connectome Project (HCP) dataset as provided by the WU-Minn Consortium [54, 55]. The HCP acquired four separate resting state scans using a multiband pulse sequence (multiband factor = 8, FA = 52°), each scan having 1200 time points (14.4 minutes, TR: 720ms, TE: 33ms, FOV: 208mm x 180mm, Matrix: 104 x 90, with 72 slices) [56]. For each scan, an equal number of volumes with left-right and right-left phase encoding directions were acquired. To correct the fMRI scans for gradient distortions, the HCP also acquired two spin echo EPI images with reversed phase encoding directions, with a unique pair of spin echo EPI images for each of the resting state acquisitions

[64]. These spin echo EPI images were then used to accurately spatially normalize the fMRI volumes to the structural scans. For the remainder of this document, we label and refer to the four resting-state runs as “REST1_LR”, “REST1_RL”, “REST2_LR”, and “REST2_RL”).

The resting state scans had already been preprocessed by the HCP Consortium using the HCP minimal preprocessing pipeline [56]. In addition to gradient distortion correction, this preprocessing includes fMRI denoising using FIX [65], masking, motion correction, registration and interpolation of the timeseries onto the cortical surface (the CIFTI “grayordinate” space). Included in this preprocessing was the application of FreeSurfer to generate mesh representations of the the cortical surface. These FreeSurfer surface meshes were upsampled to $\sim 164k$ vertices, and subsequently downsampled to $\sim 32k$ vertices, resulting in an average vertex spacing of roughly 2mm at the cortical midthickness level. The resting state time series were then mapped onto this low-resolution 32k vertex mesh. This surface-mapped data consists of 91,282 “grayordinates” ($\sim 60k$ surface points mapping the cortical surfaces, along with $\sim 30k$ subcortical voxels). In this manuscript, we restrict our attention to the 2D cortical surface mesh. Note, however, that DMD could be similarly applied to 3D voxelized data (see Discussion).

All data used in this manuscript are freely available for download from the HCP. Additionally, scripts for downloading the correct data, running all analyses described within this paper, and figure generation are openly available at <https://github.com/kunert/DMD-RSN>.

4.2. Dynamic Mode Decomposition

This section describes the specific DMD algorithm implemented in this manuscript [49, 50]. Each scan is broken into windows of data, and data from each window is collected into a data matrix \mathbf{X} , where each row represents one of the n grayordinates and each column \vec{x} is one of the m time snapshots in that particular window:

$$\mathbf{X} = \begin{bmatrix} | & | & \dots & | \\ \vec{x}_1 & \vec{x}_2 & \dots & \vec{x}_{m-1} \\ | & | & \dots & | \end{bmatrix}. \quad (1)$$

Here \vec{x}_k is a vector in grayordinates space giving the BOLD signal at time index k . Next, we also define the time-shifted data matrix \mathbf{X}' , which is defined similarly but with each column advanced forward by a single timepoint:

$$\mathbf{X}' = \begin{bmatrix} | & | & \dots & | \\ \vec{x}_2 & \vec{x}_3 & \dots & \vec{x}_m \\ | & | & \dots & | \end{bmatrix}. \quad (2)$$

The goal of DMD is to describe the matrix \mathbf{A} that best maps \vec{x}_k onto \vec{x}_{k+1} , namely by solving for the eigenvalues and eigenvectors of \mathbf{A} . In other words, we treat dynamics of the system as approximately linear and seek the \mathbf{A} that best solves

$$\mathbf{X}' = \mathbf{A}\mathbf{X}. \quad (3)$$

However, this \mathbf{A} matrix is a $n \times n$ square matrix; for the HCP data, $n = 91,282$ grayordinants, so \mathbf{A} has approximately 10^{10} elements — this is a tremendously large number of unknowns, and they are poorly constrained by the limited data in \mathbf{X} and \mathbf{X}' . Further, we hypothesize that many of these brain areas have strong correlations, so their dynamics are relatively low rank and may be explained by far fewer modes.

The DMD algorithm takes advantage of the singular value spectrum (SVD) of the data matrix \mathbf{X} to obtain the dominant eigendecomposition of \mathbf{A} without actually computing \mathbf{A} . We first take the singular value decomposition (SVD) of the data matrix $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, which decomposes the \mathbf{A} into the product of unitary matrices \mathbf{U} , \mathbf{V} and diagonal matrix $\mathbf{\Sigma}$. Because the singular values and singular vectors are ordered by decreasing energy, this decomposition can be used truncated by taking the first r columns of \mathbf{U} to form a $n \times r$ matrix \mathbf{U}_r . We may similarly form a $r \times r$ matrix $\mathbf{\Sigma}_r$ and a $m \times r$ matrix \mathbf{V}_r .

This truncation gives the optimal rank- r reconstruction of the data matrix, $\mathbf{X} \approx \mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^*$. The choice of the number of SVD modes r is equal to the resultant number of DMD modes r and is used as a parameter throughout this manuscript.

Taking the SVD allows us to then calculate the pseudoinverse $\mathbf{X}^\dagger = \mathbf{V}_r\mathbf{\Sigma}_r^{-1}\mathbf{U}_r^*$ and simply solve for \mathbf{A} :

$$\mathbf{A} \approx \mathbf{X}'\mathbf{X}^\dagger \approx \mathbf{X}'\mathbf{V}_r\mathbf{\Sigma}_r^{-1}\mathbf{U}_r^*. \quad (4)$$

Given that the number of grayordinates is large, we do not consider the $n \times n$ \mathbf{A} matrix directly, but instead project into the reduced-dimensional \mathbf{U}_r basis:

$$\tilde{\mathbf{A}} \equiv \mathbf{U}_r^* \mathbf{A} \mathbf{U}_r = \mathbf{U}_r^* \mathbf{X}' \mathbf{V} \mathbf{\Sigma}^{-1}. \quad (5)$$

Additionally, we scale each mode according to how strongly it is present in the original data, and each modes is scaled by the singular values $\mathbf{\Sigma}_r$ as in [50]:

$$\hat{\mathbf{A}} \equiv \mathbf{\Sigma}^{-1/2} \tilde{\mathbf{A}} \mathbf{\Sigma}^{1/2}. \quad (6)$$

