Title

A novel hidden Markov approach to studying dynamic functional connectivity states in human neuroimaging

Authors and Affiliations

Sana Hussain¹, Jason Langley², Aaron R. Seitz³, Megan A. K. Peters^{1,4†}, Xiaoping P. Hu^{1,2†}

1. Department of Bioengineering, University of California Riverside, Riverside, CA
2. Center for Advanced Neuroimaging, University of California, Riverside, Riverside, CA
3. Department of Psychology, University of California Riverside, Riverside, CA
4. Department of Cognitive Sciences, University of California Irvine, Irvine, CA

† These authors contributed equally to this work

Sana Hussain ORCID: 0000-0003-0986-2654; email: shuss006@ucr.edu
Jason Langley ORCID: 0000-0001-9412-1265; email: jason.langley@ucr.edu
Aaron R. Seitz ORCID: 0000-0003-4936-9303; email: aseitz@ucr.edu
Megan A. K. Peters ORCID: 0000-0002-0248-0816; email: megan.peters@uci.edu
Xiaoping P. Hu ORCID: 0000-0001-7014-5983; email: xhu@engr.ucr.edu

Word count: 10,096 (excluding abstract, references, figure captions, and appendices)
Figures/Tables: 8 Figures/0 Tables (excluding appendices)
Date submitted: 02/02/22

† Correspondence To:

Xiaoping P. Hu, Ph.D.
Provost Fellow
Professor and Chair Department of Bioengineering University of California, Riverside
Materials Science and Engineering 205
Phone: (951) 827-2925
Fax: (951) 827-6416
E-mail: xhu@engr.ucr.edu

Megan A. K. Peters, Ph.D.
Assistant Professor of Cognitive Sciences University of California, Irvine
Social & Behavioral Sciences Gateway 2314
Phone: (949) 824-6692
Fax: (949) 824-2307
E-mail: megan.peters@uci.edu

Abstract

How can we better understand the underlying neural mechanisms of brain dynamics? Many groups have examined brain states as a means of identifying recurring patterns of activity or connectivity. Although there is an abundance of methods of establishing and interpreting brain states, hidden Markov models are becoming an increasingly popular choice to extract recurring patterns of intensity or connectivity in neuroimaging data. These models not only recognize spatial patterns of brain states, but also ascertain their temporal progression. An assortment of hidden Markov model instantiations has arisen for diverse purposes, and these have been applied to a variety of neuroimaging datasets. However, most of these instantiations have focused on intensity-based states, i.e. states defined by the activity levels of one or more nodes, rather than connectivity-based states, i.e. states defined by patterns of functional connectivity between nodes. The intensity-based approach is problematic if we want to understand connectivity dynamics, since there is no reason to believe that the resultant states can provide any useful information about dynamic connectivity patterns. Here we aimed to remedy this methodological challenge by introducing a new hidden Markov model approach based on identifying states defined as full functional connectivity profiles among brain regions, and applying this approach to directly extract connectivity-based states in functional magnetic resonance imaging (fMRI) data. We then empirically explore the behavior of this new model in comparison to existing approaches based on intensity-based states and summed functional connectivity states, utilizing the widely-available HCP unrelated 100 functional magnetic resonance imaging “resting state” dataset. Results show that our newly-introduced ‘full functional connectivity’ model discovered connectivity states with more distinguishable patterns than those derived from previously-employed approaches, and demonstrated clear superiority in recovering simulated connectivity-based states. These findings suggest that if our goal is to extract and interpret connectivity states in neuroimaging data, our new model can reveal more insights than previous methods, and intensity-based or summed functional connectivity-based approaches miss crucial information about the evolution of functional connectivity in the brain.

Keywords

Resting state fMRI; hidden Markov model; state patterns; neuroimaging; functional connectivity

Highlights

- Hidden Markov models can be used to investigate brain states noninvasively
- Previous models “recover” connectivity from intensity-based hidden states
- Other previous models “recover” connectivity from connectivity ‘summed’ across nodes
- Only full connectivity-based models can reveal true connectivity hidden states

1. Introduction

A useful framework to characterize the brain is to view it as a dynamical system of interacting and interchanging brain states (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019; Lurie et al. 2020). Brain states are patterns of activity levels or connectivity strengths that characterize and quantify network interactions. These states can be further split into two broad categories: states defined by the activity levels of one or more nodes (brain areas), which we call *intensity-based* states, and states defined by patterns of functional connectivity between nodes, which we call *connectivity-based* states. In contrast to intensity-based methods, connectivity-based states and their dynamics remain relatively underexplored. However, these dynamic connectivity approaches have the potential to reveal fundamental insights into complex correlative relationships among brain regions; developing and benchmarking new methods for extracting and characterizing these states is therefore of critical importance.

Outside the context of state-based analysis, examining dynamic changes in functional connectivity is often employed in noninvasive neuroimaging in humans, particularly via functional magnetic resonance imaging (fMRI). Previous methods to identify connectivity-based states range from simple to complex. On the simple end, one can calculate pairwise dynamic functional connectivity (dFC), i.e. the correlation in activity between pairs of nodes in a brain network and how such correlations change across time. Dynamic functional connectivity is typically calculated using a sliding time window approach, and has previously been used to characterize these states using fMRI to generate insight into the underlying neural mechanisms of brain dynamics (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Lurie et al. 2020) or disease mechanisms (Ou et al. 2015; Díez-Cirarda et al. 2018; Fiorenzato et al. 2019). In the sliding time window approach, information contained within a certain window length is correlated between time series of different brain regions, which is moved a certain number of time points until a general picture of functional connectivity shifts is obtained (Lurie et al. 2020); this forms the trajectory through dFC space.

One promising approach commonly used to characterize evolution of *intensity-based* brain states across time is through employing hidden Markov models (HMMs). HMMs utilize probabilistic models to determine a hidden state sequence path not directly observable in data (L. R. Rabiner 1989; Sean R. Eddy 2004; S. R. Eddy 1996; Jurafsky and Martin 2009) and have been shown to adequately characterize fluctuating dynamics of temporally changing intensity-based brain states in a data-driven manner (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; L. R. Rabiner 1989; S. R. Eddy 1996). To achieve this, HMMs employ Markov chains that infer the underlying intensity-based states where the probability of residing in any one of these states depends only on the previous state (L. R. Rabiner 1989; S. R. Eddy 1996; Sean R. Eddy 2004). HMMs uses three different algorithms (the forward algorithm, the Viterbi algorithm, and the Baum-Welch algorithm) in conjunction to find the most likely sequence of hidden states, transition probabilities, and emission probabilities based on an observable sequence of data (L. R. Rabiner 1989; L. Rabiner and Juang 1986; Jurafsky and Martin 2009). HMMs are useful to investigate brain state dynamics in fMRI datasets because (1) they do not need assumptions about the relationships among brain states (Chen et al. 2016), and (2) spatial and temporal information are inherent to the model. Because of these properties, HMMs have been used to identify latent brain states from brain signals acquired from fMRI and magnetoencephalography (Chen et al. 2016;

Vidaurre, Smith, and Woolrich 2017; Vidaurre, Abeysuriya, et al. 2018; Vidaurre, Hunt, et al. 2018; Baker et al. 2014; Eavani et al. 2013; Vidaurre et al. 2016).

Intensity-based HMMs have successfully been used to define spatial patterns of latent, intensity-based brain activity states, as well as to recognize the transitions and time spent in those states (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019; Vidaurre, Hunt, et al. 2018). To study connectivity-based states, several groups have examined covariance values extracted from such intensity-based HMMs, transforming them into Pearson correlations to create connectivity-like states to examine (Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019). However, it is unclear to what extent such transformed covariances reflect true underlying *connectivity-based states* rather than simply the connectivities that the *intensity-based* states happened to exhibit. A more direct approach seems desirable.

Thus, if we wish to characterize a trajectory through *connectivity-based* state space, defining an HMM using functional connectivity instead of signal intensity as an input seems a natural and perhaps simpler approach. Recently, Ou et al. 2015 implemented such a method where results from a dFC sliding window analysis were summed into a representative “connectivity vector” -- describing a given node’s total connectivity to all other nodes in the network -- for every time point, and then fitted with an HMM (Ou et al. 2015). This method may appear advantageous to examining the trajectory through connectivity state space because it results in HMM-derived connectivity-based states which allow analysis of temporal aspects of connectivity-based states in a diseased population (Ou et al. 2015). Critically, however, because of the summing step, this method sums over -- and thus potentially obscures -- dynamic changes in pairwise connectivity, and therefore may average over important information, leading to erroneous or less than pertinent results. For example, increased connectivity between the source node and one target node might be balanced by decreased connectivity to another node, such that no change is observed in the overall sum of connectivity. Importantly, despite the sensitivity of HMMs, interpretation of the output of a fitted HMM depends strongly on the inputs and assumptions used to develop the model. That is, the states resulting from Ou et al. 2015’s method may have drastically different results than one in which an HMM was fitted to *all* pairwise correlation values -- between *all* pairs of nodes -- resulting from the sliding window analysis (Ou et al. 2015).

Here we tackled these concerns head-on by developing and evaluating a new HMM-based method that fits all correlation values obtained from a dFC sliding window analysis. We comprehensively compared this novel *full functional connectivity* HMM (FFC HMM) to two previously-reported methods used to examine functional connectivity states in neuroimaging data: (1) a standard *intensity-based* HMM (IB HMM) (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019), and (2) the *summed functional connectivity* HMM (SFC HMM) described above (Ou et al. 2015). We fitted each of these models to a widely-available existing dataset, the Human Connectome Project Unrelated 100 (Van Essen et al. 2013) resting state functional MRI dataset. Our findings highlight the advantages of our new FFC HMM in characterizing functional connectivity states, as well as cautioning against assuming that meaningful connectivity patterns can be derived from models fitted to alternative (intensity-based) or functional connectivity inputs summed across nodes.

2. Methods

2.1 HCP dataset and networks

All analyses described below were performed on the Human Connectome Project (HCP) Unrelated 100 (a subset of the S500 release) dataset (Van Essen et al. 2013), obtained from the HCP database (<https://ida.loni.usc.edu/login.jsp>). The HCP project (Principal Investigators: Bruce Rosen, M.D., Ph.D., Martinos Center at Massachusetts General Hospital; Arthur W. Toga, Ph.D., University of Southern California, Van J. Weeden, MD, Martinos Center at Massachusetts General Hospital) is supported by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH) and the National Institute of Neurological Disorders and Stroke (NINDS). HCP is the result of efforts of co-investigators from the University of Southern California, Martinos Center for Biomedical Imaging at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota.

