UNIVERSITY OF CALIFORNIA
RIVERSIDE

Exploring Hidden Markov Models in Human Functional Magnetic Resonance Imaging Data With Applications to the Locus Coeruleus Circuit

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Bioengineering

by

Sana Ali Hussain

September 2021

Dissertation Committee:
Dr. Xiaoping Hu, Chairperson
Dr. Megan A.K. Peters
Dr. Aaron Seitz

The Dissertation of Sana Ali Hussain is approved:

Committee Chairperson

University of California, Riverside

Acknowledgements

I would first like to thank my family. Thank you to my grandparents for their sacrifices long ago which allowed me to pursue higher education. Thank you to my grandparents, my aunt, and my uncle for always supporting my achievements.

I would also like to deeply thank Dr. Jason Langley. I am very grateful for the time, effort, and teachings he has devoted to me throughout my research experience.

Thank you to my dissertation committee member, Dr. Aaron Seitz, for all of your help and guidance throughout this project. Thank you to all of my LC collaborators and lab members for our discussions over the years.

I am also very grateful to Dr. Megan A.K. Peters for taking me under her wing towards the beginning of my journey at UC Riverside. She has helped me immensely in becoming a better researcher and bioengineer.

Lastly, I would like to thank my primary advisor, Dr. Xiaoping Hu. Thank you for accepting me into your lab, and for all of your help and guidance throughout graduate school.

This work was supported in part by the UCR NASA MIRO Fellowships and Internships for Extremely Large Datasets program.

Dedication

I dedicate this dissertation, and all my accomplishments, to my parents, Azra and Ali Hussain. It is because of their unconditional and unyielding love and support that I have succeeded with this feat.

ABSTRACT OF THE DISSERTATION

Exploring Hidden Markov Models in Human Functional Magnetic Resonance Imaging Data With Applications to the Locus Coeruleus Circuit

by

Sana Ali Hussain

Doctor of Philosophy, Graduate Program in Bioengineering
University of California, Riverside, September 2021
Dr. Xiaoping Hu, Chairperson

There has been recent growth in using hidden Markov models (HMMs) to chronicle the latent spatiotemporal dynamics of brain activity acquired from functional magnetic resonance imaging (fMRI). This technique can be used to model resting state data or dynamic processing, such as attention. The locus coeruleus (LC), a small subcortical structure, is the main source of norepinephrine throughout the brain and therefore is involved in modulating attention and arousal. An HMM can be used to identify latent fMRI-based brain states comprised of a combination of networks, and to determine how the behavior of these brain states change as a function of attentional perturbations.

The purpose of this investigation is two pronged. We aimed to create a comprehensive theoretical and empirical overview of various HMM subtypes in effort to make informed decisions about which should be used under different circumstances for future investigations. We then fit this probabilistic model to an fMRI dataset optimized to

image the effects of the LC to gain insight into its dynamic relationship with attention and arousal.

To accomplish this, we first theoretically contrasted three HMM subtypes, then applied them to an fMRI resting state dataset to obtain an empirical understanding of the strengths and weaknesses of each one. One model type, an activation-based HMM, was applied to a pseudo-resting state dataset where LC activity was noninvasively up-regulated via a handgrip task. This aimed to analyze how HMM-related measure focusing on attentional networks changed as a function of actively squeezing a squeeze-ball.

We found that an activation-based HMM is ideal for capturing temporal dynamics and is preferred if activation and connectivity state patterns are to be analyzed in conjunction. An HMM where functional connectivity values were summed at each time point is advantageous if a study's goal is to examine the general connectedness between nodes. We also noted that the HMM subtype fitted to all unique correlation values from a sliding window analysis is preferred if an investigation wishes to explore specific nodal connectivity. Our findings from HMM-based measures extricated from the LC-focused dataset show potential evidence of norepinephrine depletion and attentional allocation.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

The brain is a dynamic organ where activity never ceases, even in the absence of a task (1,2). It is comprised of a myriad of networks that interact and govern many different behaviors including, but not limited to, attention and perception (3–6). Various neuroimaging techniques have been used to image the brain and quantify network interactions thereby working towards a better understanding of the brain's dynamic behavior. Although there is a plethora of neuroimaging modalities available from which to acquire brain data, functional magnetic resonance imaging (fMRI) is commonly used because of its excellent spatial resolution and because it can image subcortical structures (7–10). Using this data acquisition modality, many groups have fragmentized the brain into a series of states comprised of different combinations of networks (11–22). One emerging technique is a hidden Markov model (HMM) because it intrinsically identifies spatial patterns of latent brain states and temporal patterns of the interactions between those brain states (23–26). An HMM can identify latent brain states pertaining to a wide variety of phenomena that the brain modulates (11–13,13,27,28). For the purposes of this dissertation, we focused on attention.