We then compute the eigendecomposition of this scaled, reduced dimensional $\hat{\mathbf{A}}$ matrix:

$$\hat{\mathbf{A}} \hat{\mathbf{W}} = \hat{\mathbf{W}} \mathbf{\Lambda}, \quad (7)$$

where $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues λ_j , and the columns of $\hat{\mathbf{W}}$ are the eigenvectors of $\hat{\mathbf{A}}$. These eigenvectors can be used to compute the dynamic modes:

$$\Phi \equiv \mathbf{X}' \mathbf{V} \mathbf{\Sigma}^{-1/2} \hat{\mathbf{W}}, \quad (8)$$

where the j -th column of Φ is the dynamic mode $\vec{\phi}_j$. From the dynamic modes and corresponding eigenvalues, we can approximate the dynamics of the system as:

$$\vec{x}_k \approx \sum_{j=1}^r b_j \lambda_j^k \vec{\phi}_j, \quad (9)$$

where the weights b_j can be fit as initial conditions for at $k = 0$. We can also write this in terms of a continuous time t by writing the dynamics in terms of a complex exponential:

$$\vec{x}(t) \approx \sum_{j=1}^r b_j \exp((\gamma_j + i2\pi f_j) \cdot t) \vec{\phi}_j. \quad (10)$$

In the exponent, γ_j is the real-valued growth/decay constant, and f_j is the real-valued oscillation frequency of the mode in cycles per second (hertz). These are calculated from the real and imaginary components of the corresponding DMD eigenvalues:

$$\gamma_j = \frac{\text{real}(\ln(\lambda_j))}{\Delta t}, \quad f_j = \frac{\text{imag}(\ln(\lambda_j))}{2\pi \Delta t}, \quad (11)$$

where Δt is time between measurements, in seconds. This change of units does not carry any additional information over the complex eigenvalue λ_j , but γ_j and f_j are readily interpretable as standard growth/decay constants and oscillation frequencies.

Like the eigenvalues, the dynamic modes $\vec{\phi}_j$ can be complex valued. In other words, each element of $\vec{\phi}_j$ has both a magnitude and a phase. In this work, we consider only the magnitude of the dynamic mode $|\vec{\phi}_j|$; generally, when we refer to the “dynamic mode”, or when we do calculations such as spatial correlations, etc., we are referring to and using $|\phi_j|$. However, the complex-valued ϕ_j also contains the relative phase between regions, which is of potential interest for future analyses. For real-valued data \vec{x}_k , oscillatory modes appear as conjugate pairs $\{\vec{\phi}_j, \vec{\phi}_{j'}\}$ with the same spatial magnitudes $|\phi_j| = |\phi_{j'}|$; thus the number of unique spatial patterns extracted from a data matrix will be $\leq r$.

4.3. Independent Component Analysis

fMRI data is often processed using Independent Component Analysis (ICA) [66]. Just as above, if we collect all data from a scan into a data matrix \mathbf{X} , then we can formulate ICA as an attempt to solve the following:

$$\mathbf{X}^T = \mathbf{A} \mathbf{S}^T. \quad (12)$$

The goal of spatial ICA is to decompose the data into non-orthogonal, statistically independent source signals, which are the columns of the $n \times r$ matrix \mathbf{S} . The $m \times r$ mixing matrix \mathbf{A} gives the mixture of these signals, at each timepoint, which best approximates the original data.

However, because individuals subjects have their own unique time courses, comparing ICA components estimated at the subject-level is difficult, and consequentially, how to make inference about group-level components is not immediately clear. As an alternative, fMRI analyses often employ multi-subject ICA methods to estimate 1) population-based components, or 2) single-subject components informed at the group level. A variety of multi-subject methods have been developed, including temporal or spatial concatenation group ICA approaches [67, 68], higher-order tensor decomposition methods [40], and dual regression-based approaches [69, 70].

In this analysis, we use the FastICA algorithm as implemented in scikit-learn 0.18 [71]. FastICA attempts to maximize the statistical independence of the signals by maximizing the non-Gaussianity of the projected data, as detailed in [72]. In order to directly compare DMD to ICA, we compare both algorithms at two spatial component levels: using components informed at the single-subject level and using components informed at the group level. For the single-subject-level comparison, we compute single-subject ICA components using sklearn’s FastICA. For the group-level comparison, we use the group ICA and dual regression results generated previously by the HCP using Melodic’s Incremental Group-PCA algorithm (MIGP) [73]. MIGP is an iterative approach that incrementally incorporates time-series data for single subjects and updates a current running estimate of the set of spatial eigenvectors that best describes the time series data for the current set of incorporated subjects. MIGP generates a very close approximation to group-ICA components generated by classic temporal concatenation group-ICA approaches, but without the large computer memory requirements. The HCP then used dual regression to regress each individual subject’s time-courses on the MIGP-generated group-ICA components.

4.4. Sliding Window Mode Calculations

The rs-fMRI scans analyzed each consist of 1200 individual timesteps, which we break up into sliding windows. We use a simple square window (as opposed to a scheme which weights timepoints depending upon their position in the window), such that our calculations are performed on a submatrix of the full data matrix.

There are three parameters to consider when computing DMD and/or ICA modes within a sliding window: T , the length of each window; dT , the number of timesteps by which to slide the window; and r , the number of modes to compute within each window. We observed that our results are relatively robust to these parameters. Within the results shown in the main manuscript, we have chosen $T = 32$, $dT = 4$, and $r = 8$. In seconds, this corresponds to roughly 22.4s windows which are translated in increments of about 2.8s. As discussed in Section 4.2, oscillatory DMD modes appear in conjugate pairs, and thus $\leq r = 8$ modes will be extracted within each window. For example, the gDMD set used in Fig. 2 (the modes calculated from the REST1_LR scans of 120 subjects) consists of 160,756 total modes.

4.5. Mode Visualization

DMD modes are computed on the full 59,412 -dimensional grayordinate space. However, this space is unnecessarily large for visualization purposes, for which we primarily want to visualize the macroscopic structure of individual modes. It is also disadvantageous for the clustering of modes; we want to cluster modes based upon their overall structural similarity, and applying some method of spatial smoothing is helpful in minimizing spurious, noise-driven correlations. For visualization and clustering purposes, we therefore bin the modes, computing the average magnitude of grayordinates within a particular patch of space.