The portion of data used here includes 100 subjects (age = 22–36, gender = 54 female/46 male) who underwent a 14.4-minute resting state scan (repetition time = 720ms, flip angle = 52°, voxel size = 2mm³, echo time = 33ms, field of view = 208mm x 180mm). We preprocessed the data using the HCP minimally preprocessing pipeline: distortion correction, motion correction, alignment to standard space, and surface projection (Glasser et al. 2013). During early development and fitting of the HMMs described below, one subject was found to remain in one single, subject-specific state that was only visited by two other subjects for one timepoint each. Removing that subject did not affect the number of hidden states chosen (or any other parameter; data not shown), so we opted to conduct all analyses described below on the remaining 99 subjects.

Following previous work (Deshpande, Santhanam, and Hu 2011; Raichle 2011), BOLD signal was extracted from regions of interest (ROIs) from four brain networks previously associated with resting state: the default mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), and salience network (SN). The nodes comprising each network (29 nodes in total) were defined using anatomical coordinates specified in literature (Table S1). Talairach coordinates for DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and converted to Montreal Neurological Institute (MNI) coordinates (Deshpande, Santhanam, and Hu 2011; Lancaster et al. 2007; Laird, Lancaster, and Fox 2005; Brett, Johnsrude, and Owen 2002). MNI coordinates for SN were taken directly from Raichle 2011 (Raichle 2011). MNI coordinates for all ROIs can be found in **Appendix A.1 (Table A1)**. After labeling each ROI with a 5mm³ isotropic marker, the BOLD signal was extracted from each voxel and averaged across all voxels in the ROI, producing a single time series representing the behavior of the ROI as a whole. This procedure was repeated for every ROI in a network for a total of 29 ROIs (9 from DMN, 7 from FPCN, 6 from DAN, and 7 from SN) per subject (Deshpande et al. 2008; 2009; Stilla et al. 2007).

2.2 Hidden Markov models

To critically evaluate the behavior of our novel *full functional connectivity* HMM (FFC HMM), we compared it to two previously-reported methods. The differences among these methods are

defined by their inputs. That is, we compared our FFC HMM, which takes as input a time series of pairwise dFCs between all pairs of ROIs to models taking as input: (1) BOLD time series (*intensity-based* HMM; IB HMM) (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019), or (2) time series of dFC summed across all nodes that a given node is connected to (*summed functional connectivity* HMM; SFC HMM) (Ou et al. 2015). Each of these HMMs is based on the same fundamental assumptions and are detailed in Rabiner and Juang 1986, Rabiner 1989 and Jurafsky and Martin 2009 (L. R. Rabiner 1989; L. Rabiner and Juang 1986; Jurafsky and Martin 2009). Each HMM was implemented using the hmmlearn python library (Pedregosa et al. 2011) and additional details for implementation of each HMM are described in **Appendix A.2**.

Below, we describe these established approaches and their behavior in order to provide context for the methods and results of our novel FFC HMM.

2.2.1 Intensity-based HMM

The first model we compare to our novel FFC HMM approach is the standard *intensity-based* HMM (IB HMM). IB HMM has been used to extract brain states from neuroimaging data (Chen et al. 2016; Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019). Here, BOLD signals from ROIs (described in Table S1) were extracted, preprocessed, z-scored, and concatenated across subjects (**Fig. A1**). These fMRI time series were concatenated timewise across all subjects to create a matrix of size (time * # subjects) x (# ROIs) and submitted as input to the hmmlearn package to be fit with standard procedures described elsewhere (Pedregosa et al. 2011). That is, the forward and Viterbi algorithms were used in conjunction to identify the most likely sequence of hidden states given the observable BOLD signal. The Baum-Welch algorithm was then implemented to calculate the transition and emission probabilities of a given state (L. R. Rabiner 1989; L. Rabiner and Juang 1986; Jurafsky and Martin 2009).

2.2.2 Summed functional connectivity HMM

The second method we compare to our novel FFC HMM is the *summed functional connectivity* HMM (SFC HMM). Here, SFC HMM was adapted from Ou et al. 2015 (Ou et al. 2015). First, a sliding time window analysis (window length $\Delta t=36$ seconds for primary analyses) was used on the z-scored BOLD signal to obtain an ROI x ROI connectivity matrix within each time window (Ou et al. 2015) (**Fig. A2**). This generated a connectivity time series of length (# TRs - Δt) representing the dynamics of functional connectivity over time. These square connectivity matrices were summed across one dimension to create a 1 x (# ROIs) vector depicting the total overall connectedness of each ROI to all other ROIs. Repeating this for every time window provided a “summed dFC time series” containing a (# time windows * # subjects) x ROI data matrix.

We examined the behavior of the SFC HMM under different window lengths for calculating dFC. Lurie et al. 2020 postulated that a window length less than 60 seconds may be optimal (Lurie et al. 2020). Thus, for the primary analysis, a window length of 50 time points (36 seconds) was selected. However, to assess the extent to which the qualitative behavior of the model is dependent on window length, we also fit the model under varying window length with sizes of 30 time points (21.6 seconds), 40 time points (28.8 seconds), 50 time points (36

seconds), 60 time points (43.2 seconds), and 80 time points (57.6 seconds). A histogram of the distribution of R^2 values was used to evaluate the state patterns across different time windows.

2.2.3 Full functional connectivity HMM

To define our novel *full functional connectivity* HMM (FFC HMM), we aimed to remedy the shortcomings of SFC HMM. That is, while SFC HMM is an established method, the HMM in that model is fit to a summed statistic (summed connectivity values within each ROI x ROI matrix from the “dFC time series”) rather than to individual pairwise connectivity values. This results in the HMM identifying changes in this measure rather than in the raw correlations. Thus, not only is the user forced to compute the full connectivity themselves rather than the model directly outputting it, but the states themselves -- and their timecourse -- also do not respect pairwise shifts in functional connectivity that are obscured by the summation step.

To rectify this, we developed our FFC HMM approach. Our FFC HMM is fitted to all correlation values in the lower (or, equivalently, upper) triangle of the dFC matrix in every time window. Thus, FFC HMM differs from SFC HMM in that it takes as input the pairwise correlation between all pairs of ROIs, rather than the summation (i.e., total absolute) connectivity of each ROI to all other ROIs. Therefore, as before, a sliding window correlation analysis was performed on the z-scored BOLD signal with window length Δt = approximately 36 seconds (50 time points for primary analyses; see SFC HMM description), but instead of summing across one of the dimensions of the ROI x ROI matrices, the lower (or, equivalently, upper) triangle of R^2 Pearson correlation values was restructured into a $1 \times \frac{(\# ROI)^2 - (\# ROI)}{2}$ vector (**Fig. A3**). Repeating this for every time window gives a $(\# \text{ time windows}) \times \frac{(\# ROI)^2 - (\# ROI)}{2}$ data matrix for every subject containing the time series of all pairwise connectivities between all pairs of ROIs. These results were then concatenated subject-wise as before, such that the final input to FFC HMM was a $(\# \text{ time windows} * \# \text{ subjects}) \times \text{ROI data matrix}$ for every subject. As with SFC HMM, we also evaluated the model’s behavior under different lengths of time windows, ranging across 30 time points (21.6 seconds), 40 time points (28.8 seconds), 50 time points (36 seconds), and 60 time points (43.2 seconds), and 80 time points (57.6 seconds).

2.3 Preliminary model fitting and analysis

2.3.1 Determining the number of hidden states for each model

A commonality across both our FFC HMM and the previously-published approaches (IB HMM and SFC HMM) is that they are all fitted with an *a priori* defined number of hidden states. That is, the number of hidden states is not a free parameter but must be specified by the experimenter. To determine the optimal number of states for each model, we adopted the Ranking and Averaging Independent Component Analysis by Reproducibility (RAICAR) method, which uses the stability of recovered hidden state patterns to determine the appropriateness of the number of states (Chen et al. 2016; Yang et al. 2008).

We examined the stability of models with 3 to 15 hidden states. For each HMM with each number of hidden states, three sets of state patterns were obtained by three different realizations: one with uniform starting probability of residing in all states, and two with randomly assigned

starting probabilities. Because the labeling of states is arbitrary, there is no reason to believe State 1 will match across realizations even if a given number of hidden states is optimally stable. Therefore, state patterns were matched across realizations via Pearson correlations, such that e.g. State 1 from Realization 2 was relabeled as State 2 just in case the Pearson correlation between that state and State 2 from Realization 1 was higher than any other pairwise correlation. Thus, after relabeling, each state label across realizations universally corresponded to the same spatial pattern to the maximal extent possible within stability constraints. Within each state assignment, the matched state patterns were then Pearson correlated between all realizations to obtain $(\# \text{ states})! / 2! * (\# \text{ states} - 2)!$ Pearson's R^2 values, thereby determining the maximal degree of similarity between the matched patterns for the present number of hidden states being tested. Finally, these values were then averaged, sorted from largest to smallest, and plotted as a function of the number of states in the model.

We repeated this process for a range of number of hidden states, generating a dataset of pattern similarity (R^2) organized by pattern label; these maximal achievable pattern similarities were compared against a predetermined threshold of stability of 0.9. Previous groups that used the RAICAR-based method examined more ROIs (236 from Chen et al. 2016 and 162 independent components from Yang et al. 2014) and employed a stability threshold of 0.8 (Chen et al. 2016; Yang et al. 2008). As we explored only 29 ROIs, we opted to appoint a more conservative threshold: 0.9. Cases where R^2 values began to dip below this threshold indicated that this model was unstable with that number of hidden states, because the states were not matching sufficiently (Chen et al. 2016; Yang et al. 2008). This procedure is also outlined in greater detail in **Appendix A.3** and **Fig. A4**.

2.3.2 Model recovery

An important step towards validating our FFC HMM and benchmarking it against previous approaches is to verify the degree to which it can recognize ground truth by successfully recovering predefined connectivity states. Thus, we manipulated the HCP data to induce “connectivity-defined states” as follows. First, the preprocessed data extracted from our 29 ROIs were randomly permuted across time to create a noisy time series. To create “connectivity-defined” ground truth states, we replaced the scrambled data from these same time points with values drawn from a multivariate gaussian distribution using $\mu = 0$ and an ROI x ROI σ matrix consisting of 0.1 on the off-diagonals where we aimed to induce connectivity, and ones on the diagonals. For example, to induce a “DMN (9 ROIs) connectivity state”, a 9 x 9 σ matrix was made with 0.1 on all off diagonals and 1 on the diagonals, which was then used to seed a time series of simulated BOLD response within the time points for which this connectivity state was to be induced; the result was then used to replace the temporally scrambled data. Connectivity state patterns and a hidden state sequence were acquired from FFC HMM outputs and compared against outputs from IB HMM and SFC HMM to establish to what extent each model is capable of recovering ground truth.