We were interested in better understanding how attention shifts during resting state, i.e., in the absence of a task. Furthermore, we concentrated on attention in order to examine the locus coeruleus (LC) circuit which has an established relationship with attention and arousal (29). The LC is a 2mm in diameter subcortical structure sitting atop the fourth ventricle in the brain (7,8). It is the main source of norepinephrine throughout the brain thereby making it the principal modulator of attention, the sleep-wake cycle, and stress

(9,10,29–32). Analyzing this structure is important because degradation of this structure has been linked to Alzheimer's and Parkinson's diseases (33,34). Thus, in order to comprehend LC dysfunction, we must first understand normal LC function. Because LC is associated with arousal and attention, examining its relationship with these trends provides major insight into its general function. We fit an HMM to a dataset optimized to image the LC to observe spatial patterns of attention-related brain states and discern how its measures changed over time in response to performing a handgrip task.

In Chapter 2, three different HMM subtypes were theoretically compared in effort to understand which model-type should be used to examine a certain research topic for future investigations. To accomplish this, we investigated three HMM subtypes that could be used to assess various neuroimaging-related topics: an activation-based HMM (AB HMM), a summed functional connectivity HMM (SFC HMM) derived from Ou et al. 2014, and a full functional connectivity HMM (FFC HMM) (11–13,28). The theoretical basis of each subtype was explained in great detail, as well as how model inputs were formatted and how model outputs were acquired. We discussed theoretical similarities and differences amongst all three subtypes as well as how the temporal resolution was affected by each model type's input. Two different yet synergistic methods of model order determination were introduced along with local and global analyses—all of which were utilized in Chapters 2, 3, and 4 (11,35). Local analyses focused on state patterns and characteristics. We examined different methods used to acquire activation and connectivity state patterns from each model type and explained why we expected them differ within and across each one. Global analyses described evaluation of measures

directly outputted from all three model types, or of ones that could be derived from the outputs. Specifically, we theoretically compared and contrasted the Viterbi path, transition probability matrix, switching rate, and fractional occupancy correlation across each model type. Two overarching questions emerged from our theoretical analyses: (1) What kinds of questions does each model type answer? (2) When is it useful to employ each model type? We attempted to obtain as much insight as possible into these inquiries in Chapter 3.

Chapter 3 focused on applying all HMM techniques described in Chapter 2 to an fMRI dataset to obtain empirical insight into the strengths and weaknesses of AB HMM, SFC HMM, and FFC HMM. From these results, we intended to deduce which HMM subtype is best used when exploring a particular research focus in an investigation. We analyzed a resting state dataset from the Human Connectome Project (HCP) as well as four specific attention-related networks: default-mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), and salience network (SN). After determining the ideal number of states to pursue in each model type, all activation and connectivity spatial patterns were closely contrasted. The Viterbi paths, transition probability matrices, switching rate, and fractional occupancy correlation from AB HMM, SFC HMM, and FFC HMM were also compared. We ended the chapter by proposing which model should be used in certain scientific inquiries based on the empirical evidence presented in the chapter.

In Chapter 4, we applied one HMM subtype (AB HMM) to a dataset targeting the LC circuit. The goal of this chapter was to understand how LC activity up-regulation

affected HMM-related measures (i.e., local and global analyses) in effort to obtain insight into LC's dynamic modulation of attention. Only AB HMM was chosen because the LC-dataset was a relatively short and employed a block design, so utilizing a windowed analysis would have reduced the temporal resolution to the point where no temporal dynamics could be extricated. Furthermore, conclusions from Chapter 2 indicated that SFC HMM and FFC HMM are not optimal for temporal dynamics, and because Chapter 4 aimed to dissect latent brain states over time, we opted not to use connectivity-based HMMs. Thus, a standard HMM was pursued and fitted to a modified handgrip experimental paradigm. Squeezing has previously been shown to stimulate sympathetic arousal, the underlying neural dynamics of which involve the solitary tract which innervates the LC (36–38). Thus, squeezing is thought to up-regulate LC activity, and its manifestation in perceived changes in downstream attentional network behavior was analyzed via an HMM. We focused on attention by fitting the HMM to attention-related ROIs, i.e., ones from DMN, FPCN, DAN, SN, and rostral and caudal LC. The same stability, local, and global analyses described in Chapters 2 and 3 were implemented here. We analyzed activation and connectivity patterns as well as the Viterbi path, switching rate, and transition probability matrix. Additionally, the fractional occupancy of a state, average state duration, and Fano factor were also included in the global analyses to evaluate how they changed as a function of active squeeze. Pupillometry data were also analyzed as a noninvasive proxy measure of fluctuations in LC activity.

In this chapter we have introduced the broad scientific topics discussed in Chapters 2-4. We provided background, rationale, and significance for studying HMMs and for

applying them to an LC-focused dataset. Over the next three chapters we delve deeper into these research aspects by exploring the related methods, results and discussion. Finally, we conclude the dissertation by summarizing our findings and their implications.

Chapter 2: Theoretical Foundations of Different Hidden Markov Model Types in Neuroimaging

2.1 Abstract

In order to better understand the underlying neural mechanisms of brain dynamics, many groups have examined brain states as a means of identifying recurring patterns of activation or connectivity. Although there is an abundance of methods of establishing and interpreting brain states, hidden Markov models are becoming an increasingly popular choice in neuroimaging. These models not only recognize the spatial patterns of the brain states, but also evaluate their temporal progression. Because it is so practical, an assortment of instantiations has arisen for diverse purposes and have been applied to data acquired from various neuroimaging modalities. Here we investigated more closely the methodology and theory involved in the most popular ones. We worked to attain clearer insight into their nature in order to make informed decisions about model choice in future analyses. Three hidden Markov models were introduced, and their inputs and outputs defined. We theoretically compared them by analyzing each subtype's assumptions, how they were applied to a dataset, and their potential applications. We discussed different types of analyses (local and global) that we intend to execute when applying the three model types to a neuroimaging dataset. We also posed two overarching questions that we expect to obtain insight into from empirical analyses.

