



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

James M. Kunert-Graf^{1*}, Kristian M. Eschenburg², David J. Galas¹, J. Nathan Kutz³, Swati D. Rane⁴ and Bingni W. Brunton⁵

¹ Pacific Northwest Research Institute, Seattle, WA, United States, ² Department of Bioengineering, University of Washington, Seattle, WA, United States, ³ Department of Applied Math, University of Washington, Seattle, WA, United States, ⁴ Department of Radiology, University of Washington, Seattle, WA, United States, ⁵ Department of Biology, University of Washington, Seattle, WA, United States

OPEN ACCESS

Edited by:

Dezhong Yao,
University of Electronic Science and
Technology of China, China

Reviewed by:

Tuo Zhang,
Northwestern Polytechnical University,
China
Junfeng Sun,
Shanghai Jiao Tong University, China
Yin Tian,
Chongqing University of Posts and
Telecommunications, China

*Correspondence:

James M. Kunert-Graf
jkunert@pnri.org

Received: 04 April 2019

Accepted: 11 October 2019

Published: 31 October 2019

Citation:

Kunert-Graf JM, Eschenburg KM,
Galas DJ, Kutz JN, Rane SD and
Brunton BW (2019) Extracting
Reproducible Time-Resolved Resting
State Networks Using Dynamic Mode
Decomposition.
Front. Comput. Neurosci. 13:75.
doi: 10.3389/fncom.2019.00075

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 min. 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. After validating the method on a synthetic dataset, we analyze 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.

Keywords: resting state network (RSN), dynamic mode decomposition (DMD), RS-fMRI, human connectome project (HCP), individualized networks, network dynamics

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 (Raichle et al., 2001; Damoiseaux et al., 2006; Fox and Raichle, 2007). 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 min) (Lowe et al., 1998; Van Den Heuvel and Pol, 2010). 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) (Beckmann et al., 2005) 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 (Van Den Heuvel and Pol, 2010). The assumption that brain states are static for many minutes has the effect of averaging over the spatial and temporal variability of networks (Hutchison et al., 2013b).

Although the spatial structure of these RSN patterns is similar throughout the population, their exact structure and dynamics in time vary considerably between individuals (Song et al., 2008; van den Heuvel et al., 2009; Adelstein et al., 2011; Wei et al., 2011). Further, even within individuals, the modes reconfigure dynamically on timescales ranging from seconds to years (Honey et al., 2009; Shehzad et al., 2009; Chang and Glover, 2010; Meindl et al., 2010; Van Dijk et al., 2010; Bassett et al., 2011). 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 (Damaraju et al., 2014; Ma et al., 2014; Rashid et al., 2014, 2016; Miller et al., 2016), bipolar disorder (Rashid et al., 2014, 2016), autism (Falahpour et al., 2016; de Lacy et al., 2017), depression (Demirtaş et al., 2016), post-traumatic stress disorder (Li et al., 2014), attention deficit/hyperactivity disorder (Ou et al., 2014), and mild cognitive impairment (Chen et al., 2016). 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 (Garritty et al., 2007; Jafri et al., 2008; Woodward et al., 2011). Patients with schizophrenia also show abnormalities in dynamic dwell times in large-scale RSNs (Damaraju et al., 2014).

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 (Pruessmann et al., 1999; Feinberg and Setsompop, 2013; Barth et al., 2016). 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 (Chang and Glover, 2010; Sakouglu et al., 2010; Handwerker et al., 2012; Jones et al., 2012; Hutchison et al., 2013b; Allen et al., 2014; Preti and de Ville, 2017) or ICA (Beckmann and Smith, 2005; Hutchison et al., 2013a). Windows are then shifted in time and the measures recalculated, leading to methods for characterizing


dynamic connectivity (Calhoun et al., 2014; Calhoun and Adali, 2016; Keilholz et al., 2017; Preti and de Ville, 2017). These most common RSN extraction methods make no explicit assumptions about the intrinsic temporal structure of the fMRI data (Kiviniemi et al., 2011; Calhoun and de Lacy, 2017). 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 (Chang and Glover, 2010). 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 (Vidaurre et al., 2017).

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 (Brunton et al., 2016; Kutz et al., 2016a) ideally suited to extract coherent modes from rs-fMRI data. Similar to ICA, DMD decomposes a signal into constituent coupled spatiotemporal modes (Rowley et al., 2009; Schmid, 2010; Tu et al., 2013). 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. DMD is computed within a short sliding window of data, and the spatial components of the DMD, known as DMD modes, represent coherent spatial correlations (**Figure 1A**). We first consider group-DMD (gDMD), where hierarchical clustering discovers stereotypical modes present within the full population (**Figure 1B**). We also apply DMD independently to single scans (sDMD), as illustrated in **Figure 1C**, which allows us to extract subject-level modes and their corresponding dynamics. To validate this method, we consider a synthetic dataset consisting of two dynamic, intermittently active, partially overlapping modes. gDMD outperforms both traditional and windowed ICA at extracting these modes from noisy data (**Figure 2**), and sDMD extracts the correct temporal dynamics to a temporal resolution matching the length of the sliding window (**Figure 3**).

We then analyze rs-fMRI data from 120 individuals in the Human Connectome Project (HCP) dataset (Glasser et al., 2013; Smith et al., 2013; Van Essen et al., 2013). The gDMD modes of the rs-fMRI data are strongly clustered, with clusters and sub-clusters corresponding directly to canonical RSNs. On a single-subject level, sDMD 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,

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 s windows and slide in 3 s steps over each 15 min scan. **(B)** Group-DMD (gDMD) clusters the full set of modes from all 120 subjects to extract group-averaged RSNs. **(C)** Subject-level DMD (sDMD) clusters only the subset of modes from a single subject, which yields both individualized RSNs and their corresponding dynamics.