2.3.3 Model robustness to data size

We are also interested to discover whether FFC HMM requires large amounts of data to produce robust results, in comparison to previously established approaches. Therefore, to determine the ability of each model to recover stable states regardless of the length of the resting state scan, we

performed an additional reproducibility analysis. To evaluate whether each model could stably discover the same states using a significantly reduced dataset, we fit the IB, SFC, and FFC HMMs to only the first half of the resting state scan for all subjects (that is, the first 7.2 minutes of a 14.4 minute scan) and computed the similarity between the discovered states for this ‘half’ dataset and the full dataset used in the main analyses. Similarity was computed according to the same process used in the RAICAR-based stability analysis (**Methods Section 2.3.1**), except that states were matched between the half- and full-dataset HMMs (fit with the uniform initial conditions) rather than between different realizations. See **Appendix A.4** for more details.

(Intensity-defined state induction and model validation was also performed but is not included in the main text for brevity; see **Appendix B** for details.)

2.4 Analysis of model outputs

2.4.1 Connectivity state pattern analysis

The primary outcome metric of interest is the recovered connectivity states. The defining feature of FFC HMM in comparison to previous methods is that it is fit to full connectivity profiles, so we would hope that it would recover more distinct full connectivity profiles than either IB HMM or SFC HMM.

To examine this behavior, we analyzed the obtained connectivity states for all three HMMs at their most stable number of states. The extraction method for these connectivity states differed by model, but after extraction, analyses were similar across models. Here we detail those extraction processes for each HMM, as well as the analytics employed.

Full connectivity state patterns (“connectivity state patterns” or “connectivity states” from this point on, to differentiate them from summed connectivity vectors) depict the correlation strength between all pairs of nodes within a given state (i.e., each state consists of a 29x29 matrix of Pearson R^2 values). For IB HMM, connectivity states corresponding to each intensity state were acquired by mathematically transforming the covariance matrices outputted by the model fitting procedure into Pearson correlation values, as done previously (Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019; Eavani et al. 2013). For SFC HMM, the model’s direct outputs are vectors of the mean state patterns describing mean summed correlation values representing global nodal strength during a particular time window which cannot be directly ‘unpacked’ into full pairwise connectivities among all nodes; therefore, we defined connectivity state patterns for SFC HMM by averaging connectivity matrices across time points when its Viterbi path labeled a state to be active. For FFC HMM, the connectivity state patterns are directly outputted from the model corresponding to the $1 \times \frac{(\# ROI)^2 - (\# ROI)}{2}$ correlation vector inputted for every time window, which are then reformatted back into a symmetric ROI x ROI matrix to constitute the connectivity states; we also additionally calculated the connectivity state patterns from the Viterbi paths for the purpose of direct comparison with the other two models.

All connectivity analyses were performed on *differential functional connectivity maps* from the IB, SFC, and FFC HMM connectivity states: connectivity matrices that “highlight” the unique functional connectivity characteristics of each state (Stevner et al. 2019). SFC and FFC HMM

state differential functional connectivity maps were computed using Eq. 1 where X_i is the original raw functional connectivity matrix for state i , H_i gives the differential functional connectivity maps of X_i , and j gives all state assignments (given the number of states determined for the model) excluding the value of i (Stevner et al. 2019).

$$H_i = X_i - \left(\frac{1}{7} \sum_{j \neq i} X_j \right) \quad (1)$$

Note that by this definition, since we examined state differential functional connectivity maps rather than the raw states themselves, the values in the connectivity states represent connectivity levels relative to baseline, not the Pearson R^2 values themselves. That is, for all state differential functional connectivity maps, negative values are associated with below baseline correlations, not anticorrelations.

A critical question is to what extent FFC HMM's recovered connectivity states are adequately captured by previous methods (IB HMM and SFC HMM). Therefore, we examined the similarities among IB, SFC, and FFC HMMs by comparing the connectivity states they revealed, especially to see whether SFC states matched FFC states. If SFC HMM is able to capture the same states as FFC HMM (because it uses the same connectivities found via the sliding window approach), then the additional complexity of FFC HMM is unnecessary and SFC HMM is adequate for examining trajectories through connectivity state space as well as connectivity states themselves. The state patterns can be compared by Pearson correlation (Chen et al. 2016) or computing the Euclidean distance (Guggenmos, Sterzer, and Cichy 2018) between all pairs of states, similar to our procedure for the stability analysis (**Methods Section 2.3.1**). For direct comparison with our stability analysis to discover the ideal number of hidden states for each model, we used Pearson correlations and sought to discover a 'stability threshold' (similar to the threshold of 0.9 used in **Methods Section 2.3.1**) that would result in unique, one-to-one pairwise matching between states recovered by two different models.

Finally, we also want to see whether FFC HMM can recover summed connectivity vectors as discovered by SFC HMM, which would add to its utility and avoid the need to fit two models to ask questions about both mean connectivity profiles as well as full connectivity states. Therefore, we similarly computed the summed functional connectivity state patterns corresponding to all connectivity state differential functional connectivity maps for IB, SFC, and FFC HMMs. To create the summed functional vectors for IB and FFC HMMs, we summed the full 29x29 matrix of Pearson R^2 values across one of the dimensions; for SFC, we utilized the summed connectivity vectors directly outputted by the model fitting procedure.

2.4.2 Viterbi path analysis

In addition to examining the connectivity states themselves, we are also interested in examining trajectories through state space and how FFC HMM might differ in such dynamics from the two previous methods. The Viterbi path, or hidden state sequence, is directly outputted from all three HMMs. We therefore examined a number of characteristics of the Viterbi path outputs from FFC HMM in comparison to IB HMM and SFC HMM, including switching rate, proportion of time spent in each state, the average duration of a state, and fractional occupancy correlation (Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019).

3. Results

3.1 Preliminary model fitting

3.1.1 Determining the number of hidden states for each model

The RAICAR-based stability analysis results for IB, SFC, and FFC HMMs are shown in **Fig. 1**, limited to models with 6 to 10 hidden states for clarity (full results for models with 3-15 hidden states are shown in **Figs. A5** and **A6** and do not change the conclusions presented in the main text). IB HMM and SFC HMM were both most stable at 8 hidden states (**Figs. 1a & 1b**). We therefore examined the stability of FFC HMM with 8 states, and determined it to be higher than a 9-state model (**Fig. 1c**); thus, we assigned FFC HMM to have 8 states as well.

Setting all models to have the same number of hidden states allowed for a level comparison of local and global analyses across all HMMs, setting FFC HMM on equal footing to the two established models in terms of complexity of outputs. However, we note that an identical number of hidden states between any two HMMs may not always occur and users of any of these HMM approaches should not assume that the number of hidden states that is appropriate for one model is appropriate for another. Specifically, there is no mathematical or conceptual basis to suggest that FFC HMM should have the same number of hidden states as SFC HMM or IB HMM especially (nor that SFC HMM and IB HMM should share the same number of states). The fact that we determined 8 states to be best for all models in the present analysis is therefore incidental, and may have occurred because we focused on a relatively small number of ROIs, which greatly narrowed the model complexity.

Figure 1. RAICAR-based stability analysis results for (A) IB HMM, (B) SFC HMM, and (C) FFC HMM. Eight hidden states were selected as the best size for all HMMs.

3.1.2 Model recovery

Because FFC HMM is a more complex model than the previous approaches, it is important to verify that our model fitting procedures are able to recover ground truth, i.e. that the states outputted by each model are likely to reflect real variation in the dataset and not noise. To validate our procedures, we therefore simulated connectivity-based “states” and used identical data preparation and model fitting procedures to discover the extent to which such states could be effectively recovered by our approach; see **Methods Section 2.3.2** for details.

When examining these induced connectivity states, we found that FFC HMM recovered simulated state patterns quite cleanly (**Fig. 2d**), with mean correlation (R^2) between true simulated and recovered Viterbi paths across subjects of 0.5738 ± 0.1301 . FFC HMM was in fact more precise in its state recovery than SFC HMM even to the naked eye (**Fig. 2c**) as well as quantitatively (mean R^2 across all subjects between simulated and recovered Viterbi paths of 0.3337 ± 0.1650). Unsurprisingly, IB HMM failed to adequately recover simulated connectivity states (**Fig. 2b**; mean R^2 between simulated and recovered Viterbi paths of 0.1741 ± 0.1569), although it appeared to recognize when two networks were “turned on” in conjunction. This may have occurred because there was more connectivity for the model to recognize: both within- and between-network connectivity.

FFC HMM also showed strong performance in recovering the actual connectivity states themselves. A paired t-test on the Fisher-z transformed R^2 values between SFC and FFC HMMs indexing how similar each model’s recovered connectivity states were to the induced states showed that FFC HMM was significantly better at recovering induced states than SFC HMM ($t(98) = 12.8745$, $p = 8.6087e-23$). Temporal blurring from the sliding time window analysis employed by both connectivity-based HMMs may explain the low correlation values between the simulated and recovered state sequences from SFC and FFC HMMs.

Figure 2. Verification of HMM connectivity-based states. **(A)** The artificially induced state sequence depicted which networks exhibited slightly increased within- and/or between-network connectivity. Outputted state sequences from **(B)** IB HMM, **(C)** SFC HMM, and **(D)** FFC HMM when connectivity states were induced. FFC HMM recovers simulated states better than the other two models; see main text for details.

3.2.3 Model robustness to data size

Because FFC HMM is a more complex model than either IB HMM or SFC HMM – that is, it must fit more complex full connectivity patterns rather than summed connectivity vectors or intensity-based states – we also need to assess its robustness to the amount of data available to train the model. We also want to compare this to the volume of data necessary to adequately train the comparison models, IB HMM and SFC HMM. To determine the degree to which the states recovered by each of the IB, SFC, and FFC HMMs were dependent on the exact length of the input data, we repeated all fitting procedures on only the first half of the resting state scan (first 7.2 minutes; see **Methods Section 2.3.3**). When we halved the dataset, aligned the connectivity states by their similarity, and then computed their overall maximal similarity, we found that the connectivity state patterns largely remained the same for FFC HMM and SFC HMM in particular, but (again unsurprisingly) not for IB HMM (**Fig. 3**). Therefore, it appears that a resting state scan as short as 7.2 minutes may be sufficient to recover the connectivity states discovered by FFC HMM in a scan twice that long.