2.2 Introduction

Analyzing the brain as a series of interacting and interchanging brain states is very popular in neuroimaging (11–14). Brain states are patterns of activation levels or

connectivity strengths that work to characterize the brain and quantify interactions amongst networks. Because the brain is not stationary, there lies an interest in studying brain states as a noninvasive way of gaining a deeper understanding of the underlying neural mechanisms of brain dynamics (11,13,14).

Other groups have derived brain states using coactivation patterns (CAPs), independent component analysis (ICA), or structural equation modelling (SEM) (15–20). CAPs is a modified point-process method that allows for extraction of network information from only a fraction of the data, and has shown that analysis of certain time points suggests the presence of multiple spatially distinct coactivation patterns (15,16). ICA decomposes thousands of voxels of data into components of networks and/or regions (17,18). SEM uses effective connectivity to identify nodal connections and determines potential causal dependencies between latent variables and their indicators (19). While CAPs, ICA, and SEM are useful methods for acquiring the spatial patterns of brain states, they simply identify the states themselves and do not report how they interact or change temporally.

To study the temporal evolution of brain states, the dynamic functional connectivity (dFC) is typically examined via a sliding time window approach (14,21,22). Information contained within a certain window length is correlated between time series, which is moved a certain number of time points until a general picture of connectivity shifts is obtained (14). This method can be beneficial because it retains sequential dynamic temporal information (11), and because the window size can be tailored to contain as much information from the time series as desired (14,21,39–41). However, brain states may be changing faster than what is allowed to be captures using a sliding window. Furthermore,

a sliding window correlation analysis is restricted to only recognizing state spatial patterns and how they change over time and may ignore important dynamic aspects of those parameters in a dataset (28).

Hidden Markov models (HMMs) have been shown to adequately characterize the fluctuating dynamics of temporally changing brain states in a data-driven manner (11,12,23,25). HMMs are probabilistic models that determine a hidden state sequence path not directly observable from collected data (23,25,26,42). It uses Markov chains to statistically infer the underlying states where the probability of residing in any one of them depends only on the previous state, and where each one of those hidden states generates the observable data collected (23,25,26,42). This model 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 (23,24,42).

In the past, HMMs have been applied to climate data for weather predicting (43), in computational biology to predict introns, exons, and slicing sites (26), and in finance to forecast stock market price trends (44). More recently HMMs have been used in neuroimaging to define spatial patterns of latent brain states as well as to recognize the transitions and time spent in those states (11–13,45). HMMs are useful in our investigation of a functional magnetic resonance imaging (fMRI) dataset because they do not form hypotheses about the relationship between brain states (11) and because analysis of spatial and temporal changes are inherent to the model. Specifically, HMMs can identify latent brain states from brain signals acquired from neuroimaging modalities such as fMRI and

magnetoencephalography (11,12,27,46–49). For fMRI, HMMs have shown sensitivity to changes in both the blood oxygen level dependent signal (11–13) and dynamic functional connectivity (28,49–51), and have been used in analyses where data were extracted from voxels (52–54), parcellations (12,13), or large-scale predefined regions of interest (11,28). With so many different types of HMMs available, choosing which model type is best for specific investigations can be challenging.

In this chapter we aimed to theoretically compared and contrasted three general HMM subtypes that can be fitted to any of the neuroimaging data types previously mentioned: (1) an activation based HMM (AB HMM) which is a standard HMM, (2) a summed functional connectivity HMM (SFC HMM) which has been applied to dFC analyses in the past and sums the Pearson correlations within each window into a representative connectivity vector (28), (3) a full functional connectivity HMM (FFC HMM) which is an extension of SFC HMM that fits all possible Pearson correlations from the dFC analysis The goal of these analyses is to identify strengths and limitations of each subtype in order to make informed decisions about which one should be selected during future investigations to analyze a specific research topic.

2.3 Methods

2.3.1 Overview

In this section we define two different categories of HMMs in three different instantiations for comparison. We focus on one activation-based and two connectivity-based HMMs. The activation-based model is termed AB HMM while the connectivity-based HMMs are termed SFC HMM (summed functional connectivity HMM) (28), and

FFC HMM (full functional connectivity HMM). After establishing the background for these models, we examine differences in the temporal resolution of inputted data, explain why we cannot directly compare the state patterns across the model types, and consider an alternative method for acquiring state spatial patterns for comparison with the direct model output counterpart. We then explore two different stability analysis methods that can be applied to all three HMMs as a means of determining the optimal model order. We then probe local and global analyses that can be performed to collate model subtype behavior and outputs. Then we pose two main questions that we aim to obtain empirical insight into by fitting all three HMM subtypes to an fMRI dataset. We finish this section by introducing and detailing a dataset and four predefined networks of interest.