FIGURE 2 | In a synthetic dataset with two “cat” modes and two “chicken” modes, windowed DMD performs better than group ICA and windowed ICA to extract the true underlying modes. **(A)** The synthetic cat and chicken modes. **(B)** The simulated timeseries data comprises intermittently pulsing cat/chicken modes plus white noise. **(C)** A comparison of the modes extracted from the simulated data using group ICA, sliding window ICA, and sliding window DMD. In (i), traditional group ICA does not separate the modes, which are neither fully spatially or temporally independent. In (ii), ICA modes computed in sliding windows and averaged over windows distinguishes each mode, but the results are noisy and contains shadows from the other modes. In (iii), DMD modes computed over sliding windows yield modes that are significantly closer to the true modes. In (ii) and (iii), the absolute value of each mode was taken before averaging, so that the colormap is different from that of (i).

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.

2. RESULTS

Here we present results using our method combining dynamic mode decomposition (DMD) and unsupervised clustering to extract resting state networks (RSNs) from our synthetic dataset and the Human Connectome Project (HCP) data. We show results at both the group level (gDMD) and the subject level (sDMD), as summarized in **Figure 1**. The sDMD results extract reproducible, time-resolved RSNs, and their transition dynamics


for individual scans. Details of our approach are described in the section 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. Synthetic Data

We first consider a synthetic dataset containing “cat” modes and “chicken” modes, shown in **Figure 2A**. Each mode has two variants, and we conduct a total of four simulations corresponding to the different combinations of Cat A/B and Chicken A/B. In each simulation, the chicken mode and cat mode are each intermittently pulsing in the presence of a high level of white noise, as depicted in **Figure 2B**.

We use three methods to extract the underlying modes from these noisy simulations, with the results shown in **Figure 2C**. We first consider Group ICA, which concatenates the four simulations into a single dataset from which the ICA modes are computed. In this synthetic case wherein the modes are neither fully spatially nor temporally independent, this approach lumps the underlying patterns together and fails to separate them spatiotemporally.

To resolve modes in time, we instead utilize a windowed approach: we compute ICA/DMD modes from short sliding windows of data within each individual scan. Specifically, we compute modes within 32-frame windows sliding by 4 frames, producing a total over 293 overlapping windows per each 1,200-frame scan. We average together subsets of the resultant modes to calculate the optimal reconstruction of the original underlying patterns. The ICA/DMD approaches shown here are identical except for the one line of code in which each technique is used to extract modes from a window of data. In addition to DMD



mode extraction being considerably faster than ICA [taking (1.7 ± 0.1) s per scan for windowed DMD and (7.0 ± 1.0) s per scan for windowed ICA], **Figure 2C** shows that windowed DMD modes yields a cleaner result than windowed ICA, the results of which are noisier and contain shadows from other modes. This windowed DMD technique forms the basis of both our sDMD and gDMD approaches.

Rather than correlating modes against a known pattern, the gDMD/sDMD approaches use unsupervised clustering to group together similar modes from different windows. This yields the time dynamics of each mode, since we know the time windows in which each mode does or does not appear. In this synthetic data, we can compare these inferred time dynamics against the ground truth, as in **Figure 3A**. The inferred dynamics match the ground truth, though they are smoothed over the time-scale of the window such that more rapid transitions are not captured. We can further use these inferred timeseries to estimate the transitions between states of the simulation, as in **Figure 3B**. The true state dynamics are governed by a Markov network which switches between only the cat mode being active, only the chicken mode being active, and both modes being active simultaneously. The window-smoothing effect results in an undercounting of rapid cat↔chicken transitions, which further highlights that the

resolution of the inferred dynamics is limited by the window size; otherwise, the inferred transition probabilities are close to the underlying ground truth.


2.2. rs-fMRI Data: 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 1,200 frames acquired at a temporal resolution of ~ 0.72 s. The data had been pre-processed with the HCP minimal pre-processing pipeline Glasser et al. (2013) 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.

As we did for the synthetic data, we apply DMD to short sliding windows within each individual scan. In the results shown here (**Figure 4**), we use windows of 32 frames sliding by 4 frames, producing a total over 293 overlapping windows per each 1,200-frame scan. In time, this corresponds to ~ 23 s windows sliding in ~ 3 s increments. We truncate to the first 8 modes from each window. Our results presented here are relatively robust to the number of modes calculated within each window, as described further in Appendix C (**Supplementary Material**).

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 (Bar-Joseph et al., 2001; Mullner, 2011) (as implemented in the `cluster.hierarchy` package of SciPy 1.0.0; Jones et al., 2001). 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 **Figure 4A**, where the modes have been reordered so that this matrix is strongly block diagonal. In other words, groups of modes cluster naturally. **Figure 4B** shows the correlation of each DMD mode to each of a set of reference RSNs (Gordon et al., 2016). 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 **Figure 4C**. 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`; Jones et al., 2001) of 0.955, a relatively coarse grouping of the modes into 13 clusters whose averages are shown in **Figure 4D**. 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 **Figure 4D**. 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 (**Supplementary Material**).

2.3. 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. **Figure 5** shows the results from hierarchically clustering modes from a single 15 min scan. Similarly to the group results, this similarity matrix is strongly block diagonal, and the blocks correspond to different canonical RSNs (**Figure 5A**, bottom, correlates these modes with the parcellation of Gordon et al., 2016). Note that we use a slightly lower threshold on the cophenetic distance than in **Figure 4** (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 **Figure 5B**. 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.

2.4. 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 1,200 individuals then uses dual regression to adapt each group mode to the heterogeneous structure seen within a particular scan.