Each grayordinate can be mapped onto a 3D coordinate in space. From this information, we first classify all grayordinates as on the left/right cortex and on the lateral/medial side. Each of these groups of grayordinates are then projected onto the sagittal plane and divided into 40 bins in each direction. This choice of bin granularity is arbitrary but was chosen heuristically to yield good visualization and clustering performance. Bins which contain no grayordinates are discarded, resulting in a total of 3856 bins.

Given a 59,412 -dimensional DMD mode ϕ_j , we may average over the grayordinates within each bin to yield a 3,856-dimensional binned mode m_j . All modes visualized within this paper have been binned in this fashion. Furthermore, modes have been binned before being clustered or correlated against each other, which was seen to substantially improve the performance of clustering.

4.6. Hierarchical Clustering

Having performed sliding window DMD mode extraction on the full set of 120 individuals, we wish to cluster modes based upon their spatial similarity on both the single-subject and whole-group level. In both cases we use the same hierarchical clustering pipeline, described as follows, applied to either the set of modes extracted from a single scan (the sDMD case) or from the full set of scans (the gDMD case).

A few pre-processing steps were found to be helpful in generating clean, robust clusters. First, modes are spatially binned as described in Section 4.5. This reduces the dimensionality of modes for clustering purposes, averaging over noise while preserving larger-scale spatial structure. We then wish to filter out modes which lack large-scale continuous spatial structure. A quick heuristic for the spatial continuity of a mode is calculated and thresholded upon as follows: for each binned mode, we count the number of bins with a z-score ≥ 2 which also have a neighbor with z-score ≥ 2 . Thresholding on c keeps only those modes which possess a certain level of spatial continuity (e.g. in Fig. 2 we cluster only modes with $c \geq 25$).

To group this filtered set of modes into a small set of interpretable clusters, we use hierarchical clustering as implemented in SciPy 1.0.0[58]. There are many metrics of spatial similarity which can be used in such a procedure. Qualitatively, we found the best performance (as indicated by the formation of tight, discrete clusters which group modes with similar large-scale spatial structure) by thresholding modes at a defined z-score, z_t , to generate spatial masks (in Fig. 2 we use $z_t = 2.5$). These masks were then clustered based upon the average spatial correlation within a group of modes.

We then form flat clusters that have a cophenetic distance of no greater than a defined threshold t (i.e. using the ‘distance’ threshold option of the `scipy.cluster.hierarchy` function `fcluster`). Fig. 2 shows the clusters which are formed using $t = 0.955$, whereas Fig. 3 shows finer-grained subclusters formed using $t = 0.950$. The choice of t will affect the size and granularity of the cluster assignment, and can be varied to obtain smaller subclusters of any particular cluster. Note that in Figs. 2 and 3 we have only visualized clusters containing a minimum number of modes (≥ 400 in the gDMD case of Fig. 2, and ≥ 5 for the sDMD case of Fig. 3).

We compare these modes against a set of canonical RSNs, using the parcellation of Gordon et al. [57]. In Fig. 2(B) we show the correlation of each mode against each of these canonical masks, where the ordering of the modes has been defined by the hierarchical clustering. The clustering succeeds at grouping modes which resemble similar RSNs, and many of the clusters appear to correspond to distinct networks: for instance, the bottom-rightmost cluster (cluster number thirteen) strongly corresponds to the “SMmouth” network. When the average of the modes within each cluster is plotted in Fig. 2(D), it indeed has the expected spatial structure.

4.7. Subject-Level DMNs

The subject-level Default Mode Networks (DMNs) shown in Fig. 4 are found by calculating the sDMD clusters for modes from a particular subject, and taking the average of the cluster which most resembles the DMN. The optimal choice of z_t and t will vary from scan to scan; we therefore perform the clustering procedure several times for values of $z_t \in (1.5, 3.0)$ and $t \in (0.65, 0.98)$, saving the output cluster average which has the highest correlation with the gICA result. Though the choices of z_t and t afford relatively little control over the final spatial structure of the result, we must be mindful of the fact that we are selecting the best-performing result out of several different runs. For a fair comparison, we afford ICA the same opportunity: allowing the number of modes to vary in the range $r \in (8, 10, 12, 16)$, calculating ICA multiple times, and saving only the result which correlates most highly with the gICA result. This was seen to be necessary, as the output of the ICA algorithm is not deterministic and will sometimes yield a low-quality result (often yielding correlations of ~ 0.1 , as in the long tail of Fig. 4(B)).

Each of these approaches were compared against the group-ICA and dual regression approach, the results of which are calculated and provided for download by the HCP [70]. This approach first calculates group-ICA, which is ICA performed on the temporally-concatenated full dataset of scans from 1200 individuals. This calculates high-quality, population-level averaged gICA modes. The gICA mode corresponding to the DMN is then regressed against the scans for a single individual, to calculate an approximate timecourse of DMN activation within a particular individual’s scans. There is then a second regression step of the timecourses of all spatial coordinates against the averaged-DMN timecourse. This yields a spatial map of coordinates which have similar time dynamics to the population-level DMN mode, which in effect yields a

subject-level spatial map of the DMN. This results in clean, high-quality maps as seen in Fig. 4, but has the drawback of requiring a reference mask calculated from the entire set of population data.

Fig. 4(B) plots how strongly each subject-level ICA/sDMD DMN correlates with both (i) the full set of gICA modes, and (ii) the gICA mode calculated for the same subject. These are calculated by thresholding the modes at a z -value of $z = 2$ and then calculating the correlation coefficients between the thresholded masks.

4.8. Characterizing Cluster Dynamics

Assigning time dynamics to individual clusters (as in Fig. 3(B)) follows trivially from the windowing process: each cluster consists of a collection of modes, each calculated in a particular window corresponding to a particular time. The time dynamics of a cluster are simply given by the times of the windows in which its constituent modes were found. Notably, this gives a binary measure of mode activation (a cluster either contains a mode within a particular time window or it doesn't). This is distinct from other methods of assessing the temporal dynamics of modes, such as the common technique of correlating a spatial mode with each frame of a scan; such techniques yield a continuous measure of mode activation, which does not fully disambiguate the activity of a mode from that of overlapping spatial patterns or random fluctuations. This disambiguation is accomplished in our approach through the assignment of modes into discrete clusters.