2.3.2 Mathematical Background

Each of the HMMs are based on the same fundamental assumptions and equations. Here we feature the main mathematical equations and concepts, and for a more detailed discussion please refer to Rabiner and Juang 1986, Rabiner 1989 and Jurafsky and Martin 2009 (23,24,42). All equations and theory summarized below are based on their work (23,42).

HMMs employ two major assumptions. The first is a Markovian assumption where the probability of residing in a state depends only on the previous state (23,42). The second is an Output Independence assumption where the probability of an observation depends only on the state associated with generating that observation (23,25,26,42). Both are written in mathematical form below where Q is a set of N states ($Q = q_1q_2\ldots q_N$) and O is the sequence of observations ($O = o_1o_2\ldots o_N$) (23,42).

Markovian Assumption: $P(q_i|q_1 \dots q_{i-1}) = P(q_i|q_{i-1})$ Eq. 1

Output Independence Assumption:

$$P(o_i|q_1 \dots q_i, q_T, o_1, \dots, o_i, \dots, o_T) = P(o_i|q_i)$$ Eq. 2

An HMM uses two algorithms in combination to identify the most likely sequence of hidden states given an observable sequence of data. The forward algorithm determines the probability of the observed data given the most probable sequence of hidden states. The corresponding mathematical equation is written out in Eq. 3 which shows that the probability of being in state *j* at time *t* is found by summing the probability of the most likely state for all previous time points as weighted by their transition and emission probabilities. α_t gives the probability of state *j* occurring at time point *t* based on the information from the previous time step, a_{ij} represents the probability of transitioning out of one state, *i*, and into another, *j*, and $b_j(o_t)$ represents the probability that the current state generates the observation seen at time point *t* (23,42).

Forward Algorithm: $\alpha_t(j) = \sum_{i=1}^{N} \alpha_{t-1}(i) a_{ij} b_j(o_t)$ Eq. 3

The most probable sequence of states (Viterbi path) and the corresponding probability of that sequence given the observed data is decoded by the Viterbi algorithm. It also makes use of a backtrace algorithm which monitors the path of hidden states that led to each state. Eq. 4 gives the maximum probability, *v*, of a state *j* occurring at time point *t* while Eq. 5 reports the actual sequence of states. The Viterbi algorithm is essentially the forward algorithm with one important distinction: it computes the maximum probability of previous path probabilities whereas the forward algorithm calculates the sum of this parameter over all time points. This method is computationally efficient and distinct

because it computes the maximum likelihood for each time point, *t* (and total time *T)*, for each state in succession. This is in place of estimating the probability of the observable data given the entire state sequence and recursively changing the state sequence until the maximum observation likelihood is reached, and the model converges to the most likely hidden state sequence. The latter alternative technique is computationally costly and requires exponentially more calculations since there can be up to *T* states (23,42).

$$\text{Viterbi Algorithm: } v_t(j) = \max_{1 \le i \le N} v_{t-1}(i) a_{ij} b_j(o_t) \qquad \text{Eq. 4}$$

$$\text{Backtrace: } bt_t(j) = \operatorname*{argmax}_{1 \le i \le N} \alpha_{t-1}(i) a_{ij} b_j(o_t) \qquad \text{Eq. 5}$$

The HMM uses the Baum-Welch algorithm (forward-backward algorithm) to calculate the transition and emission probabilities of a given state. That is, for a given observable sequence of data, the HMM uses this algorithm to compute the parameters A (transition probability matrix where A = a_{11}…. a_{ij}…. a_{NN}) and B (sequence of observation likelihoods or emission probabilities where B = $b_i(o_t)$) by starting with an initial estimate then recursively changing it to compute a better one. The Baum-Welch algorithm is a recursive algorithm resulting in the output variable for one time point becoming an input variable for a succeeding time point (23,42).

$$\text{Baum-Welch Algorithm: } \beta_i(i) = \sum_{j=1}^{N} a_{ij} b_j(o_{t+1}) \beta_{t+1}(j) \qquad \text{Eq. 6}$$

In the information summarized above, the method for acquiring state patterns is not mentioned. HMMs have usually been fitted to categorial data where the hidden states are already known and where the motivation for implementing HMMs lie in acquiring the sequence of hidden states, transition probabilities, and emission probabilities (23,26,42,43). Only recently have HMMs been applied to neuroimaging data resulting in

different groups using different methods to obtain the spatial patterns of their hidden states. Vidaurre et al. 2017 and Stevner et al. 2019 utilized the HMMMAR toolbox to acquire activity patterns of each state (12,13). The hmmlearn python library outputs a mean parameter which is interpreted as state spatial patterns directly outputted by the model (55). Chen et al. 2016 acquired spatial patterns by averaging the z-scores of all time frames wherever the Viterbi algorithm decoded a state to be active (11). As discussed later in this chapter, we also acquired state spatial patterns in the same manner and compared them to the model output counterpart to ensure that they spatially match.