Figure 6 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 (i.e., the DMN from Gordon et al., 2016). 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 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 **Figure 6A**, 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. **Figure 6B** shows 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.

2.5. 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 each cluster is active, and how often pairs of clusters are active together (either simultaneously, or with one cluster 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 clusters with similar averaged 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 clusters. 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.



Clusters were paired between the four independent scans based upon the spatial similarity of their averaged modes; this matching procedure, as well as visualizations of all extracted average modes, is detailed in Appendix A (Supplementary Material). Figure 7A shows the six most highly correlated clusters. Figure 7B shows the occupancy matrix O_{ij} of these 6 clusters, where the value in the i th row and j th column corresponds to how often clusters i and j are simultaneously active. The diagonal O_{ii} shows how often each cluster i is active, and the matrix is symmetric. Next, Figure 7C shows a transfer matrix T_{ij} , which contains the probability that, if cluster i is active, then cluster j will be active 30 s later. This matrix not symmetric, and it encodes rich dynamical information that can be interpreted as a dynamic network model, as shown schematically in Figure 7D.

We next use these same 6 cluster 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 cluster was broken up into how many of its modes are from each individual scan, and Figure 8 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 Figure 8 reports the median correlation across all scan comparisons and the corresponding p -value. For five of the six clusters, occupancy results from different scans of the same individual are positively correlated, with the strongest and most significant correlations observed for the most active clusters. Note that the high p -value of Cluster 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 the activity of that RSN 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 use DMD to extract coherent spatiotemporal modes from short, sliding windows of data, and underlying patterns are revealed by unsupervised clustering. We validate this approach on a synthetic dataset, in which we find that this windowed DMD approach separates the underlying spatiotemporal modes more cleanly than windowed ICA, and that it correctly infers the underlying dynamics down to the time-scale of the sliding window length.

We then demonstrate the application of gDMD and sDMD to rs-fMRI using data from 120 subjects from the Human Connectome Project (HCP). 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 **Figure 4**) but have drawn no conclusions from them beyond noting that they are plausible compared to the known literature (i.e., generally < 0.1 Hz). However, the preliminary analyses in Appendix B (**Supplementary Material**) suggest that networks may have distinct and reproducible frequency content which could be characterized by future studies.

Another aspect of the DMD modes which we have not explored fully here is the phase information in each mode and how it relates to resting state network organization. The phase information is particularly interesting given that DMD modes are allowed to spatially overlap. Future studies should consider if there is a systematic difference in the phases of overlapping regions, and whether this is related to the RSN dynamics (e.g., the network in **Figure 7**). If the activation of the overlapping region was phase-lagged behind the rest of the mode, and if this was followed by the activation of the overlapping RSN, then these overlaps could potentially be interpretable as mediating regions the transition between the modes.

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 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 [62]) 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, as it distinguishes intrinsically-driven dynamics (i.e., downstream excitement of modes by other modes) from externally-driven dynamics (i.e., those driven by the external task). Rather than the simple correlations captured by the networks extracted in the current paper, this future development should allow for the extraction of networks which capture the causal relationships between RSNs. Another approach, multi-resolution DMD (mrDMD; Kutz et al., 2016b) 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 (Askham and Kutz, 2017) is an implementation of the DMD algorithm which is more computationally expensive than Exact DMD (the algorithm we use here), 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 (**Supplementary Material**).

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 (Smith et al., 2013; Van Essen et al., 2013). The HCP acquired four separate resting state scans using a multiband pulse sequence (multiband factor = 8, FA = 52°), each scan having 1,200 time points (14.4 min, TR: 720 ms, TE: 33 ms, FOV: 208 × 180 mm, Matrix: 104 × 90, with 72 slices; Glasser et al., 2013). 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 (Sotiropoulos et al., 2013). 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 (Glasser et al., 2013). In addition to gradient distortion correction, this preprocessing includes fMRI denoising using FIX (Salimi-Khorshidi et al., 2013), 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 cortical surface. These FreeSurfer surface meshes were upsampled to ~164 k vertices, and subsequently downsampled to ~32 k vertices, resulting in an average vertex spacing of roughly 2 mm at the cortical midthickness level. The resting state time series were then mapped onto this low-resolution 32 k vertex mesh. This surface-mapped data consists of 91,282 “grayordinates” (~60 k surface points mapping the cortical surfaces, along with ~30 k 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 section 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. Synthetic Data Simulations

4.2.1. Synthetic Data Modes

The synthetic data X_{ij} at time t_k consists of a sum of the intermittently pulsing images of cat i and chicken k with noise:

$$X_{ij}(t_k) = a_c \sin(2\pi f_c t + \phi_c) \cdot C_i + a_k \sin(2\pi f_k t + \phi_k) \cdot K_j + N \quad (1)$$

where $C_i \in \{C_A, C_B\}$, $K_j \in \{K_A, K_B\}$ are the 75×50 pixel cat/chicken images. Amplitudes a_c and a_k are 1 or 0 depending on the system state. Each mode has a frequency f_c, f_k and phase ϕ_c, ϕ_k . These results use $f_c = 0.045$ and $f_k = 0.071$ cycles per timestep (if $\delta t \approx 0.73$ s, as for rs-fMRI data, these are about 0.06 and 0.1 Hz, within the appropriate range for RSNs) and $\phi_c = 1.6$ rad and $\phi_k = 0.7$ rad. N is a 75×50 matrix of uniform random noise with amplitude 5.

Each simulation consists of 1,200 timesteps, chosen to match the length of an HCP scan. The state of the system is governed by an underlying Markov network, with the possible states being $(a_c, a_k) \in \{(1, 0), (0, 1), (1, 1)\}$ (i.e., just the cat, just the chicken, or both). The transition probabilities at each timestep are shown in **Figure 3B**.