The gDMD pipeline was performed independently on the four sets of scans for all 120 individuals (the sets “REST1_LR”, “REST1_RL”, “REST2_LR”, and “REST2_RL”). This resulted in four sets of gDMD modes, which were correlated against each other to find the modes which were most similar across all sets. We chose the top six most consistent modes, as described in Appendix A, with the resulting modes visualized in Fig. 5(A). Each set of modes has associated time dynamics; as examples of how to characterize these dynamics we calculate both the occupancy matrix O_{ij} and transfer matrix T_{ij} . The occupancy matrix O_{ij} indicates the number of windows in which both modes i and j are active. The diagonal elements O_{ii} count the total number of times mode i was active (with or without the presence of other modes). The transfer matrix T_{ij} shows the probability that if mode i is present, then mode j is present 30s later.

We assess the reproducibility of our approach by comparing the occupancies of each mode by each individual between different scans. This rests upon the assumption that if an individual shows increased activity in a particular mode in one scan, they are then more likely to be more active in the same mode within another scan. Note that we do not expect a perfect correlation near 1.0, but we do anticipate a significant one; for example, if an individual has a particularly active DMN in one scan, we may expect that that same individual is more likely to have an active DMN in another scan, but not that the DMN should be active for the exact same amount of time. An example of such a comparison is plotted in Figure 6, where each point represents a different individual, and we do indeed observe a positive correlation. We correlated all combinations of the four scans against each other for all different modes, and report the median correlation and occupancy within the table of Fig. 6. A significant correlation (in the sense that $p < 0.05$) is clear for all modes except for mode 4, indicating that our characterization of the dynamics indeed encodes underlying dynamic properties specific to different individuals.

5. Acknowledgements

The authors thank Nikita Sakhanenko for useful discussions. JMK and DJG acknowledge support from the Bill and Melinda Gates Foundation, and the Pacific Northwest Research Institute. JNK acknowledges support from the Air Force Office of Scientific Research (FA9550-17-1-0329). BWB acknowledges support from the National Science Foundation (award 1514556), the Alfred P. Sloan Foundation, and the Washington Research Foundation. SDR acknowledges support from the National Institutes of Health/National Institute on Aging (K01AG055669).

- [1] Marcus E Raichle, Ann Mary MacLeod, Abraham Z Snyder, William J Powers, Debra A Gusnard, and Gordon L Shulman. A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682, 2001.

- [2] JS Damoiseaux, SARB Rombouts, F Barkhof, P Scheltens, CJ Stam, Stephen M Smith, and CF Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the national academy of sciences*, 103(37):13848–13853, 2006.
- [3] Michael D Fox and Marcus E Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature reviews neuroscience*, 8(9):700, 2007.
- [4] MJ Lowe, BJ Mock, and JA Sorenson. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *Neuroimage*, 7(2):119–132, 1998.
- [5] Martijn P Van Den Heuvel and Hilleke E Hulshoff Pol. Exploring the brain network: a review on resting-state fmri functional connectivity. *European neuropsychopharmacology*, 20(8):519–534, 2010.
- [6] R Matthew Hutchison, Thilo Womelsdorf, Joseph S Gati, Stefan Everling, and Ravi S Menon. Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. *Human brain mapping*, 34(9):2154–2177, 2013.
- [7] Ming Song, Yuan Zhou, Jun Li, Yong Liu, Lixia Tian, Chunshui Yu, and Tianzi Jiang. Brain spontaneous functional connectivity and intelligence. *Neuroimage*, 41(3):1168–1176, 2008.
- [8] Martijn P van den Heuvel, Cornelis J Stam, René S Kahn, and Hilleke E Hulshoff Pol. Efficiency of functional brain networks and intellectual performance. *Journal of Neuroscience*, 29(23):7619–7624, 2009.
- [9] Jonathan S. Adelstein, Zarrar Shehzad, Maarten Mennes, Colin G. DeYoung, Xi-Nian Zuo, Clare Kelly, Daniel S. Margulies, Aaron Bloomfield, Jeremy R. Gray, F. Xavier Castellanos, and Michael P. Milham. Personality is reflected in the brain’s intrinsic functional architecture. *PLOS ONE*, 6(11):1–12, 11 2011.
- [10] Luqing Wei, Xujun Duan, Yang Yang, Wei Liao, Qing Gao, Ju-rong Ding, Zhiqiang Zhang, Weixi Zeng, Yuan Li, Guangming Lu, and Huafa Chen. The synchronization of spontaneous bold activity predicts extraversion and neuroticism. *Brain research*, 1419:68–75, 2011.
- [11] CJ Honey, O Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.
- [12] Zarrar Shehzad, AM Clare Kelly, Philip T Reiss, Dylan G Gee, Kristin Gotimer, Lucina Q Uddin, Sang Han Lee, Daniel S Margulies, Amy Krain Roy, Bharat B Biswal, Eva Petkova, F. Xavier Castellanos, and Michael P. Milham. The resting brain: unconstrained yet reliable. *Cerebral cortex*, 19(10):2209–2229, 2009.
- [13] Thomas Meindl, Stefan Teipel, Rachid Elmouden, Sophia Mueller, Walter Koch, Olaf Dietrich, Ute Coates, Maximilian Reiser, and Christian Glaser. Test–retest reproducibility of the default-mode network in healthy individuals. *Human brain mapping*, 31(2):237–246, 2010.
- [14] Koene RA Van Dijk, Trey Hedden, Archana Venkataraman, Karleyton C Evans, Sara W Lazar, and Randy L Buckner. Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. *Journal of neurophysiology*, 103(1):297–321, 2010.
- [15] Catie Chang and Gary H Glover. Time–frequency dynamics of resting-state brain connectivity measured with fmri. *Neuroimage*, 50(1):81–98, 2010.
- [16] Danielle S Bassett, Nicholas F Wymbs, Mason A Porter, Peter J Mucha, Jean M Carlson, and Scott T Grafton. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18):7641–7646, 2011.
- [17] Eswar Damaraju, Elena A Allen, Aysenil Belger, JM Ford, S McEwen, DH Mathalon, BA Mueller, GD Pearson, SG Potkin, A Preda, JA Turner, JG Vaidya, TG van Erp, and VD Calhoun. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clinical*, 5:298–308, 2014.