Figure 3. Pearson correlation values (R^2) between connectivity state patterns identified from whole (14.4 minutes) and half (7.2 minutes) of the HCP Unrelated 100 fMRI dataset for **(A)** IB HMM, **(B)** SFC HMM, and **(C)** FFC HMM.

3.2 Analysis of model outputs

With 8 states selected as the best number of hidden states for FFC HMM and the comparison models (IB and SFC HMMs), we next proceeded to evaluate the output metrics: state analysis, time window length analysis, and Viterbi path analyses including switching rate, proportion of

time spent in each state, the average duration of a state, and fractional occupancy correlation. States from each model are distinguished with subscripts corresponding to the HMM from which they stem, i.e., $S1_{FFC}$ corresponds to State 1 from FFC HMM. In the following, we present results from IB and SFC HMMs first to provide context for the differences in behavior exhibited by our FFC HMM.

3.2.1 Connectivity state pattern analysis

First, we examine IB HMM as a baseline, which we do not expect to recover connectivity states well. Unsurprisingly, the IB HMM connectivity state differential functional connectivity maps stemming from the output covariance matrices did not display strong distinguishing patterns either within or between the states even when state differential functional connectivity maps were assessed (top row, **Fig. 4**). However, there were a few slight deviations from mean connectivity overall. For example, DAN in $S7_{IB}$ appears to be relatively disconnected from all other networks and associated with activated DMN, FPCN, DAN, and SN. We also observed that ROIs within DMN and DAN appear to have slightly above average connectivity in $S4_{IB}$ and are linked to an attention state where FPCN and DAN show the highest levels of activation. Finally, $S8_{IB}$ shows a relative disconnection between ROIs in FPCN and DAN while DAN exhibits slightly elevated within-network connectivity and is paired with all four networks' deactivation. Nevertheless, these results show that a model not trained on connectivity states is, expectedly, not able to recover connectivity states.

The critical benchmark for our FFC HMM is therefore SFC HMM's behavior. Can a model fitted to summed connectivity vectors nevertheless adequately recover full connectivity profiles? As described above (**Methods Section 2.2.2**), the states for SFC HMM were found by averaging the ROI x ROI connectivity matrices wherever the Viterbi path determined a state to be active, because we cannot recover full connectivity profiles directly from the summed connectivity vectors outputted by the model. (For reference, the summed vectors against the connectivity state differential functional connectivity maps of dimensions ROI x ROI are also plotted in **Fig. 4**.)

SFC HMM extracted distinct functional connectivity profiles across all 8 states (middle row, **Fig. 4**). From visual examination, we can see that $S1_{SFC}$ showed below-baseline correlations among all networks, particularly within and between DAN and SN, indicating a disconnect within the attentional system. Elevated correlations within and between all networks except DAN were observed in $S2_{SFC}$ and DAN does not interact with the rest of the networks within this state. Slightly higher than average correlation values were seen within $S3_{SFC}$ with the exception of several DMN exhibiting below-average connectivity with the rest of the networks. The strongest connectivity within and between the attentional networks was seen in $S4_{SFC}$. $S5_{SFC}$ had a disconnect among DMN, FPCN, and SN. In $S6_{SFC}$, below baseline connectivity was seen within DAN and between DAN and the other networks; SN exhibited similar behavior. These properties indicate a disconnect within and between these attentional networks. $S7_{SFC}$ showed a slight disconnect between DMN and all other networks indicating a disconnect in the resting state networks. Elevated connectivity within DMN and between DMN and all other networks was seen in $S8_{SFC}$.

Figure 4. State pattern differential functional connectivity maps for SFC HMM (top row), FFC HMM (middle row), and IB HMM (bottom row). The summed connectivity vectors (summed across one dimension) are displayed below each state. The summed values for SFC were directly outputted from the model while those for IB and FFC were calculated as described in **Methods Section 2.4.1**.

3.2.1.a Connectivity state comparisons

The final question, then, is to what extent the connectivity states recovered by FFC HMM differ from those found via SFC HMM. If FFC HMM's recovered states look quite a lot like SFC HMM's, then the additional complexity and computing power required to fit FFC HMM may not be warranted. However, once again we can see differences between FFC HMM and SFC HMM (as well as with IB HMM) just by eye (bottom row, **Fig. 4**). $S1_{\text{FFC}}$ exhibited no distinguishable state differential functional connectivity maps since all connectivity values were around baseline while $S2_{\text{FFC}}$ showed slightly higher than baseline connectivity within and between all networks except for DAN. Slightly above average correlations within and between all networks were seen in $S3_{\text{FFC}}$. $S4_{\text{FFC}}$ showed all networks to have above average correlations with one another and $S5_{\text{FFC}}$ to have slightly above baseline correlations between DAN and all other networks. $S6_{\text{FFC}}$ exhibited below baseline correlations between DAN and all other networks as well as between SN and all other networks. DMN was disconnected from all other networks and SN and DAN have slightly above average correlation within and between each other in $S7_{\text{FFC}}$. Finally, reduced correlations between DAN and all other networks were seen in $S8_{\text{FFC}}$.

Let us now engage in a more quantitative comparison between states recovered by each model. The most important comparison for examining FFC HMM's behavior is to compare it directly to SFC HMM's recovered states. To evaluate such potential similarities, we directly compared the connectivity-based states output by SFC HMM and FFC HMM to each other in two ways. First,

we computed the pairwise Pearson correlations among all pairs of states (**Fig. 5a**). A one-to-one match in states would be illustrated with one large correlation coefficient (one orange/yellow) square and seven small correlations (seven green/blue squares) in each row. Yet this phenomenon is not observed, as it appears there are several states that might ‘match’ across the two models from visual inspection.

To quantitatively assess the degree to which there might be one-to-one state matching across FFC HMM and SFC HMM, we took inspiration from the RAICAR analysis described above and sought to identify whether there was a threshold at which there would be a one-to-one state matching across all eight states found by both models. (Recall that in the RAICAR analysis, **Section 3.1.1**, we used a threshold of 0.9.) In this analysis, we are looking for cases where a particular correlation threshold leads to exactly one SFC HMM state matching each of the FFC HMM states. We examined whether any possible correlation threshold in the range of 0-1 in steps of 0.05 would lead to exactly one SFC HMM state matching each FFC HMM state by counting the number of SFC states that exceeded each possible threshold. Visually, this accounts to counting how many ‘squares’ in each row of **Fig. 5a** surpass a particular value, with the goal being exactly one for all rows. This quantitative analysis confirmed the visual inspection results: There was no threshold for which all – or even most – FFC HMM states could achieve a unique mapping with exactly one SFC HMM state (**Fig. 5d**). This finding supports the interpretation that our novel FFC HMM recovers functional connectivity states that are distinct from those recovered by SFC HMM.

Figure 5. Pairwise comparisons between connectivity states discovered by each HMM show that each model recovered eight unique states. A one-to-one match between two states recovered by two different models would appear as a single orange/yellow square (high Pearson correlation) among seven green/blue squares (low Pearson correlation) in those two states' row or column combination. However, (A) FFC HMM's recovered states showed no unique correspondences based on similarity to those recovered by SFC HMM by visual inspection, and (D) no stability threshold (see the RAICAR analysis to discover number of hidden states, **Sections 2.3.1 & 3.1.1**) can lead to any semblance of a one-to-one match between FFC HMM's recovered states and those recovered by SFC HMM. Similar results were found for pairwise comparisons between FFC HMM and IB HMM states (panels B and E) and between SFC HMM and IB HMM states (panels C and F).

For completeness, we also repeated this analysis for comparisons between FFC HMM and IB HMM (**Figs. 5b & 5e**), as well as comparing SFC HMM and IB HMM (**Figs. 5c & 5f**). Results from these pairwise comparisons mirror those from the critical FFC vs SFC comparison: All three models appear to discover unique connectivity states, as there is no visual or quantitative correspondence between the states discovered by each model. Note that the comparisons with IB HMM are particularly informative, as they confirm that both FFC HMM and SFC HMM recovered states that reflected changes in connectivity patterns (in SFC HMM's case, summed connectivity vectors) and were robust to fluctuations in overall BOLD signal magnitude (i.e., the connectivity-based models did not end up simply discovering 'connectivity' states based on fluctuations in average intensity). See **Appendix B** for fuller discussion of the intensity-based states recovered by IB HMM, and how they compare to intensity-based states derived from SFC and FFC HMMs.

3.2.1.b Time window analysis (connectivity-based HMMs)

To conclude that there are consistent differences between FFC HMM's recovered connectivity states and those recovered by previous approaches – SFC HMM in particular – we need to evaluate to what extent the recovered connectivity states themselves are robust to choices made early in the data processing pipeline. One choice that has the potential to greatly impact the states recovered by FFC HMM and SFC HMM is the selection of the window length to use in the sliding-window computation of dynamic functional connectivity (Lurie et al. 2020).

To evaluate the impact of window length used to calculate dynamic functional connectivity on the connectivity states recovered by FFC and SFC HMMs, we therefore examined how connectivity states would change as window length varied from 30 to 80 time points (36-57.6 seconds). Importantly, neither SFC nor FFC HMMs showed a strong effect of such a change in window length. FFC state differential functional connectivity maps showed minimal differences across window lengths. All states matched to their counterparts with $R^2 \geq \sim 0.5$ with 1 exception: matching $S3_{\text{FFC}} 40\text{tp}$ to $S3_{\text{FFC}} 80\text{tp}$ ($R^2 = 0.3469$). Likewise, all SFC state differential functional connectivity maps matched their counterparts across window sizes with Pearson $R^2 \geq 0.5$ except for two instances: (1) correlating $S4_{\text{SFC}} 30$ time points (tp) to $S4_{\text{SFC}} 80\text{tp}$ ($R^2 = 0.3084$), and (2) $S3_{\text{SFC}}$ of any time length to 80tp (tp=30, $R^2 = -0.3587$; tp=40, $R^2 = -0.2588$; tp=50, $R^2 = -0.3062$;

tp=60, $R^2 = -0.2563$). In fact, although the difference is slight, SFC states may have shown more variability across different time windows than FFC states because of the summing factor, i.e. the definition of connectivity in SFC as the sum of connectivity from one node to all other nodes. As seen in **Figs. 6** and **7**, poor pattern matching occurred when correlating state differential functional connectivity maps to differential functional connectivity maps from a window size of (80 tp or ~60 seconds).