Windowed ICA, windowed DMD, and sDMD were calculated identically to how they were calculated for the rs-fMRI data, which is discussed in detail throughout the rest of the Methods. The only computation unique to the synthetic dataset took place in the mode reconstruction procedure, which yielded the results in **Figure 2C** (ii) and (iii). In principle, we could have applied the mode clustering procedure of gDMD to both methods, but this would yield a potentially ambiguous comparison for two reasons: first, it would not clarify whether the DMD modes were intrinsically cleaner, or if they were simply better suited to this particular clustering procedure; second, it would require an intensive parameter optimization to ensure that both methods were returning their optimal results. We instead use the following procedure:

1. Calculate the DMD modes (or ICA modes) from every window: $M = \{M_1, M_2, \dots\}$. To allow for the averaging of modes, take their absolute value.
2. Correlate every mode in M against a ground-truth mode C (i.e., C is one of the cat/chicken modes).
3. Re-order M in descending order of the magnitude of the correlation with C (i.e., M_1 is now the mode which most strongly resembles C , M_2 is the second strongest, etc.)
4. Average together the first k modes, choosing k to maximize the average mode's correlation with C .

Clearly, this is a supervised method which requires correlation against a reference mode, rather than the unsupervised approach we take when clustering the rs-fMRI results. However, it provides an upper bound on the quality of the averaged mode which clustering could return. It is thus representative of the

best possible outcome of using the gDMD approach using either ICA or DMD to calculate the modes in each window, thereby establishing that DMD does a better job at extracting spatiotemporally decoupled modes from the short windows of noisy data considered here.

4.2.2. Synthetic Data Dynamics

To assess whether the inferred dynamics are reasonable, we compared them to the known transition probabilities within the underlying Markov network. The reconstruction of the dynamic traces used the same analysis pipeline which we developed for analyzing single-subject rs-fMRI data, as detailed in sections 4.3, 4.5, and 4.7. In brief, we do the following:

1. Calculate the DMD modes within all sliding windows for a single subject.
2. Hierarchically cluster the modes, setting the clustering threshold so as to form two large clusters. In practice, these are seen to be the Cat and Chicken modes corresponding to that particular run. Modes within each clusters are correspondingly labeled as cat modes or chicken modes.
3. For each timepoint, tally how many of its corresponding windows contain a cat and/or chicken mode. This generates the dynamic traces as seen in **Figure 3**.

Note that since we are using overlapping sliding windows, each timepoint belongs within multiple windows; in **Figure 3**, the cat and/or chicken modes are considered to be active at a given timepoint if any of the corresponding windows contain a cat and/or chicken mode, which has the effect of smoothing the dynamics over the time-scale of the sliding window.

The reconstruction of the dynamics thus far is identical to our method for the fMRI data, but diverges on two key points when comparing the dynamics against the ground truth: first, the fMRI transition matrix (as in **Figure 7**) compares the state at a given time t against the state a full 30 s (or 40 timesteps) later. The synthetic data analysis, however, compares the state at a given time t to the state only one timestep later. This is necessary to compare against the underlying Markov probabilities. Second, the synthetic dataset is Markovian, with precisely one active state at each timestep, where “both” states being active is an explicit state of the system; the network in **Figure 7D**, however, does not have this constraint, and should not be interpreted as a Markov network.

4.3. Dynamic Mode Decomposition

This section describes the specific DMD algorithm implemented in this manuscript (Brunton et al., 2016; Kutz et al., 2016a). 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:

$$\mathbf{X} = \begin{bmatrix} | & | & \dots & | \\ \vec{x}_1 & \vec{x}_2 & \dots & \vec{x}_{m-1} \\ | & | & \dots & | \end{bmatrix}. \quad (2)$$

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:

$$\mathbf{X}' = \begin{bmatrix} | & | & & | \\ \tilde{x}_2 & \tilde{x}_3 & \dots & \tilde{x}_m \\ | & | & & | \end{bmatrix}. \quad (3)$$

The goal of DMD is to describe the matrix $\mathbf{A}$ that best maps $\tilde{x}_k$ onto $\tilde{x}_{k+1}$, 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

$$\mathbf{X}' = \mathbf{A}\mathbf{X}. \quad (4)$$

However, this $\mathbf{A}$ matrix is a $n \times n$ square matrix; for the HCP data, $n = 91,282$ grayordinates, 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 $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^*$, which decomposes the $\mathbf{A}$ into the product of unitary matrices $\mathbf{U}, \mathbf{V}$ and diagonal matrix Σ . 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 Σ_r and a $m \times r$ matrix $\mathbf{V}_r$.

This truncation gives the optimal rank- r reconstruction of the data matrix, $\mathbf{X} \approx \mathbf{U}_r\Sigma_r\mathbf{V}_r^*$. 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 $\mathbf{X}^\dagger = \mathbf{V}_r\Sigma_r^{-1}\mathbf{U}_r^*$ and simply solve for $\mathbf{A}$:

$$\mathbf{A} \approx \mathbf{X}'\mathbf{X}^\dagger \approx \mathbf{X}'\mathbf{V}_r\Sigma_r^{-1}\mathbf{U}_r^*. \quad (5)$$

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:

$$\tilde{\mathbf{A}} \equiv \mathbf{U}_r^*\mathbf{A}\mathbf{U}_r = \mathbf{U}_r^*\mathbf{X}'\mathbf{V}\Sigma^{-1}. \quad (6)$$

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 Σ_r as in Brunton et al. (2016):

$$\hat{\mathbf{A}} \equiv \Sigma^{-1/2}\tilde{\mathbf{A}}\Sigma^{1/2}. \quad (7)$$