The mathematical equations and concepts described above are inherent to the hmmlearn python library, so the data simply needs to be structured into the right format prior to inputting it (55). All three HMM subtypes involve standard preprocessing the fMRI data, extracting the blood oxygen level dependent (BOLD) signal from chosen regions of interest (ROIs), and z-scoring them within each subject. The resulting BOLD signal is then organized differently for each HMM of interest as reviewed in detail in later sections. Once the data has been properly restructured and is ready to be fitted by an HMM, it is concatenated timewise across all subjects to create a matrix of size (time * # subjects) x (# ROIs). The hmmlearn library in python outputs the same four pieces of information for all three HMM subtypes used in our investigation: mean state patterns, covariance matrices for each state, a transition probability matrix, and a Viterbi path defining which state is active at each time point for every subject (55). The emission probability is not listed as a measure of interest because, in this case, it speaks more to the stability of the model, i.e., how well does the fitted model describe the observable dataset. Previous

investigations of HMMs applied to neuroimaging datasets explored techniques involving state spatial pattern robustness and reproducibility to assess the model's stability. To remain consistent with these established methods, we also employed these neuroimaging-based stability analyses described in additional detail in Section 3.3.1 (11,35).

2.3.3 Models to Compare

In this section we describe the procedure for implementing the HMMs in greater detail. We explain how to organize the data for input and how to acquire state patterns for all three instantiations.

2.3.3.1 Model Type 1: Activation Based Hidden Markov Model (AB HMM)

As aforementioned, this is a standard HMM that has been fitted to data acquired from various neuroimaging modalities (11–13). The BOLD signal from various ROIs are extracted, z-scored, and concatenated across subjects. That information is then subjected mathematical algorithms intrinsic to the HMM (as described in section 2.3.2) which

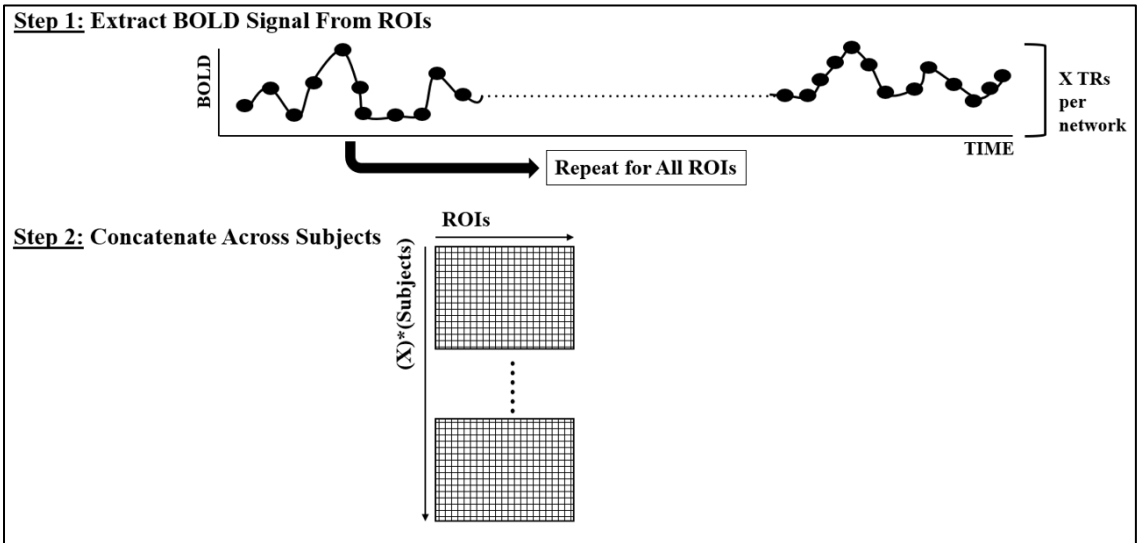

Figure 1: Illustration of the procedure used to implement AB HMM. Once the BOLD signal is extracted from all predefined ROIs, it is concatenated across all subjects and fitted with an HMM from the python hmmlearn library (55).

compute mean state patterns, a Viterbi path, a transition probability matrix, and covariance matrices for every state. This procedure is outlined in Fig. 1.

2.3.3.1.1 Inputs

AB HMM requires z-scored BOLD signal as an input, so no data reformatting beyond the description in previous sections is necessary (Fig. 1).

2.3.3.1.2 Outputs

Mean activation state patterns are obtained as a direct output from the model. As explained later in this chapter, we also acquired state patterns using the method Chen et al. 2016 employed to ensure that our outputted states were spatially consistent using two different techniques of procurement (11). Connectivity states corresponding to each activation state were acquired by mathematically transforming the covariance matrices into Pearson correlation values. (12,13,49).

2.3.3.2 Model Type 2: Summed Functional Connectivity Hidden Markov Model (SFC HMM)

Although converting the covariance matrices acquired from activation-based HMMs into Pearson correlations is an instituted method of defining connectivity states, this technique did not satisfy our interest of comparing BOLD-based HMM outputs to connectivity-based HMM outputs. To account for this, we discuss the first of our two connectivity-based HMMs: SFC HMM. This method is adapted from Ou et al. 2014 and is an established method appropriate for comparison to AB HMM because it acquires connectivity states stemming from connectivity inputs (28).