These results show that a window size containing less than approximately 60 seconds of data is preferred for these two connectivity-based HMMs, consistent with the ideas Lurie et al. 2020 presented (Lurie et al. 2020). The histograms of connectivity values across the state differential functional connectivity maps (far right columns, **Figs. 6** and **7**) support this interpretation, since the distributions of connectivity values became less separated as the window size increased; this reduction in distributional separability suggests that the state patterns' individual distinctiveness -- i.e., lack of resemblance to their counterparts -- was larger in smaller window sizes and was reduced as window length increased. This was particularly noticeable in the 80tp window size, where the distributions of connectivity values appear closest together, meaning the state pattern differential functional connectivity maps are minimally distinct from one another.

Thus, the window length used in the primary analyses presented above (50 time points, or ~36 seconds) produces similarly unique states as other reasonable window lengths. Together, these analyses confirm that FFC HMM recovers connectivity states that are distinct from those recovered by SFC HMM (and IB HMM), and that the differences between FFC HMM's connectivity states and those from other models are not trivially due to somewhat arbitrary choices about window length for calculating dynamic functional connectivity measures.

Figure 6. SFC HMM connectivity state pattern differential functional connectivity maps for

window sizes of 30tp, 40tp, 50tp, 60tp, and 80tp used in the sliding window correlation analysis. The histograms (rightmost column) display the overall distribution of R^2 values of $S1_{\text{SFC}}-S8_{\text{SFC}}$ differential functional connectivity maps for each window size. As window size increases, the histograms become less separated, indicating that the states become more similar in pattern as window size increases. These results show that any window size containing less than ~60 seconds of data (60tp or less) produces similarly distinct connectivity profiles for SFC HMM.

Figure 7. FFC HMM connectivity state pattern differential functional connectivity maps for window sizes of 30tp, 40tp, 50tp, 60tp, and 80tp used in the sliding window correlation analysis. The histograms (rightmost column) display the overall distribution of R^2 values of $S1_{\text{FFC}}-S8_{\text{FFC}}$ differential functional connectivity maps for each window size. As window size increases, the histograms become less separated, indicating that the states become more similar in pattern as window size increases. These results show that a window size containing less than ~60 seconds of data (60tp or less) produces similarly distinct connectivity profiles for FFC HMM, and that differences between FFC HMM's connectivity states and those from other models cannot trivially be explained by arbitrary choices about window size.

3.2.2 Viterbi path analysis

Our final analyses sought to evaluate whether FFC HMM's recovered trajectory through connectivity state space was particular to this model, or could be recovered through fitting a simpler previously-reported model. These questions can be answered through examining each model's Viterbi path across the whole dataset's time series, which indicate which state was active at every time point for every subject.

The Viterbi paths can be visualized by assigning each state a color and plotting them for every person as a function of time point (TR in the fMRI time series, IB HMM, **Fig. 8a**) or of time window (anchored on the first TR of the time window, SFC HMM [**Fig. 8b**] and FFC HMM [**Fig. 8c**]) to show which states are active at each time point. These visualizations show that, compared to IB HMM, both connectivity-based HMMs appear ‘smoothed’ over time, i.e. transition among states more slowly; this is especially dramatic for FFC HMM. Autocorrelation and temporal discrepancy in input resolution between the intensity- versus connectivity-based HMMs -- i.e., IB HMM had a temporal resolution equal to that of the fMRI TR, while SFC and FFC HMMs had an effective sampling frequency on the order of 1 sample per 36 seconds -- likely contributing to this smoothing.

Interestingly, this temporal ‘smoothing’ appears even more pronounced for FFC HMM than for SFC HMM. We suspect that FFC HMM fitting so many more components than SFC HMM could have played a role in this model behavior. That is, compared to SFC HMM, FFC HMM fitted a more holistic picture of the information in a sliding window analysis and therefore necessitated fitting an extremely large number of connectivity values per time window (406 connectivity values). While the FFC HMM was stable (the Viterbi paths converged to the same solution for every initialization, indicating that 406 components per time window was not an irrational amount of data to fit), this level of complexity was significantly higher than SFC’s 29 components per time window. This increased complexity may have caused FFC HMM to take longer in recognizing switches from one state to another: FFC HMM was required to find significant changes in 14 times as many connectivity components as SFC HMM. Fitting many values at once in combination with a very noisy dataset could have caused the noise to outweigh the signal in FFC HMM, thereby resulting in fewer state transitions. Moreover, the severe smoothing may have resulted from constraining FFC HMM to bear the same number of hidden states as SFC HMM -- a choice we made for the sake of meaningful comparisons between SFC and FFC. It is possible that because FFC HMM had more variables that it could have stably identified more states. However, a comprehensive exploration of these and other potential explanations for this behavior is beyond the scope of the present project, so we leave these explorations to future work.

As a final check, we also confirmed that these Viterbi paths were robust to different realizations, i.e. different initial conditions for model fitting. Recall that, although the stability analyses determined 8 states was ideal for all three models, state assignment (e.g. labeling a state as “S1” vs “S8”) was initially arbitrary across all HMMs. Nevertheless, following state labeling alignment across initializations (see **Methods Section 2.3.1**), we observed that the Viterbi path was reproducible across different realizations of the HMMs ($R^2 \geq \sim 0.84$) for all models tested. That is, for all initializations, all models recognized the same connectivity states to be prevalent during the same time windows and the same switches between states.

Figure 8. Viterbi Paths for (A) IB HMM, (B) SFC HMM, and (C) FFC HMM. The Viterbi paths for the SFC and FFC HMMs are much “smoother” (i.e., more spread out in time) than those of IB HMM. Within the connectivity-based HMMs, FFC HMM’s Viterbi path exhibits fewer and less frequent switches than SFC HMM’s, which may have occurred because of the selected number of hidden states or the total number of components fitted. See main text for detailed discussion.

4. Discussion

Here, we introduced the ‘full functional connectivity’ hidden Markov model (FFC HMM) and investigated its ability to extract meaningful functional connectivity states in resting state fMRI data. Using the HCP Unrelated 100, we fitted the FFC HMM to a sliding window of functional connectivity from 29 ROIs across four networks (DMN, FPCN, DAN, and SN) and compared the latent functional connectivity states that it discovered to those identified by a SFC HMM and an IB HMM. We found that FFC HMM’s discovered full connectivity states were starkly different from those recovered by SFC and IB HMMs. Moreover, and importantly, FFC HMM recovered simulated connectivity-based states much more faithfully than either IB or SFC HMM. This suggests that any full, pairwise functional connectivity profiles that are recovered from either IB or SFC HMMs must be interpreted with caution, as they may not resemble the true functional connectivity states exhibited in the data. We also observed that FFC HMM (as well as SFC HMM) showed strong robustness to choice of window size in the sliding window calculation of dynamic functional connectivity, as well as a moderate insensitivity to the length of the data (50% vs 100% of available resting-state data). Together, these results support the conclusion that FFC HMM is the appropriate choice when the goal of a study is to examine and interpret hidden functional connectivity profiles in functional neuroimaging data.

4.1 Relation to previous literature

Our results drive home the notation that, perhaps unsurprisingly, the input to a model matters when interpreting its output. More specifically, if we wish to discover and interpret full functional connectivity states in resting-state fMRI data, we should not assume that connectivity states derived from intensity-based state discovery methods (i.e., IB HMM; (Vidaurre, Smith, and Woolrich 2017; Stevner et al. 2019; Eavani et al. 2013)) are in any way equivalent to the true functional connectivity states in a dataset. Moreover, even previous methods using summations of functional connectivity vectors (SFC HMM; (Ou et al. 2015)) cannot be used to recover full functional connectivity profiles, as they are insensitive to changes in each node's functional connectivity profile and instead only seek overall changes in average connectivity for each node examined. Recall that this previous summed functional connectivity vector fitting method did have the appeal that the summing step greatly reduces the computing power necessary to consider all connectivity values in fitting the HMM; however, here we showed that the increase in necessary computing is worth the challenge, because FFC HMM's recovered states and Viterbi paths are quite different from those recovered via the SFC HMM method.

Our initial fitting and model recovery results (**Results Section 3.1.2**) suggest that these differences are not idiosyncratic, but are robust to fitting procedures. First, we showed that FFC HMM is better suited to extracting true functional connectivity profiles than SFC HMM (and certainly better than IB HMM). We observed that the correlations between induced and discovered states were on average much higher for FFC HMM than for SFC HMM (mean R^2 of 0.5738 ± 0.1301 for FFC HMM but of 0.3337 ± 0.1650 for SFC HMM). Second, we also showed that connectivity state patterns also largely remained the same for FFC and SFC HMMs even when the length of the functional times series was reduced by half. Specifically, we found functional time series with durations as short as 7.2 minutes were sufficient to provide reproducible functional connectivity states. Third, we also observed that the differences between SFC and FFC HMMs were preserved regardless of the size of the sliding window used for the calculation of dynamic functional connectivity. Connectivity state patterns appeared to be robust to window size indicating that there was some flexibility in selecting this parameter in sliding window analyses performed on big datasets with a high temporal resolution and large numbers of subjects. A window size containing less than 60 seconds of data was found to be ideal because the spatial patterns began to deviate from those observed when using longer windows. This is consistent with Lurie et al. 2020, who postulated that a window length less than 60 seconds may be the optimal window size (Lurie et al. 2020). Our time window analyses (**Figs. 6 & 7**) showed that a window less than 60 seconds makes little difference in the distinctiveness of the connectivity states recovered, either by SFC HMM or FFC HMM. Thus, our selection of 50 timepoints for the main analysis played no role in the observed differences in recovered connectivity states between the two models.

We also observed that, despite differences in the switching rate between SFC and FFC HMMs, the summed connectivity vectors for each matched state were similar (**Fig. 2**) as indicated by high values for the absolute value of the inner product of the summed connectivity vectors for matched states from each model (mean inner product across all states = 0.84). These results suggest that FFC HMM is able to recover the summed connectivity vectors that would be sought by SFC HMM, i.e. that both connectivity-based HMMs result in states with similar average

nodal connectivity. These findings are encouraging for the use of FFC HMM in place of SFC HMM, as examining FFC HMM's outputs allows us to answer all of the questions addressed by SFC HMM.

Interestingly, our results revealed that FFC HMM was slower to change states even with the same sliding window size as SFC. Given that FFC HMM examines pairwise connectivity of each node with all other nodes under consideration whereas SFC HMM examines the average connectivity of each node, we posit that the slower switching rate in FFC HMM is due to the increased number of components fitted in FFC HMM compared to SFC HMM. The increased number of components fitted in the FFC HMM may reduce state switching by requiring broader changes in functional connectivity across multiple nodes to register a change in functional connectivity state. It is possible this behavior may also make FFC HMM less sensitive to noise than SFC HMM, although a comprehensive demonstration of this consequence is beyond the scope of the present paper. Future work should comprehensively examine the underlying causes for the slower switch rates in FFC HMM as they may relate to these and other possible scenarios.