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

$$\hat{\mathbf{A}}\hat{\mathbf{W}} = \hat{\mathbf{W}}\Lambda, \quad (8)$$

where Λ 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:

$$\Phi \equiv \mathbf{X}'\mathbf{V}\Sigma^{-1/2}\hat{\mathbf{W}}, \quad (9)$$

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:

$$\tilde{x}_k \approx \sum_{j=1}^r b_j \lambda_j^k \vec{\phi}_j, \quad (10)$$

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:

$$\tilde{x}(t) \approx \sum_{j=1}^r b_j \exp((\gamma_j + i2\pi f_j) \cdot t) \vec{\phi}_j. \quad (11)$$

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:

$$\gamma_j = \frac{\text{real}(\ln(\lambda_j))}{\Delta t}, \quad f_j = \frac{\text{imag}(\ln(\lambda_j))}{2\pi \Delta t}, \quad (12)$$

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 $|\vec{\phi}_j|$. However, the complex-valued $\vec{\phi}_j$ also contains the relative phase between regions, which is of potential interest for future analyses. For real-valued data $\tilde{x}_k$, oscillatory modes appear as conjugate pairs $\{\vec{\phi}_j, \vec{\phi}_j^*\}$ with the same spatial magnitudes $|\vec{\phi}_j| = |\vec{\phi}_j^*|$; thus the number of unique spatial patterns extracted from a data matrix will be $\leq r$.

4.4. Independent Component Analysis

fMRI data is often processed using Independent Component Analysis (ICA) (Beckmann et al., 2005). 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:

$$\mathbf{X}^T = \mathbf{A}\mathbf{S}^T. \quad (13)$$

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 (Calhoun et al., 2001; Lukic et al., 2001), higher-order tensor decomposition methods (Beckmann and Smith, 2005), and dual regression-based approaches (Beckmann et al., 2009; Filippini et al., 2009).

In this analysis, we use the FastICA algorithm as implemented in scikit-learn 0.18 (Pedregosa et al., 2011). FastICA attempts to maximize the statistical independence of the signals by maximizing the non-Gaussianity of the projected data, as detailed in Hyvärinen and Oja (2000). 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) (Smith et al., 2014). 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.

In practice, DMD should prove faster than ICA for short windows of high-dimensional data. Consider a data matrix X with m variables and n timepoints, where we have $m \gg n$ (for example, the rs-fMRI data has $m = 59,412$ grayordinates and windows of length $n = 32$). We compute r DMD modes in each window, where $m \gg r$ (our rs-fMRI analysis uses $r = 8$). Exact DMD consists of an SVD [with time-complexity $\mathcal{O}(m^2n)$], two matrix multiplications [with $\mathcal{O}(mnr + mr^2)$], and an eigendecomposition [with $\mathcal{O}(r^3)$]. Since $m \gg n, r$, the time-complexity is dominated by the term $\mathcal{O}(m^2n)$.

By contrast, FastICA has a time-complexity of $\mathcal{O}(m(m+1)nx)$ Laparra et al. (2011), where x is the number of iterations required for convergence. In our case, this is approximately $\mathcal{O}(m^2nx)$. In other words, the time-complexity of DMD scales similarly to a single iteration of FastICA, which can lead to considerable speed-up as FastICA can take many iterations to converge (by default, the scikit-learn implementation of FastICA will allow up to 200 iterations).

4.5. Sliding Window Mode Calculations

The rs-fMRI scans analyzed each consist of 1,200 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 22.4 s windows which are translated in increments of about 2.8 s. As discussed in section 4.3, 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 Figure 4 (the modes calculated from the REST1_LR scans of 120 subjects) consists of 160,756 total modes.

4.6. 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 3,856 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.

4.7. 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 4.6. 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 **Figure 4** 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 (Jones et al., 2001). 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 **Figure 4** we use $z_t = 2.5$; the clustering procedure is fairly robust to the choice of z_t , as demonstrated in Appendix D (**Supplementary Material**)]. 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`). **Figure 4** shows the clusters which are formed using $t = 0.955$, whereas **Figure 5** 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 [the robustness of clustering to the choice of t is also explored in Appendix D (**Supplementary Material**)]. Note that in **Figures 4, 5** we have only visualized clusters containing a minimum number of modes (≥ 400 in the gDMD case of **Figure 4**, and ≥ 5 for the sDMD case of **Figure 5**).

We compare these modes against a set of canonical RSNs, using the parcellation of Gordon et al. (2016). In **Figure 5B** 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 13) strongly corresponds to the “SMmouth” network. When the average of the modes within each cluster is plotted in **Figure 4D**, it indeed has the expected spatial structure.

4.8. Subject-Level DMNs

The subject-level Default Mode Networks (DMNs) shown in **Figure 6** 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 **Figure 6B**).

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 (Filippini et al., 2009). This approach first calculates group-ICA, which is ICA performed on the temporally-concatenated full dataset of scans from 1,200 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 **Figure 6**, but has the drawback of requiring a reference mask calculated from the entire set of population data.

Figure 6B 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.

4.9. Characterizing Cluster Dynamics

Assigning time dynamics to individual clusters (as in **Figure 5B**) 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 (**Supplementary Material**), with the resulting modes visualized in **Figure 7A**. 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 30 s later.

We assess the reproducibility of our approach by comparing the occupancies of each cluster by each individual between different scans. This rests upon the assumption that if an individual shows increased activity in a particular cluster in one scan, they are then more likely to be more active in the same cluster 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 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 8**, 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 **Figure 8**. 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.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: <https://db.humanconnectome.org>.