- [18] Sai Ma, Vince D Calhoun, Ronald Phlypo, and Tülay Adalı. Dynamic changes of spatial functional network connectivity in healthy individuals and schizophrenia patients using independent vector analysis. *Neuroimage*, 90:196–206, 2014.
- [19] Robyn L Miller, Maziar Yaesoubi, Jessica A Turner, Daniel Mathalon, Adrian Preda, Godfrey Pearlson, Tulay Adali, and Vince D Calhoun. Higher dimensional meta-state analysis reveals reduced resting fmri connectivity dynamism in schizophrenia patients. *PloS one*, 11(3):e0149849, 2016.
- [20] Barnaly Rashid, Eswar Damaraju, Godfrey D Pearlson, and Vince D Calhoun. Dynamic connectivity states estimated from resting fmri identify differences among schizophrenia, bipolar disorder, and healthy control subjects. *Frontiers in human neuroscience*, 8:897, 2014.
- [21] Barnaly Rashid, Mohammad R Arbabshirani, Eswar Damaraju, Mustafa S Cetin, Robyn Miller, Godfrey D Pearlson, and Vince D Calhoun. Classification of schizophrenia and bipolar patients using static and dynamic resting-state fmri brain connectivity. *Neuroimage*, 134:645–657, 2016.
- [22] Maryam Falahpour, Wesley K Thompson, Angela E Abbott, Afroz Jahedi, Mark E Mulvey, Michael Datko, Thomas T Liu, and Ralph-Axel Müller. Underconnected, but not broken: Dynamic functional connectivity mri shows underconnectivity in autism is linked to increased intra-individual variability across time. *Brain connectivity*, 6(5):403–414, 2016.
- [23] N de Lacy, D Doherty, BH King, S Rachakonda, and VD Calhoun. Disruption to control network function correlates with altered dynamic connectivity in the wider autism spectrum. *NeuroImage: Clinical*, 15:513–524, 2017.
- [24] Murat Demirtaş, Cristian Tornador, Carles Falcon, Marina López-Solà, Rosa Hernández-Ribas, Jesus Pujol, Jose M Menchon, Petra Ritter, Narcis Cardoner, Carles Soriano-Mas, and Gustavo Deco. Dynamic functional connectivity reveals altered variability in functional connectivity among patients with major depressive disorder. *Human brain mapping*, 37(8):2918–2930, 2016.
- [25] Xiang Li, Dajiang Zhu, Xi Jiang, Changfeng Jin, Xin Zhang, Lei Guo, Jing Zhang, Xiaoping Hu, Lingjiang Li, and Tianming Liu. Dynamic functional connectomics signatures for characterization and differentiation of ptsd patients. *Human brain mapping*, 35(4):1761–1778, 2014.
- [26] Jinli Ou, Zhichao Lian, Li Xie, Xiang Li, Peng Wang, Yun Hao, Dajiang Zhu, Rongxin Jiang, Yufeng Wang, Yaowu Chen, Jing Zhang, and Tianming Liu. Atomic dynamic functional interaction patterns for characterization of adhd. *Human brain mapping*, 35(10):5262–5278, 2014.
- [27] Xiaobo Chen, Han Zhang, Yue Gao, Chong-Yaw Wee, Gang Li, and Dinggang Shen. High-order resting-state functional connectivity network for mci classification. *Human brain mapping*, 37(9):3282–3296, 2016.
- [28] Madiha J Jafri, Godfrey D Pearlson, Michael Stevens, and Vince D Calhoun. A method for functional network connectivity among spatially independent resting-state components in schizophrenia. *Neuroimage*, 39(4):1666–1681, 2008.
- [29] Neil D Woodward, Baxter Rogers, and Stephan Heckers. Functional resting-state networks are differentially affected in schizophrenia. *Schizophrenia research*, 130(1):86–93, 2011.
- [30] Abigail G Garrity, Godfrey D Pearlson, Kristen McKiernan, Dan Lloyd, Kent A Kiehl, and Vince D Calhoun. Aberrant "default mode" functional connectivity in schizophrenia. *American journal of psychiatry*, 164(3):450–457, 2007.
- [31] Eswar Damaraju, Elena A Allen, Aysenil Belger, JM Ford, S McEwen, DH Mathalon, BA Mueller, GD Pearlson, SG Potkin, A Preda, et al. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clinical*, 5:298–308, 2014.
- [32] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, Peter Boesiger, et al. Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine*, 42(5):952–962, 1999.