2.3.3.2.1 Inputs

As seen in Fig. 2, Ou et al. 2014 performed a sliding time window analysis (window length Δt) to obtain an ROI x ROI connectivity matrix within each time window (28). This generated a "dFC time series" of length (# TRs – Δt) representing the dynamics of functional connectivity over time. These connectivity matrices were summed across a dimension to create a vector depicting the overall connectedness of each node to all other

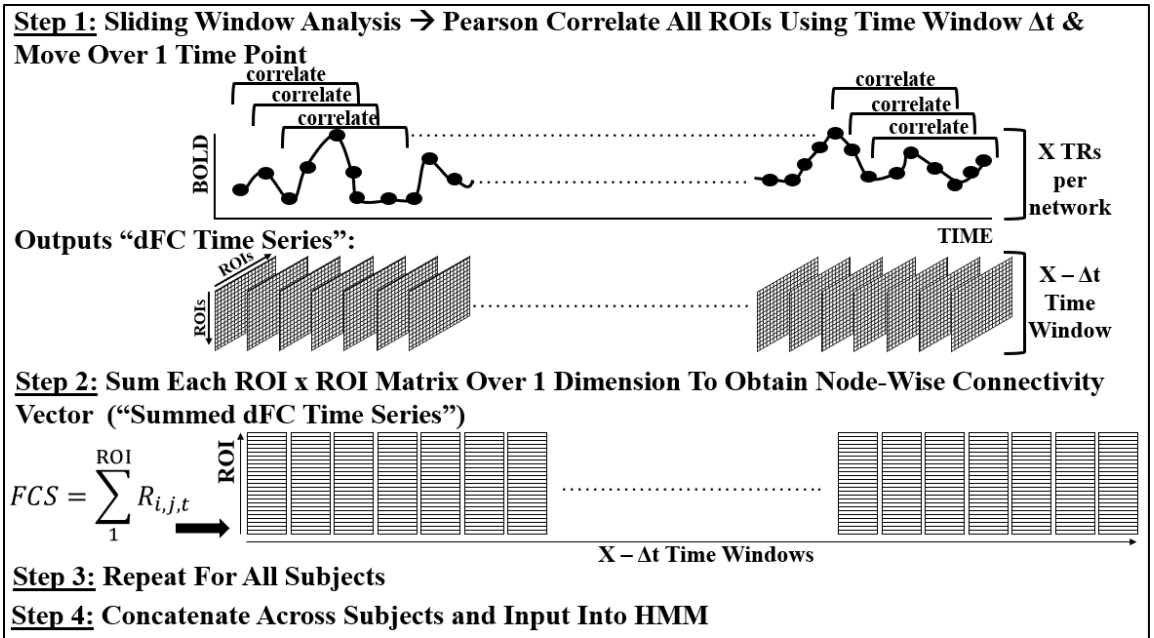

Figure 2: Illustration of the procedure used to implement SFC HMM as followed by the method described in Ou et al. 2014 (33). Once a sliding time window correlation analysis is performed, the connectivity matrix in each time window is summed across a dimension into a representative nodal connectivity 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 (55).

nodes. Repeating this for every time window provided a "summed dFC time series" containing a (# time windows) x ROI data matrix for every subject. Ou et al. 2014 then performed a two-stage hierarchical clustering analysis prior to concatenating this information across subjects to reduce the computational cost of performing the HMM (28).

One decision about this model left to the discretion of the user is the time window length (Δt) employed in the sliding window correlation analysis. Ou et al. 2014 empirically determined that a window size of 14 time points was suitable for their dataset (28). Lurie et al. 2020 discussed that new evidence shows an ideal window size might contain less than 60 seconds of data (14). Because Δt might speak to the temporal resolution overall, we compared SFC's state patterns from a sliding window approach (and resulting "dFC time series") using a window size containing between 20 and 60 seconds of data to assess its effect on the temporal resolution. Although we inputted Pearson correlations of z-scored BOLD signal rather than that of the raw BOLD, we do not expect the model outputs were affected. Z-scoring simply removes the mean and scales the BOLD signal, so it does not influence the direction of fluctuations that an ROI's signal exhibits.

2.3.3.2.2 Outputs

Unlike those for AB HMM, SFC HMM state patterns are not directly outputted from the model. Instead, the ROI x ROI connectivity matrices from the "dFC time series" are averaged across all time windows where the SFC HMM Viterbi path labeled a state to be active. Repeating this for every state provides a connectivity state profile. As shown in Section 3.3.2.2, the raw connectivity states look very similar to one another. To identify defining points and unique properties from each state, highlights from each connectivity state were observed and analyzed. This was done by calculating the grand mean of all state patterns except for the one in question and subtracting it from the raw connectivity patterns (13). This is shown in greater detail in Fig. 14 in Section 3.3.2.2.