4.2 Limitations

Of course, our approach demonstrating FFC HMM's capabilities does have limitations. For example, one might note that when using sliding window functional correlations as inputs (SFC and FFC HMMs), there is a large amount of autocorrelation in the input connectivity profiles due to the length of each sliding window being longer than the incrementing of the window's center (essentially, each new 'window' contains a majority of the previous 'window'). HMMs in general assume independent Gaussians in the likelihood functions and three different algorithms (the forward algorithm, the Viterbi algorithm, and the Baum-Welch algorithm) in conjunction to find the most likely sequence of hidden states, transition probabilities, and emission probabilities based on an observable sequence of data (Jurafsky and Martin 2009; S. R. Eddy 1996; Sean R. Eddy 2004; L. Rabiner and Juang 1986; L. R. Rabiner 1989). Thus, autocorrelation of the functional connectivity profiles across time might cause some concerns in interpretation. Here, we have assumed independence as a simplifying assumption, following previous work (Ou et al. 2015). Future work should evaluate the degree to which autocorrelation affects the outputs of any model which takes dynamic functional connectivity as input (here, SFC and FFC HMMs).

Additionally, the analysis presented here cannot speak to the entire spectrum of types of data one might wish to analyze with hidden Markov approaches. For example, we did not evaluate the degree to which TR impacts FFC HMM's performance, nor did we evaluate the shortest resting-state scan duration required to fit FFC HMM satisfactorily. Such explorations would indeed be useful, but unfortunately are beyond the scope of this manuscript as our goal here was to present a proof of concept and empirical demonstration for the advantages of FFC HMM over other established methods. Future work should examine these factors in detail.

Finally, our results are also limited due to our selection of sliding time window computations of functional connectivity as types of inputs to our hidden Markov models. Although the sliding time window approach retains sequential dynamic temporal information (Chen et al. 2016), it is limited in that a given brain state may not persist for the entire sliding window, or signals from multiple brain states may overlap in the sliding time window. If instead one wishes to characterize some number of unique, stable states of functional connectivity -- i.e., under some

conditions, two nodes are highly correlated and under others they are not -- one can derive functional connectivity states using other dynamical approaches such as spatial independent component analysis (Beckmann et al. 2005; Smith et al. 2012) or structural equation modeling (Schlösser et al. 2003; Petridou et al. 2013). Other similar clustering methods have been designed primarily to look at intensity-based states, such as coactivation patterns (Liu and Duyn 2013; Liu, Chang, and Duyn 2013). However, while independent component analysis and structural equation modeling are able to characterize the spatial patterns of connectivity states -- i.e., stable connectivity patterns describing which nodes tend to be connected to other nodes -- they only identify the connectivity states themselves and do not account for the temporal ordering, or sequence, of fMRI times series or, consequently, the trajectory through connectivity state space that an individual might traverse. With these alternative methods, each time point in a fMRI time series is treated as an independent sample of the brain in these approaches and shuffling fMRI time series does not affect spatial patterns of brain states derived from these approaches. Thus, because these approaches are unable to analyze temporal interactions between connectivity states or how they evolve over time, as of now one of the best tools we have to study the evolution of connectivity states over time.

4.3 Conclusions

We found that the newly-introduced FFC HMM discovered connectivity states with more distinguishable patterns than those derived from HMMs with an intensity input (IB HMM) or summed connectivity input (SFC HMM). Our comprehensive state comparison across IB HMM, SFC HMM, and FFC HMM revealed the connectivity profiles as extracted by FFC HMM to be fundamentally different from the functional connectivity profiles extracted by either of the other two methods. We also demonstrated that the disadvantages of fitting FFC HMM in terms of required compute power are outweighed by the direct interpretability of the findings: Rather than hypothesize about what connectivity-based states might have been from a simpler model, FFC HMM allows for an essentially direct readout of connectivity-based states and their temporal evolution. This behavior makes FFC HMM a powerful tool for extracting, analyzing, and understanding how dynamic connectivities among brain regions may relate to different disease states or cognitive processes in future studies.

Data/Code Availability Statement

Data from HCP is available from the HCP database (<https://ida.loni.usc.edu/login.jsp>). Hidden Markov models were generated using the `hmmlearn` library in python (<https://github.com/hmmlearn/hmmlearn>).

Credit Authorship Contribution Statement

Sana Hussain: Conceptualization, Formal Analysis, Investigation, Methodology, Project Administration, Software, Validation, Visualization, Writing -- Original Draft, Writing -- Review & Editing. **Jason Langley:** Investigation, Methodology, Supervision, Validation, Visualization, Writing -- Review & Editing. **Aaron R. Seitz:** Conceptualization, Funding Acquisition,

Investigation, Methodology, Project Administration, Supervision, Validation, Visualization, Writing -- Review & Editing. **Megan A. K. Peters:** Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Validation, Visualization, Writing -- Original Draft, Writing -- Review & Editing. **Xiaoping P. Hu:** Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Validation, Visualization, Writing -- Review & Editing.

Declaration of Competing Interest

None

Acknowledgements

This work was supported in part by the UCR NASA MIRO FIELDS Fellowship (to Sana Hussain). Data collection and sharing for this project was provided by the Human Connectome Project (HCP; Principal Investigators: Bruce Rosen, M.D., Ph.D., Arthur W. Toga, Ph.D., Van J. Weeden, MD). HCP funding was provided by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH), and the National Institute of Neurological Disorders and Stroke (NINDS). HCP data are disseminated by the Laboratory of Neuro Imaging at the University of Southern California. This work was funded by NIA R01 NS108638-01 (PIs: Xiaoping P. Hu and Aaron R. Seitz) and by the Canadian Institute for Advanced Research Azrieli Global Scholars Program (PI: Megan A. K. Peters). Funding sources had no involvement in the design and methodology of the study.

References

- Baker, Adam P, Matthew J Brookes, Ieab A Rezek, Stephen M Smith, Timothy Behrens, Penny J Probert Smith, and Mark Woolrich. 2014. “Fast Transient Networks in Spontaneous Human Brain Activity.” *ELife* 3 (March): e01867. <https://doi.org/10.7554/eLife.01867>.
- Beckmann, Christian F, Marilena DeLuca, Joseph T Devlin, and Stephen M Smith. 2005. “Investigations into Resting-State Connectivity Using Independent Component Analysis.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1457): 1001–13. <https://doi.org/10.1098/rstb.2005.1634>.
- Brett, Matthew, Ingrid S. Johnsrude, and Adrian M. Owen. 2002. “The Problem of Functional Localization in the Human Brain.” *Nature Reviews Neuroscience* 3 (3): 243–49. <https://doi.org/10.1038/nrn756>.
- Chen, Shiyang, Jason Langley, Xiangchuan Chen, and Xiaoping Hu. 2016. “Spatiotemporal Modeling of Brain Dynamics Using Resting-State Functional Magnetic Resonance Imaging with Gaussian Hidden Markov Model.” *Brain Connectivity* 6 (4): 326–34. <https://doi.org/10.1089/brain.2015.0398>.
- Deshpande, Gopikrishna, Xiaoping Hu, Randall Stilla, and K. Sathian. 2008. “Effective Connectivity during Haptic Perception: A Study Using Granger Causality Analysis of Functional Magnetic Resonance Imaging Data.” *NeuroImage* 40 (4): 1807–14. <https://doi.org/10.1016/j.neuroimage.2008.01.044>.
- Deshpande, Gopikrishna, Stephan LaConte, George Andrew James, Scott Peltier, and Xiaoping Hu. 2009. “Multivariate Granger Causality Analysis of fMRI Data.” *Human Brain Mapping* 30 (4): 1361–73. <https://doi.org/10.1002/hbm.20606>.
- Deshpande, Gopikrishna, Priya Santhanam, and Xiaoping Hu. 2011. “Instantaneous and Causal Connectivity in Resting State Brain Networks Derived from Functional MRI Data.” *NeuroImage* 54 (2): 1043–52. <https://doi.org/10.1016/j.neuroimage.2010.09.024>.
- Díez-Cirarda, María, Antonio P. Strafella, Jinhee Kim, Javier Peña, Natalia Ojeda, Alberto Cabrera-Zubizarreta, and Naroa Ibarretxe-Bilbao. 2018. “Dynamic Functional Connectivity in Parkinson’s Disease Patients with Mild Cognitive Impairment and Normal Cognition.” *NeuroImage: Clinical* 17: 847–55. <https://doi.org/10.1016/j.nicl.2017.12.013>.
- Eavani, Harini, Theodore D. Satterthwaite, Raquel E. Gur, Ruben C. Gur, and Christos Davatzikos. 2013. “Unsupervised Learning of Functional Network Dynamics in Resting State fMRI.” In *Information Processing in Medical Imaging*, edited by James C. Gee, Sarang Joshi, Kilian M. Pohl, William M. Wells, and Lilla Zöllei, 7917:426–37. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-38868-2_36.
- Eddy, S. R. 1996. “Hidden Markov Models.” *Current Opinion in Structural Biology* 6 (3): 361–65. [https://doi.org/10.1016/s0959-440x\(96\)80056-x](https://doi.org/10.1016/s0959-440x(96)80056-x).
- Eddy, Sean R. 2004. “What Is a Hidden Markov Model?” *Nature Biotechnology* 22 (10): 1315–16. <https://doi.org/10.1038/nbt1004-1315>.
- Fiorenzato, Eleonora, Antonio P Strafella, Jinhee Kim, Roberta Schifano, Luca Weis, Angelo Antonini, and Roberta Biundo. 2019. “Dynamic Functional Connectivity Changes Associated with Dementia in Parkinson’s Disease.” *Brain* 142 (9): 2860–72. <https://doi.org/10.1093/brain/awz192>.
- Glasser, Matthew F., Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, et al. 2013. “The Minimal Preprocessing