- [33] Markus Barth, Felix Breuer, Peter J Koopmans, David G Norris, and Benedikt A Poser. Simultaneous multislice (sms) imaging techniques. *Magnetic resonance in medicine*, 75(1):63–81, 2016.
- [34] David A Feinberg and Kawin Setsompop. Ultra-fast mri of the human brain with simultaneous multi-slice imaging. *Journal of magnetic resonance*, 229:90–100, 2013.
- [35] M. G. Preti and D. Van de Ville. Dynamics of functional connectivity at high spatial resolution reveal long-range interactions and fine-scale organization. *Scientific Reports*, 7:12773, 2017.
- [36] Unal Sakouglu, Godfrey D Pearlson, Kent A Kiehl, Y Michelle Wang, Andrew M Michael, and Vince D Calhoun. A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 23(5-6):351–366, 2010.
- [37] David T Jones, Prashanthi Vemuri, Matthew C Murphy, Jeffrey L Gunter, Matthew L Senjem, Mary M Machulda, Scott A Przybelski, Brian E Gregg, Kejal Kantarci, David S Knopman, Bradley F Boeve, Ronald C Petersen, and Clifford R Jack Jr. Non-stationarity in the “resting brain’s” modular architecture. *PloS one*, 7(6):e39731, 2012.
- [38] Daniel A Handwerker, Vinai Roopchansingh, Javier Gonzalez-Castillo, and Peter A Bandettini. Periodic changes in fmri connectivity. *Neuroimage*, 63(3):1712–1719, 2012.
- [39] Elena A Allen, Eswar Damaraju, Sergey M Plis, Erik B Erhardt, Tom Eichele, and Vince D Calhoun. Tracking whole-brain connectivity dynamics in the resting state. *Cerebral cortex*, 24(3):663–676, 2014.
- [40] Christian F Beckmann and Stephen M Smith. Tensorial extensions of independent component analysis for multisubject fmri analysis. *Neuroimage*, 25(1):294–311, 2005.
- [41] R Matthew Hutchison, Thilo Womelsdorf, Elena A Allen, Peter A Bandettini, Vince D Calhoun, Maurizio Corbetta, Stefania Della Penna, Jeff H Duyn, Gary H Glover, Javier Gonzalez-Castillo, et al. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, 80:360–378, 2013.
- [42] V. D. Calhoun, R. Miller, G. Pearlson, and T. Adali. The chronnectome: time-varying connectivity networks as the next frontier in fmri data discovery. *Neuron*, 84(2):262–74, 2014.
- [43] Vince D Calhoun and Tulay Adali. Time-varying brain connectivity in fmri data: whole-brain data-driven approaches for capturing and characterizing dynamic states. *IEEE Signal Processing Magazine*, 33(3):52–66, 2016.
- [44] Shella Keilholz, Cesar Caballero-Gaudes, Peter Bandettini, Gustavo Deco, and Vince Calhoun. Time-resolved resting-state functional magnetic resonance imaging analysis: Current status, challenges, and new directions. *Brain connectivity*, 7(8):465–481, 2017.
- [45] Vince D Calhoun and Nina de Lacy. Ten key observations on the analysis of resting-state functional mr imaging data using independent component analysis. *Neuroimaging Clinics*, 27(4):561–579, 2017.
- [46] Vesa Kiviniemi, Tapani Vire, Jukka Remes, Ahmed Abou Elseoud, Tuomo Starck, Osmo Tervonen, and Juha Nikkinen. A sliding time-window ica reveals spatial variability of the default mode network in time. *Brain connectivity*, 1(4):339–347, 2011.
- [47] Catie Chang and Gary H Glover. Time–frequency dynamics of resting-state brain connectivity measured with fmri. *Neuroimage*, 50(1):81–98, 2010.
- [48] Diego Vidaurre, Stephen M Smith, and Mark W Woolrich. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):12827–12832, 2017.
- [49] J Nathan Kutz, Steven L Brunton, Bingni W Brunton, and Joshua L Proctor. *Dynamic mode decomposition: data-driven modeling of complex systems*, volume 149. SIAM, 2016.

- [50] Bingni W Brunton, Lise A Johnson, Jeffrey G Ojemann, and J Nathan Kutz. Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *Journal of neuroscience methods*, 258:1–15, 2016.
- [51] Clarence W Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S Henningson. Spectral analysis of nonlinear flows. *Journal of fluid mechanics*, 641:115–127, 2009.
- [52] Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.
- [53] Jonathan H Tu, Clarence W Rowley, Dirk M Luchtenburg, Steven L Brunton, and J Nathan Kutz. On dynamic mode decomposition: theory and applications. *arXiv preprint arXiv:1312.0041*, 2013.
- [54] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, and Wu-Minn HCP Consortium. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.
- [55] Stephen M Smith, Christian F Beckmann, Jesper Andersson, Edward J Auerbach, Janine Bijsterbosch, Gwenaëlle Douaud, Eugene Duff, David A Feinberg, Ludovica Griffanti, Michael P Harms, Michael Kelly, Timothy Laumann, Karla Miller, Steen Moeller, Steve Petersen, Jonathan Power, Gholamreza Salimi-Khorshidi, Abraham Z. Snyder, An T. Vu, Mark W. Woolrich, Junqian Xu, Essa Yacoub, Kamil Ugurbil, David C. Van Essen, and Matthew F. Glasser. Resting-state fmri in the human connectome project. *Neuroimage*, 80:144–168, 2013.
- [56] Matthew F Glasser, Stamatis N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, David C Van Essen, Mark Jenkinson, and Wu-Minn HCP Consortium. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.
- [57] Evan M. Gordon, Timothy O. Laumann, Babatunde Adeyemo, Jeremy F. Huckins, William M. Kelley, and Steven E. Petersen. Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cerebral Cortex*, 26(1):288–303, 2016.
- [58] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed 2/19/18].
- [59] Ziv Bar-Joseph, David K. Gifford, and Tommi S. Jaakkola. Fast optimal leaf ordering for hierarchical clustering. *Bioinformatics*, 17:S22–S29, 2001.
- [60] Daniel Mullner. Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378*, 2011.
- [61] Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.
- [62] J Nathan Kutz, Xing Fu, and Steven L Brunton. Multiresolution dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 15(2):713–735, 2016.
- [63] Travis Askham and J Nathan Kutz. Variable projection methods for an optimized dynamic mode decomposition. *arXiv preprint arXiv:1704.02343*, 2017.
- [64] Sotiropoulos SN, Jbabdi S, Xu J, Andersson JL, Moeller S, Auerbach EJ, Glasser MF, Hernandez M, Sapiro G, Jenkinson M, Feinberg DA, Yacoub E, Lenglet C, Van Essen DC, Ugurbil K, Behrens TE, and WU-Minn HCP Consortium. Advances in diffusion mri acquisition and processing in the human connectome project. *Neuroimage*, 80:125–43, 2013.
- [65] Salimi-Khorshidi G, Douaud G, Beckmann CF, Glasser MF, Griffanti L, and Smith SM. Automatic denoising of functional mri data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage*, 15(90):449–468, 2013.

- [66] Christian F Beckmann, Marilena DeLuca, Joseph T Devlin, and Stephen M Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1457):1001–1013, 2005.
- [67] Pearlson GD Pekar JJ. Calhoun VD, Adali T. A method for making group inferences from functional a method for making group inferences from functional mri data using independent component analysis. *Human brain mapping*, 14(3):140–151, 2001.
- [68] Wernick M. N. Hansen L. K. Anderson J. Lukic, A. S. and S. C Strother. A spatially robust ica algorithm for multiple fmri data sets. *IEEE ISBI*, pages 839–842, 2001.
- [69] N. Filippini C.F. Beckmann, C.E. Mackay and S.M. Smith. Group comparison of resting-state fmri data using multi-subject ica and dual regression. *OHBM*, 2009.
- [70] Nicola Filippini, Bradley J MacIntosh, Morgan G Hough, Guy M Goodwin, Giovanni B Frisoni, Stephen M Smith, Paul M Matthews, Christian F Beckmann, and Clare E Mackay. Distinct patterns of brain activity in young carriers of the apoe- ϵ 4 allele. *Proceedings of the National Academy of Sciences*, 106(17):7209–7214, 2009.
- [71] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [72] Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.
- [73] Stephen M Smith, Aapo Hyvärinen, Gael Varoquaux, Karla L Miller, and Christian F Beckmann. Group-pca for very large fmri datasets. *Neuroimage*, 101:738–749, 2014.