As seen in the previous section, the mean patterns AB HMM outputs are the mean activation values that comprise each state and are based on activation levels observed in each node during each TR. Contrarily, SFC HMM's mean pattern and covariance outputs are unsuitable for the purposes of this investigation and for our analyses. This HMM subtype's outputted mean patterns are vectors describing mean summed correlation values because they are based on the inputs illustrated in Fig. 2: vectors of summed correlations representing global nodal strength during a particular time window. However, because we aim to have full connectivity matrices as states, converting the 1 x ROI vector of mean sums outputted into an ROI x ROI matrix of R^2 values is not feasible. Because SFC HMM's outputted covariances matrices are similarly dependent on the summed correlations inputted, they cannot be converted to R^2 values and therefore cannot be used to obtain connectivity state patterns in the same manner as the AB HMM counterparts when covariances were transformed into correlations. Instead, state patterns will be obtained by averaging dFC connectivity matrices wherever the SFC HMM Viterbi path labeled a state to be active.

2.3.3.3 Model Type 3: Full Functional Connectivity Hidden Markov Model (FFC HMM)

While SFC HMM is an established method, the HMM is fit to a summary statistic (summed connectivity values within each ROI x ROI matrix from the "dFC time series") rather than to actual connectivity values. This results in the HMM identifying changes in this summary measure rather than in the raw correlations forcing the user to create the states themselves rather than the model directly outputting it. To account for this, we

discuss a third HMM model type, FFC HMM, fitted to all correlation values in the lower (or upper) triangle of the dFC matrix in every time window. FFC HMM should also aid in determining the effect that summing the R^2 values across a dimension of the matrices from the "dFC time series" has on the state patterns.

2.3.3.3.1 Inputs

The sliding window correlation analysis is still performed with window length Δt, but instead of summing across one of the dimensions of the ROI x ROI matrices, the lower (or upper) triangle of R^2 Pearson correlation values is restructured into a 1 x $\frac{(\#\ ROI)^2-(\#\ ROI)}{2}$ vector (Fig. 3). Repeating this for every time window gives a (# time windows) x $\frac{(\#\ ROI)^2-(\#\ ROI)}{2}$ data matrix for every subject. which is concatenated subject-wise.

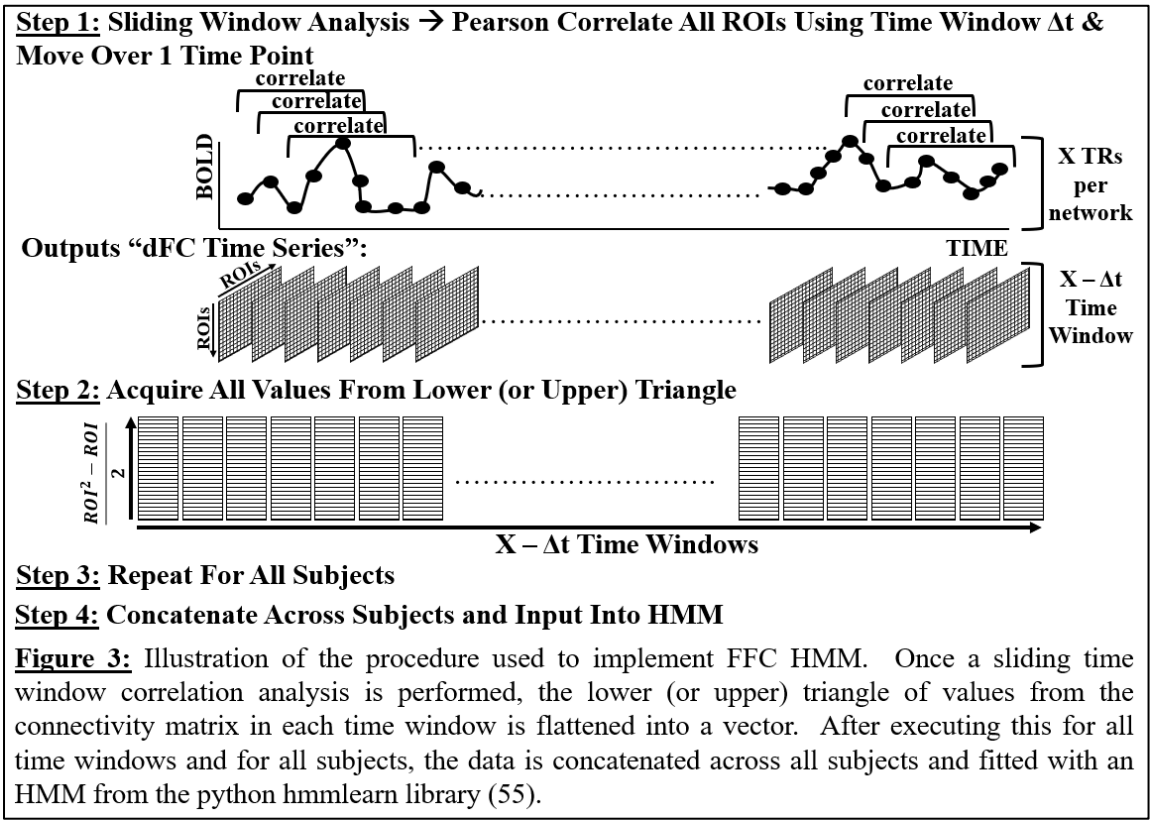

Figure 3: Illustration of the procedure used to implement FFC HMM. Once a sliding time window correlation analysis is performed, the lower (or upper) triangle of values from the connectivity matrix in each time window is flattened into a vector. After executing this for all time windows and for all subjects, the data is concatenated across all subjects and fitted with an HMM from the python hmmlearn library (55).