REFERENCES

- Adelstein, J. S., Shehzad, Z., Mennes, M., DeYoung, C. G., Zuo, X.-N., Kelly, C., et al. (2011). Personality is reflected in the brain's intrinsic functional architecture. *PLoS ONE* 6:e27633. doi: 10.1371/journal.pone.0027633
- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., and Calhoun, V. D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* 24, 663–676. doi: 10.1093/cercor/bhs352
- Askham, T., and Kutz, J. N. (2017). Variable projection methods for an optimized dynamic mode decomposition. *arXiv:1704.02343*.
- Bar-Joseph, Z., Gifford, D. K., and Jaakkola, T. S. (2001). Fast optimal leaf ordering for hierarchical clustering. *Bioinformatics* 17, S22–S29. doi: 10.1093/bioinformatics/17.suppl_1.S22
- Barth, M., Breuer, F., Koopmans, P. J., Norris, D. G., and Poser, B. A. (2016). Simultaneous multislice (sms) imaging techniques. *Magn. Reson. Med.* 75, 63–81. doi: 10.1002/mrm.25897
- Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., and Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proc. Natl. Acad. Sci. U.S.A.* 108, 7641–7646. doi: 10.1073/pnas.1018985108
- Beckmann, C.F., Mackay, C.E., Filippini, N., and Smith, S. M. (2009). Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *NeuroImage* 47(Suppl. 1), S39–S41. doi: 10.1016/S1053-8119(09)71511-3

AUTHOR CONTRIBUTIONS

All authors contributed to the conceptualization of the study. JK-G, DG, JK, and BB contributed to the mathematical conceptualization. KE and SR contributed to the domain-specific conceptualization. JK-G and BB developed the methodology, in consultation with JK and DG. KE and SR curated the data. JK-G wrote the software to perform the investigation and visualization. JK-G, SR, and BB supervised the project. All authors contributed to the writing of the manuscript.

FUNDING

JK-G and DG acknowledge support from the Pacific Northwest Research Institute. JK acknowledges support from the Air Force Office of Scientific Research (FA9550-17-1-0329). BB acknowledges support from the National Science Foundation (award 1514556), the Alfred P. Sloan Foundation, and the Washington Research Foundation. SR acknowledges support from the National Institutes of Health/National Institute on Aging (K01AG055669). This manuscript has been released as a Pre-Print at <https://www.biorxiv.org/content/10.1101/343061v1> (Kunert-Graf et al., 2018).

ACKNOWLEDGMENTS

The authors thank Nikita Sakhanenko for useful discussions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncom.2019.00075/full#supplementary-material>

- Beckmann, C. F., DeLuca, M., Devlin, J. T., and Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 1001–1013. doi: 10.1098/rstb.2005.1634
- Beckmann, C. F., and Smith, S. M. (2005). Tensorial extensions of independent component analysis for multisubject fMRI analysis. *Neuroimage* 25, 294–311. doi: 10.1016/j.neuroimage.2004.10.043
- Brunton, B. W., Johnson, L. A., Ojemann, J. G., and Kutz, J. N. (2016). Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *J. Neurosci. Methods* 258, 1–15. doi: 10.1016/j.jneumeth.2015.10.010
- Calhoun, V. D., and Adali, T. (2016). Time-varying brain connectivity in fMRI data: whole-brain data-driven approaches for capturing and characterizing dynamic states. *IEEE Signal Process. Mag.* 33, 52–66. doi: 10.1109/MSP.2015.2478915
- Calhoun, V. D., Adali, T., Pearlson, G. D., and Pekar, J. J. (2001). A method for making group inferences from functional a method for making group inferences from functional mri data using independent component analysis. *Hum. Brain Mapp.* 14, 140–151. doi: 10.1002/hbm.1048
- Calhoun, V. D., and de Lacy, N. (2017). Ten key observations on the analysis of resting-state functional MR imaging data using independent component analysis. *Neuroimaging Clin.* 27, 561–579. doi: 10.1016/j.nic.2017.06.012
- Calhoun, V. D., Miller, R., Pearlson, G., and Adali, T. (2014). The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84, 262–274. doi: 10.1016/j.neuron.2014.10.015