Appendix A Comparing Clusters Across Scans

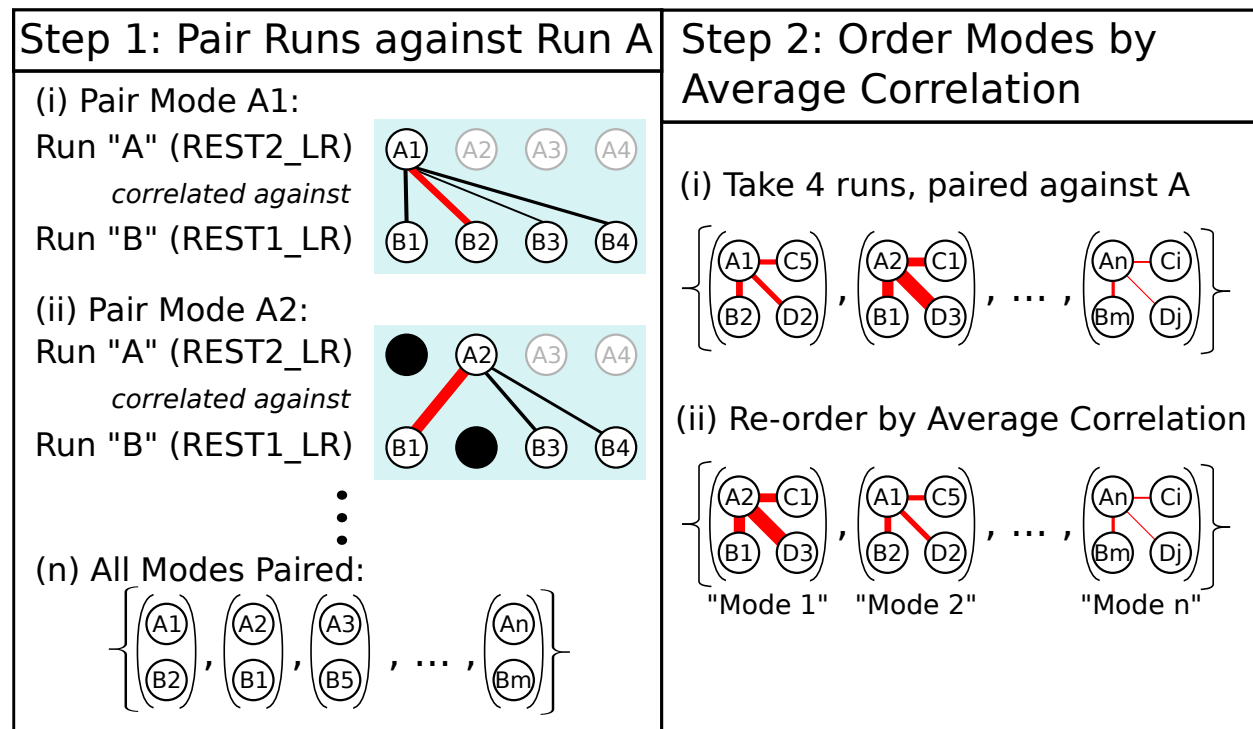


Figure 1: Procedure for pairing in ordering modes between different runs. When we process the four runs independently with identical parameters, they result in a different number of clusters, and thus a different number of corresponding averaged modes. Here we label the runs A, B, C, D; in the data, these refer to REST2_LR, REST1_LR, REST2_LR and REST1_LR, which yielded 36, 48, 49 and 42 modes respectively. We choose the run with the least number of modes (REST2_LR) as Run A, and match its modes to a unique mode from each of the other runs. **Step 1:** Iterate through the modes of Run A and match them (without replacement) to the most highly spatially correlated mode from Run B, until all modes are paired. Repeat this process for Runs C and D. **Step 2:** Calculate the average correlation between the A modes and the paired B/C/D modes, and re-order the modes by their average correlation. The modes which have the highest average correlation (i.e. which appear, on average, most similar between runs) we dub the “most consistent”. Spatial correlations of the B/C/D modes against the paired A modes in our results are plotted in Figure 2.







Figure 2: The averaged modes from each of the clusters from all four runs were paired together as described in Figure 1. As described there, these cluster-average modes are ordered by the average spatial correlation with those from the run REST2_LR. In the main manuscript, we visualize and analyze the dynamics of the top 6 modes, though this cut-off is arbitrary: similar spatial consistency appears in higher-numbered modes. For completeness, we include visualizations of the averaged modes of all clusters for all four runs on the following pages.

Appendix A.1 Averaged Modes from All Clusters







Appendix B Cluster Frequency Content

In the previous section, we paired clusters extracted from four different runs and then averaged the spatial modes in each of the clusters. Spatially, the average mode of each cluster appears quite similar across the four different runs. DMD extracts a temporal frequency to accompany each spatial mode, and thus each cluster also has a distribution of frequencies. The distribution of frequencies in each cluster in each of the four runs is shown in Figure 1.