- Pipelines for the Human Connectome Project.” *NeuroImage* 80 (October): 105–24.
<https://doi.org/10.1016/j.neuroimage.2013.04.127>.
- Guggenmos, Matthias, Philipp Sterzer, and Radoslaw Martin Cichy. 2018. “Multivariate Pattern Analysis for MEG: A Comparison of Dissimilarity Measures.” *NeuroImage* 173 (June): 434–47. <https://doi.org/10.1016/j.neuroimage.2018.02.044>.
- Jurafsky, Dan, and James H. Martin. 2009. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 2nd ed. Prentice Hall Series in Artificial Intelligence. Upper Saddle River, N.J: Pearson Prentice Hall.
- Laird, Angela R., Jack L. Lancaster, and Peter T. Fox. 2005. “BrainMap: The Social Evolution of a Human Brain Mapping Database.” *Neuroinformatics* 3 (1): 065–078.
<https://doi.org/10.1385/NI:3:1:065>.
- Lancaster, Jack L., Diana Tordesillas-Gutiérrez, Michael Martinez, Felipe Salinas, Alan Evans, Karl Zilles, John C. Mazziotta, and Peter T. Fox. 2007. “Bias between MNI and Talairach Coordinates Analyzed Using the ICBM-152 Brain Template.” *Human Brain Mapping* 28 (11): 1194–1205. <https://doi.org/10.1002/hbm.20345>.
- Liu, Xiao, Catie Chang, and Jeff H. Duyn. 2013. “Decomposition of Spontaneous Brain Activity into Distinct fMRI Co-Activation Patterns.” *Frontiers in Systems Neuroscience* 7.
<https://doi.org/10.3389/fnsys.2013.00101>.
- Liu, Xiao, and Jeff H. Duyn. 2013. “Time-Varying Functional Network Information Extracted from Brief Instances of Spontaneous Brain Activity.” *Proceedings of the National Academy of Sciences* 110 (11): 4392–97. <https://doi.org/10.1073/pnas.1216856110>.
- Lurie, Daniel J., Daniel Kessler, Danielle S. Bassett, Richard F. Betzel, Michael Breakspear, Shella Kheilholz, Aaron Kucyi, et al. 2020. “Questions and Controversies in the Study of Time-Varying Functional Connectivity in Resting fMRI.” *Network Neuroscience* 4 (1): 30–69. https://doi.org/10.1162/netn_a_00116.
- Ou, Jinli, Li Xie, Changfeng Jin, Xiang Li, Dajiang Zhu, Rongxin Jiang, Yaowu Chen, Jing Zhang, Lingjiang Li, and Tianming Liu. 2015. “Characterizing and Differentiating Brain State Dynamics via Hidden Markov Models.” *Brain Topography* 28 (5): 666–79.
<https://doi.org/10.1007/s10548-014-0406-2>.
- Pedregosa, Fabian, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. 2011. “Scikit-Learn: Machine Learning in Python.” *Journal of Machine Learning Research* 12: 2825–30.
- Petridou, Natalia, César Caballero Gaudes, Ian L. Dryden, Susan T. Francis, and Penny A. Gowland. 2013. “Periods of Rest in fMRI Contain Individual Spontaneous Events Which Are Related to Slowly Fluctuating Spontaneous Activity.” *Human Brain Mapping* 34 (6): 1319–29. <https://doi.org/10.1002/hbm.21513>.
- Rabiner, L., and B. Juang. 1986. “An Introduction to Hidden Markov Models.” *IEEE ASSP Magazine* 3 (1): 4–16. <https://doi.org/10.1109/MASSP.1986.1165342>.
- Rabiner, L.R. 1989. “A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition.” *Proceedings of the IEEE* 77 (2): 257–86. <https://doi.org/10.1109/5.18626>.
- Raichle, Marcus E. 2011. “The Restless Brain.” *Brain Connectivity* 1 (1): 3–12.
<https://doi.org/10.1089/brain.2011.0019>.
- Schlösser, Ralf, Thomas Gesierich, Bettina Kaufmann, Goran Vucurevic, Stefan Hunsche, Joachim Gawehn, and Peter Stoeter. 2003. “Altered Effective Connectivity during Working Memory Performance in Schizophrenia: A Study with fMRI and Structural

- Equation Modeling.” *NeuroImage* 19 (3): 751–63.
[https://doi.org/10.1016/S1053-8119\(03\)00106-X](https://doi.org/10.1016/S1053-8119(03)00106-X).
- Shappell, Heather, Brian S. Caffo, James J. Pekar, and Martin A. Lindquist. 2019. “Improved State Change Estimation in Dynamic Functional Connectivity Using Hidden Semi-Markov Models.” *NeuroImage* 191 (May): 243–57.
<https://doi.org/10.1016/j.neuroimage.2019.02.013>.
- Smith, S. M., K. L. Miller, S. Moeller, J. Xu, E. J. Auerbach, M. W. Woolrich, C. F. Beckmann, et al. 2012. “Temporally-Independent Functional Modes of Spontaneous Brain Activity.” *Proceedings of the National Academy of Sciences* 109 (8): 3131–36.
<https://doi.org/10.1073/pnas.1121329109>.
- Stevner, A. B. A., D. Vidaurre, J. Cabral, K. Rapuano, S. F. V. Nielsen, E. Tagliazucchi, H. Laufs, et al. 2019. “Discovery of Key Whole-Brain Transitions and Dynamics during Human Wakefulness and Non-REM Sleep.” *Nature Communications* 10 (1): 1035.
<https://doi.org/10.1038/s41467-019-08934-3>.
- Stilla, R., G. Deshpande, S. LaConte, X. Hu, and K. Sathian. 2007. “Posteromedial Parietal Cortical Activity and Inputs Predict Tactile Spatial Acuity.” *Journal of Neuroscience* 27 (41): 11091–102. <https://doi.org/10.1523/JNEUROSCI.1808-07.2007>.
- Van Essen, David C., Stephen M. Smith, Deanna M. Barch, Timothy E.J. Behrens, Essa Yacoub, and Kamil Ugurbil. 2013. “The WU-Minn Human Connectome Project: An Overview.” *NeuroImage* 80 (October): 62–79. <https://doi.org/10.1016/j.neuroimage.2013.05.041>.
- Vidaurre, Diego, Romesh Abeysuriya, Robert Becker, Andrew J. Quinn, Fidel Alfaro-Almagro, Stephen M. Smith, and Mark W. Woolrich. 2018. “Discovering Dynamic Brain Networks from Big Data in Rest and Task.” *NeuroImage* 180 (October): 646–56.
<https://doi.org/10.1016/j.neuroimage.2017.06.077>.
- Vidaurre, Diego, Laurence T. Hunt, Andrew J. Quinn, Benjamin A. E. Hunt, Matthew J. Brookes, Anna C. Nobre, and Mark W. Woolrich. 2018. “Spontaneous Cortical Activity Transiently Organises into Frequency Specific Phase-Coupling Networks.” *Nature Communications* 9 (1): 2987. <https://doi.org/10.1038/s41467-018-05316-z>.
- Vidaurre, Diego, Andrew J. Quinn, Adam P. Baker, David Dupret, Alvaro Tejero-Cantero, and Mark W. Woolrich. 2016. “Spectrally Resolved Fast Transient Brain States in Electrophysiological Data.” *NeuroImage* 126 (February): 81–95.
<https://doi.org/10.1016/j.neuroimage.2015.11.047>.
- Vidaurre, Diego, Stephen M. Smith, and Mark W. Woolrich. 2017. “Brain Network Dynamics Are Hierarchically Organized in Time.” *Proceedings of the National Academy of Sciences* 114 (48): 12827–32. <https://doi.org/10.1073/pnas.1705120114>.
- Yang, Zhi, Stephen LaConte, Xuchu Weng, and Xiaoping Hu. 2008. “Ranking and Averaging Independent Component Analysis by Reproducibility (RAICAR).” *Human Brain Mapping* 29 (6): 711–25. <https://doi.org/10.1002/hbm.20432>.

Appendices

Appendix A. Detailed methods

A.1 Network ROI MNI coordinates

We narrowed our scope of analysis to four pre-defined networks that have previously been associated with resting state: the default mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), and salience network (SN). The nodes comprising each network were defined using anatomical coordinates specified in literature and converted from Talairach to MNI coordinates when necessary (Raichle 2011; Lancaster et al. 2007; Laird, Lancaster, and Fox 2005; Brett, Johnsrude, and Owen 2002; Deshpande et al. 2008; 2009; Stilla et al. 2007). Dorsal anterior cingulate cortex and left dorsolateral prefrontal cortex ROIs in FPCN and right parahippocampal gyrus and right inferolateral temporal cortex in DMN were excluded from analyses because they overlapped or were too closely located to other ROIs. Thus, we only used 29 ROIs per subject when examining the HCP dataset (9 from DMN, 7 from FPCN, 6 from DAN, and 7 from SN).

ROI	Abbrev.	Full Name	MNI (x,y,z)	Source
Default Mode Network				
1	PCC	Posterior Cingulate Cortex	(2,54,16)	Deshpande et al. 2011
2	L pIPL	Left Posterior Inferior Parietal Lobule	(-46,-72,28)	Deshpande et al. 2011
3	R pIPL	Right Posterior Inferior Parietal Lobule	(50,-64,26)	Deshpande et al. 2011
4	OFC/vACC	Orbitofrontal Cortex/Ventral Anterior Cingulate Cortex	(4,30,26)	Deshpande et al. 2011
5	dMPFC BA 8	Dorsomedial Prefrontal Cortex Brodmann Area 8	(-14,54,34)	Deshpande et al. 2011
6	dMPFC BA 9	Dorsomedial Prefrontal Cortex Brodmann Area 9	(22,58,26)	Deshpande et al. 2011
7	L DLPFC	Left Dorsolateral Prefrontal Cortex	(-50,20,34)	Deshpande et al. 2011
8	L PHG	Left Parahippocampal Gyrus	(-10,-38,-2)	Deshpande et al. 2011
9	L ITC	Left Inferolateral Temporal Cortex	(-60,-20,-18)	Deshpande et al. 2011
Fronto-Parietal Control Network				
1	L aPFC	Left Anterior Prefrontal Cortex	(-36,56,10,)	Deshpande et al. 2011
2	R aPFC	Right Anterior Prefrontal Cortex	(34,52,10)	Deshpande et al. 2011

3	R DLPFC	Right Dorsolateral Prefrontal Cortex	(46,14,42)	Deshpande et al. 2011
4	L aINS	Left Anterior Insula	(-30,20,-2)	Deshpande et al. 2011
5	R aINS	Right Anterior Insula	(32,22,-2)	Deshpande et al. 2011
6	L aIPL	Left Anterior Inferior Parietal Lobule	(-52,-50,46)	Deshpande et al. 2011
7	R aIPL	Right Anterior Inferior Parietal Lobule	(52,-46,46)	Deshpande et al. 2011
Dorsal Attention Network				
1	L MT	Left MidThalamus	(-44,-64,-2)	Deshpande et al. 2011
2	R MT	Right MidThalamus	(50,-70,-4)	Deshpande et al. 2011
3	L FEF	Left Frontal Eye Field	(-24,-8,50)	Deshpande et al. 2011
4	R FEF	Right Frontal Eye Field	(28,-10,50)	Deshpande et al. 2011
5	L SPL	Left Superior Parietal Lobule	(-26,-52,56)	Deshpande et al. 2011
6	R SPL	Right Superior Parietal Lobule	(24,-56,-54)	Deshpande et al. 2011
Salience Network				
1	DAC	Dorsal Anterior Cingulate	(0,-22,36)	Raichle 2011
2	L aPFC	Left Anterior PFC	(-34,44,30)	Raichle 2011
3	R aPFC	Right Anterior PFC	(32,44,30)	Raichle 2011
4	L Insula	Left Insula	(-40,2,6)	Raichle 2011
5	R Insula	Right Insula	(42,2,6)	Raichle 2011
6	L LP	Left Lateral Parietal	(-62,-46,30)	Raichle 2011
7	R LP	Right Lateral Parietal	(62,-46,30)	Raichle 2011

Table A1: List of MNI coordinates used for ROIs in the default mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), and salience network (SN). Talaraich coordinates for DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and were converted to MNI using (Deshpande, Santhanam, and Hu 2011; Lancaster et al. 2007; Laird, Lancaster, and Fox 2005; Brett, Johnsrude, and Owen 2002) while MNI coordinates for SN were taken directly from Raichle 2011 (Raichle 2011).