- Chang, C., and Glover, G. H. (2010). Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage* 50, 81–98. doi: 10.1016/j.neuroimage.2009.12.011
- Chen, X., Zhang, H., Gao, Y., Wee, C.-Y., Li, G., and Shen, D. (2016). High-order resting-state functional connectivity network for MCI classification. *Hum. Brain Mapp.* 37, 3282–3296. doi: 10.1002/hbm.23240
- Damaraju, E., Allen, E. A., Belger, A., Ford, J., McEwen, S., Mathalon, D., et al. (2014). Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage* 5, 298–308. doi: 10.1016/j.neuroimage.2014.07.003
- Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S. M., et al. (2006). Consistent resting-state networks across healthy subjects. *Proc. Natl. Acad. Sci. U.S.A.* 103, 13848–13853. doi: 10.1073/pnas.0601417103
- de Lacy, N., Doherty, D., King, B., Rachakonda, S., and Calhoun, V. (2017). Disruption to control network function correlates with altered dynamic connectivity in the wider autism spectrum. *NeuroImage* 15, 513–524. doi: 10.1016/j.neuroimage.2017.05.024
- Demirtaş, M., Tornador, C., Falcon, C., López-Solà, M., Hernández-Ribas, R., Pujol, J., et al. (2016). Dynamic functional connectivity reveals altered variability in functional connectivity among patients with major depressive disorder. *Hum. Brain Mapp.* 37, 2918–2930. doi: 10.1002/hbm.23215
- Falahpour, M., Thompson, W. K., Abbott, A. E., Jahedi, A., Mulvey, M. E., Datko, M., et al. (2016). Underconnected, but not broken: dynamic functional connectivity MRI shows underconnectivity in autism is linked to increased intra-individual variability across time. *Brain Connect.* 6, 403–414. doi: 10.1089/brain.2015.0389
- Feinberg, D. A., and Setsompop, K. (2013). Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. *J. Magn. Reson.* 229, 90–100. doi: 10.1016/j.jmr.2013.02.002
- Filippini, N., MacIntosh, B. J., Hough, M. G., Goodwin, G. M., Frisoni, G. B., Smith, S. M., et al. (2009). Distinct patterns of brain activity in young carriers of the apoe-ε4 allele. *Proc. Natl. Acad. Sci. U.S.A.* 106, 7209–7214. doi: 10.1073/pnas.0811879106
- Fox, M. D., and Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711. doi: 10.1038/nrn2201
- Garrity, A. G., Pearlson, G. D., McKiernan, K., Lloyd, D., Kiehl, K. A., and Calhoun, V. D. (2007). Aberrant “default mode” functional connectivity in schizophrenia. *Am. J. Psychiatry* 164, 450–457. doi: 10.1176/ajp.2007.164.3.450
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., et al. (2013). The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124. doi: 10.1016/j.neuroimage.2013.04.127
- Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., and Petersen, S. E. (2016). Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cereb. Cortex* 26, 288–303. doi: 10.1093/cercor/bhu239
- Handwerker, D. A., Roopchansingh, V., Gonzalez-Castillo, J., and Bandettini, P. A. (2012). Periodic changes in fMRI connectivity. *Neuroimage* 63, 1712–1719. doi: 10.1016/j.neuroimage.2012.06.078
- Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.-P., Meuli, R., et al. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. U.S.A.* 106, 2035–2040. doi: 10.1073/pnas.0811168106
- Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., et al. (2013a). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* 80, 360–378. doi: 10.1016/j.neuroimage.2013.05.079
- Hutchison, R. M., Womelsdorf, T., Gati, J. S., Everling, S., and Menon, R. S. (2013b). Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. *Hum. Brain Mapp.* 34, 2154–2177. doi: 10.1002/hbm.22058
- Hyvärinen, A., and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Netw.* 13, 411–430. doi: 10.1016/S0893-6080(00)00026-5
- Jafri, M. J., Pearlson, G. D., Stevens, M., and Calhoun, V. D. (2008). A method for functional network connectivity among spatially independent resting-state components in schizophrenia. *Neuroimage* 39, 1666–1681. doi: 10.1016/j.neuroimage.2007.11.001
- Jones, D. T., Vemuri, P., Murphy, M. C., Gunter, J. L., Senjem, M. L., Machulda, M. M., et al. (2012). Non-stationarity in the “resting brain’s” modular architecture. *PLoS ONE* 7:e39731. doi: 10.1371/journal.pone.0039731
- Jones, E., Oliphant, T., and Peterson, P. (2001). *SciPy: Open Source Scientific Tools for Python*. Available online at: <https://www.scipy.org/citing.html>. (accessed February 19, 2018).
- Keilholz, S., Caballero-Gaudes, C., Bandettini, P., Deco, G., and Calhoun, V. (2017). Time-resolved resting-state functional magnetic resonance imaging analysis: current status, challenges, and new directions. *Brain Connect.* 7, 465–481. doi: 10.1089/brain.2017.0543
- Kiviniemi, V., Vire, T., Remes, J., Elseoud, A. A., Starck, T., Tervonen, O., et al. (2011). A sliding time-window ICA reveals spatial variability of the default mode network in time. *Brain Connect.* 1, 339–347. doi: 10.1089/brain.2011.0036
- Kunert-Graf, J. M., Eschenburg, K. M., Galas, D. J., Nathan Kutz, J., Rane, S. D., and Brunton, B. W. (2018). Extracting reproducible time-resolved resting state networks using dynamic mode decomposition. *bioRxiv [Preprint]*. doi: 10.1101/343061
- Kutz, J. N., Brunton, S. L., Brunton, B. W., and Proctor, J. L. (2016a). *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, Vol. 149. Philadelphia, PA: SIAM.
- Kutz, J. N., Fu, X., and Brunton, S. L. (2016b). Multiresolution dynamic mode decomposition. *SIAM J. Appl. Dyn. Syst.* 15, 713–735. doi: 10.1137/15M1023543
- Laparra, V., Camps-Valls, G., and Malo, J. (2011). Iterative gaussianization: from ICA to random rotations. *IEEE Trans. Neural Netw.* 22, 537–549. doi: 10.1109/TNN.2011.2106511
- Li, X., Zhu, D., Jiang, X., Jin, C., Zhang, X., Guo, L., et al. (2014). Dynamic functional connectomics signatures for characterization and differentiation of PTSD patients. *Hum. Brain Mapp.* 35, 1761–1778. doi: 10.1002/hbm.22290
- Lowe, M., Mock, B., and Sorenson, J. (1998). Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *Neuroimage* 7, 119–132. doi: 10.1006/nimg.1997.0315
- Lukic, A. S., Wernick, M. N., Hansen, L. K., Anderson, J., and Strother, S. C. (2001). “A spatially robust ICA algorithm for multiple fMRI data sets,” *Proceedings IEEE International Symposium on Biomedical Imaging* (Washington, DC), 839–842.
- Ma, S., Calhoun, V. D., Phlypo, R., and Adahi, T. (2014). Dynamic changes of spatial functional network connectivity in healthy individuals and schizophrenia patients using independent vector analysis. *Neuroimage* 90, 196–206. doi: 10.1016/j.neuroimage.2013.12.063
- Meindl, T., Teipel, S., Elmouden, R., Mueller, S., Koch, W., Dietrich, O., et al. (2010). Test-retest reproducibility of the default-mode network in healthy individuals. *Hum. Brain Mapp.* 31, 237–246. doi: 10.1002/hbm.20860
- Miller, R. L., Yaesoubi, M., Turner, J. A., Mathalon, D., Preda, A., Pearlson, G., et al. (2016). Higher dimensional meta-state analysis reveals reduced resting fMRI connectivity dynamism in schizophrenia patients. *PLoS ONE* 11:e0149849. doi: 10.1371/journal.pone.0149849
- Mullner, D. (2011). Modern hierarchical, agglomerative clustering algorithms. *arXiv:1109.2378*.
- Ou, J., Lian, Z., Xie, L., Li, X., Wang, P., Hao, Y., et al. (2014). Atomic dynamic interaction patterns for characterization of ADHD. *Hum. Brain Mapp.* 35, 5262–5278. doi: 10.1002/hbm.22548
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830. Available online at: <http://www.jmlr.org/papers/v12/pedregosa11a.html>
- Preti, M. G., and de Ville, D. V. (2017). Dynamics of functional connectivity at high spatial resolution reveal long-range interactions and fine-scale organization. *Sci. Rep.* 7:12773. doi: 10.1038/s41598-017-12993-1
- Pruessmann, K. P., Weiger, M., Scheidegger, M. B., and Boesiger, P. (1999). Sense: sensitivity encoding for fast MRI. *Magn. Reson. Med.* 42, 952–962. doi: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proc. Natl. Acad. Sci. U.S.A.* 98, 676–682. doi: 10.1073/pnas.98.2.676