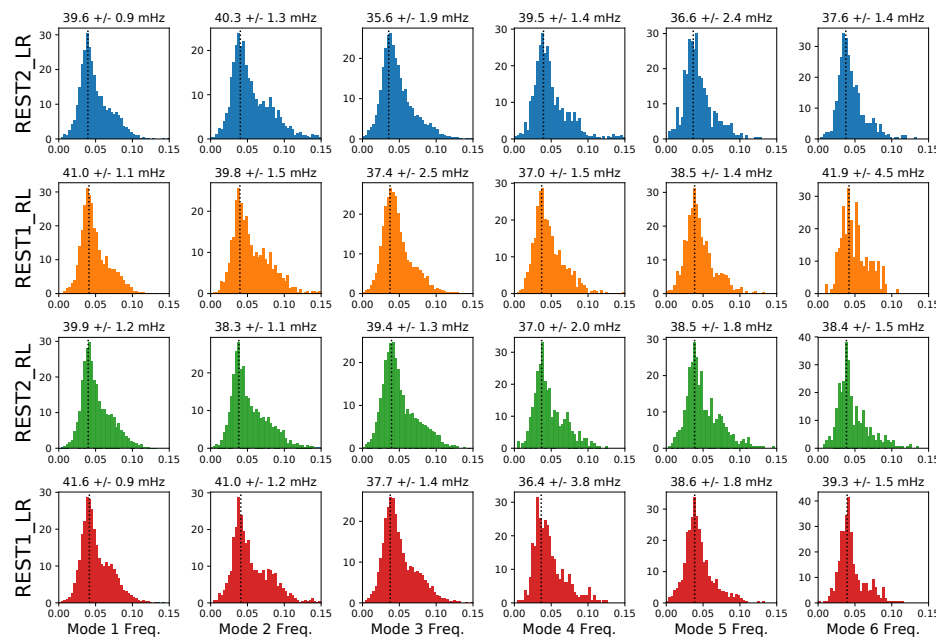


Figure 1: Frequency distributions for the modes constituting each of the six clusters in each of the four runs. Distributions appear approximately bimodal with a maximum-likelihood (peak) frequency at approximately 0.04Hz. The peak frequency (and its uncertainty) are approximated as described in this section, and are shown for each cluster.

The frequency distributions in Figure 1 are approximately bimodal; they are clearly not Gaussian, but appear to be well-approximated by the sum of two Gaussians, the first with a central frequency of $\sim 0.04\text{Hz}$ and the second, shorter Gaussian with a central frequency of $\sim 0.075\text{Hz}$. It would be un-informative for these distributions to simply compute and compare the means. Instead, we fit a Gaussian Mixture Model to the frequency data using the `scikit-learn` module `mixture`. We select the optimal number of components using the Akaike information criterion, and take the maximum likelihood of the fit distribution. We perform this fitting procedure 100 times for each cluster on randomly-selected subsets of 75% of the data, and take the mean and standard deviation of these 100 trials as our estimate of the peak frequency and its uncertainty. This estimate of the peak frequency and its uncertainty are plotted for each mode in Figure 2

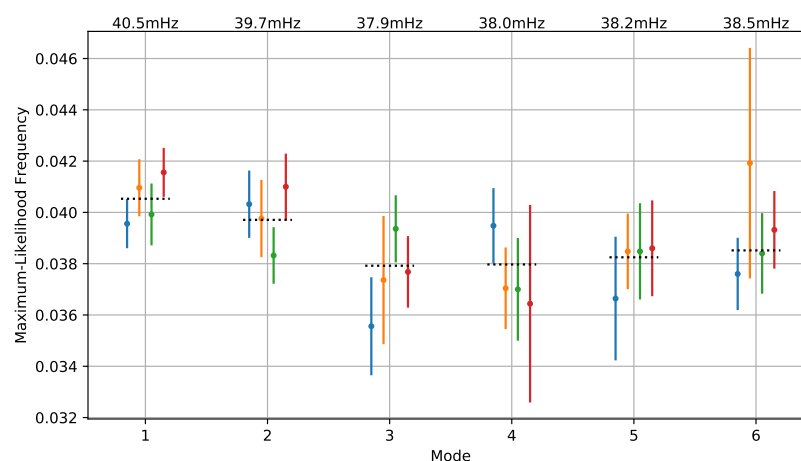


Figure 2: Peak frequencies inferred from the frequency distributions in Figure 1. For each mode, the peak was estimated for 100 random subsets of 75% of the data, and the resulting mean and standard deviation are plotted here. The dashed black lines indicate the mean frequencies of each cluster across the four runs. In short, the estimated peak frequencies are too imprecise to yield definitive conclusions, but these preliminary results suggest that the frequency content of different modes may vary reproducibly.

The distributions in Figure 1 are broad, and the estimates of peak frequency in Figure 2 are subsequently imprecise. This makes it difficult to make definitive statements about how the frequency content of the clusters are reproducibly varying between different scans. However, this preliminary analysis suggests possible reproducible variation: for example, Modes 1 and 2 appear to reliably have a peak frequency of close to 0.040Hz, whereas Modes 3-5 have a slightly lower peak frequency of 0.038Hz.

These preliminary analyses leave several open questions: does each network (corresponding to each cluster) have a distinct and reproducible frequency distribution? Does the peak frequency of each network vary from individual to individual, and could this individual variation in RSN frequency serve as a useful biomarker? In this paper we use the “exact DMD” algorithm for computing DMD, which is computationally efficient but yields a biased and imprecise estimate of the frequency. This has led to methods such as Optimized DMD ¹, which considerably improves the precision of the frequency estimate but is much more computationally expensive. Future work should consider the use of these more expensive methods, as they may be able to characterize RSN frequency to a level of precision such that RSN frequency variations could be characterized down to the level of single subjects.

¹T. Askham and J.N. Kutz. “Variable projection methods for an optimized dynamic mode decomposition.” arXiv preprint arXiv:1704.02343, 2017.

Appendix C Singular Value Spectrum

In calculating windowed DMD, we need to select the number of modes to extract from each window. Throughout the paper we have chosen $n = 8$ modes to be extracted from each of our 32-frame windows. The appropriateness of this choice can be investigated by looking at the distribution of singular values. In this section, we look at the singular values from all the windows from a single scan (Subject 102513, run REST1_LR). DMD is performed over 32-frame windows, slid over a 1200-timestep scan in increments of 4-frames, leading to 293 windows. DMD is thus performed 293 times, including 293 SVD computations. The distributions of singular values from each of these computations is shown in Figure C1.




Figure C1: Distributions of singular values over the 293 windows of a single scan (Subject 102513, run REST1_LR). Our cutoff includes the first 8 modes (indicated by the dotted red line), which sums to only $\sim 40\%$ of the energy as there is a large amount of noise. We roughly approximate a “noise floor” for singular values by taking the average minimum/maximum singular values of Modes 20-30, yielding the range indicated by the black dotted lines in the figure. Figure C2 shows a zoomed version of this figure, to better display the distributions close to our $n = 8$ cutoff.




Figure C2: The same as Figure C1, zoomed in around the $n = 8$ cutoff. After this cutoff, singular values begin to dip into the “noise floor” range. This suggests that $n = 8$ is a reasonable choice, but also suggests that similar nearby choices of n should yield similar results, as modes in this range carry only a small (though potentially non-negligible) amount of energy. Indeed, we find that our results are not particularly sensitive to the precise choice of n .