A.2 HMM models: graphical descriptions

The following figures graphically describe the methods used to implement the three HMMs discussed in the main text. **Fig. A1** corresponds to IB HMM, **Fig. A2** to SFC HMM, and **Fig. A3** to FFC HMM.

A.2.1 Intensity-based HMM

Step 1: Extract BOLD Signal From ROIs

Step 2: Concatenate Across Subjects

Step 3: Input Into HMM

Figure A1. Illustration of the procedure used to implement IB HMM. Once the BOLD signal is extracted from all predefined ROIs, it is concatenated across subjects and fitted with an HMM from the python hmmlearn library (Pedregosa et al. 2011).

A.2.2 Summed functional connectivity HMM

Step 1: Sliding Window Analysis → Pearson Correlate All ROIs Using Time Window Δt & Move Over 1 Time Point

Outputs “dFC Time Series”:

Step 2: Sum Each ROI x ROI Matrix Over 1 Dimension To Obtain Node-Wise Connectivity Vector (“Summed dFC Time Series”)

Step 3: Repeat For All Subjects

Step 4: Concatenate Across Subjects and Input Into HMM

Figure A2. Illustration of the procedure used to implement SFC HMM as followed by the method described in Ou et al. 2015 (Ou et al. 2015). Once a sliding time window correlation analysis is performed, the connectivity matrix in each time window is summed across a dimension into a representative nodal connectivity vector. After executing this for all time windows and for all subjects, the data is concatenated across all subjects and fitted with an HMM from the python hmmlearn library (Pedregosa et al. 2011).

A.2.3 Full functional connectivity HMM

Step 1: Sliding Window Analysis → Pearson Correlate All ROIs Using Time Window Δt & Move Over 1 Time Point

Step 2: Acquire All Values From Lower (or Upper) Triangle

Step 3: Repeat For All Subjects

Step 4: Concatenate Across Subjects and Input Into HMM

Figure A3. Illustration of the procedure used to implement FFC HMM. Once a sliding time window correlation analysis is performed, the lower (or upper) triangle of values from the connectivity matrix in each time window is flattened into a vector. After executing this for all time windows and for all subjects, the data is concatenated across all subjects and fitted with an HMM from the python hmmlearn library (Pedregosa et al. 2011).

A.3 Detail of the RAICAR-based method to determine number of hidden states for each model

Fig. A4 illustrates the procedure used to determine the optimal number of states for this investigation: a Ranking and Averaging Independent Component Analysis by Reproducibility (RAICAR) based method. The stability analysis results for models with 3-15 hidden states for IB HMM and SFC HMM are depicted in **Figs. A5** and **A6**, respectively.

Figure A4. Schematic illustrating the procedure for the RAICAR-based method for determining number of hidden states in a model. Each of the blocks shown represent a state where matched states are shown in the same color. A basic example of 3 states is shown where states' spatial patterns are first matched across realizations using Pearson correlation values (Chen et al. 2016; Yang et al. 2008). That is, states' spatial patterns have been matched via the highest R^2 value observed and reordered to the same state assignment. The Pearson correlation is then found amongst all pairs of states within a matched group and averaged to represent the stability of that state pattern for that number of hidden states. This procedure is then repeated for all states, with a model fitted with that particular number of hidden states.

Figure A5. Stability analysis results for IB HMM for models with 3-15 hidden states with a predetermined threshold at 0.9. It may appear that models with as many as 12 hidden states could be suitable for IB HMM because its stability values remained above the predetermined threshold. However, models with 10, 11, and 12 states contained a repeated state making them undesirable as a lack of parsimony likely occurred in the state spatial patterns. The 9-state model identified a state where mean activation equaled zero (consistent with Chen et al. 2016's findings) that was not observed in the 8-state model (cite once changes are accepted). The 10-, 11-, and 12-state models included two occurrences of this activation pattern; although these models were considered stable, they were undesirable because they contained a repeated state and a lack of parsimony. Thus, eight states were the preferred choice for IB HMM in this investigation.

Figure A6. Stability analysis results for SFC HMM for models with 3-15 hidden states with a predetermined threshold at 0.9.

A.4 Details of model robustness to data size

In **Fig. A7**, we present the connectivity-based state patterns identified (along with their corresponding summed vectors) when only half of the fMRI dataset (7.2 minutes) was fitted with an HMM.

Figure A7. State pattern differential functional connectivity maps for SFC HMM (top row), FFC HMM (middle row), and IB HMM (bottom row) when only half of the fMRI dataset (7.2 minutes) were used. The summed connectivity values across one dimension are displayed below each state.

Appendix B. Intensity-based state patterns

B.1 Pattern acquisition

Intensity state patterns are defined as a combination of activated or deactivated ROIs comprising the four aforementioned networks (i.e., each state consists of a 1x29 vector of intensity levels). Intensity states in the SFC and FFC HMMs were computed using the method of state acquisition described by Chen and colleagues (Chen et al. 2016). **Fig. B1** shows that the BOLD time series was truncated by a length of $\Delta t/2$ to match the temporal scale of the SFC and FFC Viterbi paths. Next, the BOLD signal at every TR where the SFC or FFC labeled a state to be active was averaged. Repeating this for all ROIs yielded a 1 X ROI vector of intensity for every ROI for a particular state.

Figure B1. Schematic illustration outlining the method used to recreate the intensity states using the Viterbi paths from SFC and FFC HMMs. The discrepancy in temporal resolution was accounted for by removing some TRs of the BOLD time series to ensure equal length between the BOLD signal and the connectivity-based HMMs' Viterbi paths.

B.2 Intensity-based states for all HMMs

The intensity patterns for IB HMM are seen in **Fig. B2** where the subscripts correspond to the model discussed; i.e., $S1_{IB}$ corresponds to state 1 from IB HMM) were directly outputted from the model fitting procedure. $S1_{IB}$ appears to be a DMN-dominant state since the DMN is activated, and all other networks are deactivated. $S2_{IB}$ shows both DMN and FPCN to be activated. $S3_{IB}$ and $S4_{IB}$ both appear to be attention-dominant states: DAN and SN are activated in $S3_{IB}$ while FPCN, DAN, and SN are activated in $S4_{IB}$. $S5_{IB}$ and $S7_{IB}$ both show all networks to

be activated, but $S5_{IB}$ has slightly lower activation levels compared to $S7_{IB}$. In $S6_{IB}$, DMN, DAN, and SN are deactivated and FPCN has minor positive activation levels. $S8_{IB}$ shows all networks to be deactivated. The corresponding connectivity states stemming from the output covariance matrices can be found in the bottom row of **Fig. 3**. The intensity states for SFC and FFC HMMs using the methods illustrated in **Fig. B1** are seen in **Fig. B3a** and **B3c**, respectively while the Euclidean distances between them and the intensity states directly outputted from IB HMM are seen in **Fig. B3b** and **Fig. B3d**, respectively. Comparing the two connectivity-based models' results against IB HMM's outputted intensity state patterns aimed to determine whether the model types were recognizing different states and ensured that repeated information was not acquired from different HMMs. The similarities of these states to each other were assessed via Euclidean distances (**Fig. B3**).

Figure B2. IB HMM activation state patterns. Values shown in red correspond to positive activation levels while those in blue correspond to deactivation.

Figure B3. Intensity patterns from averaging the BOLD signal according to the Viterbi path from (A) SFC HMM and (B) FFC HMM. (C) and (D), respectively, show the Euclidean distances between the activation patterns directly outputted from IB HMM (Fig. B2) and those from panels A and B.

These results indicate that IB HMM recognized fluctuations in BOLD signal, and that SFC and FFC HMMs were not sensitive to those same patterns. Each model adequately recognized changes in the data inputted into their respective model type and, therefore, did not recognize the same state patterns. Thus, each HMM is distinct in identifying changes unique to the inputted data, and consequently, in identifying states.

B.3 Intensity-based state model validation

To create “intensity-defined” ground truth states, eight intensity states (matching the best number of hidden states recovered in the primary analysis; see **Results Section 3.1.1**) were induced by adding an arbitrary value of 2 (the signal is in arbitrary units) to nodes in certain networks and at certain time points within the permuted data.

Fig. B4a depicts the time series of states we induced during the validation analysis. In accordance with the stability analysis results, 8 states were induced in our model validation analyses. For the first validation instantiation, intensity states were induced by adding a value of 2 to the specified networks and fitted with IB HMM, which adequately recovered each of the described states (**Fig. B4b**). Compared to the induced state sequence (**Fig. B4a**), the IB HMM fitted to this sequence outputted a state time course with mean R^2 across all subjects of $0.9999 \pm 4.1551e-04$. Consistent with the unique nature of intensity-based versus connectivity-based states (see **Results Section 3.2.1**), as expected neither of the connectivity-based models could recover states induced by intensity-based approaches (**Figs. B4c** and **B4d**; SFC HMM $R^2 = 0.0501 \pm 0.1003$ and FFC HMM $R^2 = 0.0739 \pm 0.0683$ respectively). This finding suggests that our IB HMM model fitting procedure is ideally suited to recover intensity-based states, and the SFC and FFC models do not ‘accidentally’ discover such states when they are present.

Figure B4. Verification of HMM intensity-based states. (A) The artificially induced state sequence depicted which networks exhibited slightly increased network intensity levels. Outputted state sequences from (B) IB HMM, (C) SFC HMM, and (D) FFC HMM when intensity states were induced.