- Rashid, B., Arbabshirani, M. R., Damaraju, E., Cetin, M. S., Miller, R., Pearlson, G. D., et al. (2016). Classification of schizophrenia and bipolar patients using static and dynamic resting-state fMRI brain connectivity. *Neuroimage* 134, 645–657. doi: 10.1016/j.neuroimage.2016.04.051
- Rashid, B., Damaraju, E., Pearlson, G. D., and Calhoun, V. D. (2014). Dynamic connectivity states estimated from resting fMRI identify differences among schizophrenia, bipolar disorder, and healthy control subjects. *Front. Hum. Neurosci.* 8:897. doi: 10.3389/fnhum.2014.00897
- Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S. (2009). Spectral analysis of nonlinear flows. *J. Fluid Mech.* 641, 115–127. doi: 10.1017/S0022112009992059
- Sakouglu, U., Pearlson, G. D., Kiehl, K. A., Wang, Y. M., Michael, A. M., and Calhoun, V. D. (2010). A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magn. Reson. Mater. Phys. Biol. Med.* 23, 351–366. doi: 10.1007/s10334-010-0197-8
- Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., and Smith, S. M. (2013). Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 15, 449–468. doi: 10.1016/j.neuroimage.2013.11.046
- Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* 656, 5–28. doi: 10.1017/S0022112010001217
- Shehzad, Z., Kelly, A. C., Reiss, P. T., Gee, D. G., Gotimer, K., Uddin, L. Q., et al. (2009). The resting brain: unconstrained yet reliable. *Cereb. Cortex* 19, 2209–2229. doi: 10.1093/cercor/bhn256
- Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., et al. (2013). Resting-state fMRI in the human connectome project. *Neuroimage* 80, 144–168. doi: 10.1016/j.neuroimage.2013.05.039
- Smith, S. M., Hyvärinen, A., Varoquaux, G., Miller, K. L., and Beckmann, C. F. (2014). Group-PCA for very large fMRI datasets. *Neuroimage* 101, 738–749. doi: 10.1016/j.neuroimage.2014.07.051
- Song, M., Zhou, Y., Li, J., Liu, Y., Tian, L., Yu, C., et al. (2008). Brain spontaneous functional connectivity and intelligence. *Neuroimage* 41, 1168–1176. doi: 10.1016/j.neuroimage.2008.02.036
- Sotiropoulos, S. N., Jbabdi, S., Xu, J., Andersson, J. L., Moeller, S., Auerbach, E. J., et al. (2013). Advances in diffusion MRI acquisition and processing in the human connectome project. *Neuroimage* 80, 125–43. doi: 10.1016/j.neuroimage.2013.05.057
- Tu, J. H., Rowley, C. W., Luchtenburg, D. M., Brunton, S. L., and Kutz, J. N. (2013). On dynamic mode decomposition: theory and applications. *arXiv:1312.0041*.
- Van Den Heuvel, M. P., and Pol, H. E. H. (2010). Exploring the brain network: a review on resting-state fMRI functional connectivity. *Eur. Neuropsychopharmacol.* 20, 519–534. doi: 10.1016/j.euroneuro.2010.03.008
- van den Heuvel, M. P., Stam, C. J., Kahn, R. S., and Pol, H. E. H. (2009). Efficiency of functional brain networks and intellectual performance. *J. Neurosci.* 29, 7619–7624. doi: 10.1523/JNEUROSCI.1443-09.2009
- Van Dijk, K. R., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., and Buckner, R. L. (2010). Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. *J. Neurophysiol.* 103, 297–321. doi: 10.1152/jn.00783.2009
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., et al. (2013). The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79. doi: 10.1016/j.neuroimage.2013.05.041
- Vidaurre, D., Smith, S. M., and Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci. U.S.A.* 114, 12827–12832. doi: 10.1073/pnas.1705120114
- Wei, L., Duan, X., Yang, Y., Liao, W., Gao, Q., Ding, J.-R., et al. (2011). The synchronization of spontaneous bold activity predicts extraversion and neuroticism. *Brain Res.* 1419, 68–75. doi: 10.1016/j.brainres.2011.08.060
- Woodward, N. D., Rogers, B., and Heckers, S. (2011). Functional resting-state networks are differentially affected in schizophrenia. *Schizophr. Res.* 130, 86–93. doi: 10.1016/j.schres.2011.03.010

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Kunert-Graf, Eschenburg, Galas, Kutz, Rane and Brunton. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.