FFC HMM will also utilize a Δt of approximately 40 seconds since it stems from the same sliding window analysis as SFC HMM. The effect of window size and temporal resolution on FFC HMM state patterns were also assessed by acquiring them when window sizes containing 20 to 60 seconds of data were used.

2.3.3.3.2 Outputs

Unlike that for SFC HMM, FFC HMM state patterns are directly produced from the model. The mean patterns that FFC HMM provides mean connectivity patterns corresponding to the 1 x $\frac{(\#\ ROI)^2-(\#\ ROI)}{2}$ correlation vector inputted for every time window. These patterns are then reformatted back into a symmetric ROI x ROI matrix to constitute the connectivity states. Because R^2 values were used as inputs, FFC HMM was capable of directly producing spatial patterns related to raw correlations from each nodal pair and eliminated the need to manually create the state patterns. As with SFC HMM, the highlights were found for each connectivity state and used for all subsequent analyses.

2.3.3.4 Summary of HMM Model Subtypes

In summary, we defined three different HMM subtypes, their inputs, and their outputs. AB HMM requires z-scored BOLD signal as its input and activation states are acquired as direct outputs of the model. SFC HMM's input is a vector of summed R^2 values within each time window after a sliding window analysis and its states are solely defined as the global average of dFC matrices from time points where its Viterbi path labels a state to be active (28). FFC HMM necessitates a vector of all R^2 values from a sliding window analysis as its input and its connectivity states are straightforwardly outputted from the model. The connectivity-based HMMs will utilize a window size of ~40 seconds in

accordance with ideas from Lurie et al. 2020 and will be moved over one time point to maximize temporal resolution (14). Furthermore, all analyses hereafters will be performed on the highlights of each connectivity state.

Using the methods outlined above, we compared different measures of interest within and across HMM model types. We theoretically analyzed differences between model types as well as stability analyses that can be applied to each. Then we delineated local and global analyses that can be applied to fMRI data.

2.3.4 Model Type Comparison

Now that we have introduced and defined each HMM subtype in detail, we discuss important differences and commonalities to all. Main differences between each model type lie in the temporal resolution and in our inability to compare state patterns, while the method of model order determination and measures of interest are common to all.

2.3.4.1 Differences Between All Model Types

2.3.4.1.1 Temporal Resolution

One major discrepancy between the activation-based and connectivity-based HMMs is the temporal resolution of the data inputted into the models. The sliding window correlation analysis forces the information inputted into the connectivity-based HMMs to contain Δt fewer time points than that of the BOLD data and AB HMM. Because we choose to move our time window over only one time point and because each window contains approximately 40 seconds of data, autocorrelation inevitably artificially smooths the data across time. Furthermore, transforming the data from normalized intensity to Pearson correlations compresses the scale of the data to strictly range from -1 to 1 rather

than allowing the inputted values to freely range as with AB HMM's inputs. This reduced scope of inputs could potentially aggravate the temporal smearing as changes in correlation values may not be as exaggerated as those in BOLD signal patterns and possibly may not be recognized by the model. Examining all three model types in parallel allows us to observe the effect that this inconsistent temporal resolution has on measures of interest.

While FFC HMM has a temporal resolution equal to that of SFC HMM, both have a poorer resolution than AB HMM, which we expect to be reflected in the HMM outputs. Because the Viterbi path shows the sequence of states that subjects reside in at each time point, this measure ought to overtly exhibit the effects of poorer temporal resolution, smoothing, and summing the connectivity values. Specifically, we expect that switching between each state is reduced and that the Viterbi path overall to show temporal smoothing resulting in reduced transitions between states in the connectivity-based HMMs. Although the FFC HMM is still expected to exhibit these characteristics, we expect it to be less overt since a fuller, more holistic, portrayal of the data is fitted on the basis that all R^2 values are fitted per time window, thereby accounting for some of the phenomena mentioned.

2.3.4.1.2. Activation vs. Connectivity State Patterns

All three model types inherently detect different patterns of underlying latent brain states. However, direct comparison between the activation states arising from AB HMM with the connectivity states stemming from SFC HMM and FFC HMM is unfeasible. With the tools and methods presented in this dissertation activation states cannot be directly linked to connectivity states purely based on the levels of either measure. Designing a mechanism to do so is outside the scope of this investigation. Furthermore, connectivity

states from SFC or FFC cannot be directly related to the covariance-based connectivity states from AB HMM. The HMM subtypes were fit to different data types so there is no a priori reason to believe that the activation-based and connectivity-based models recognize the same states. There is also no literature support or mathematical reasoning to suggest that all HMM sub-types should have the same model order, further reinforcing the idea that each identifies different states.

SFC states can potentially be linked to the FFC states. We have an a priori reason to believe that these states should show similarity because they both stem from the same ROI x ROI connectivity matrices formed in the sliding window analysis. The state patterns can be matched by Pearson correlating (11) or computing the Euclidean distance (56) between all states and determining which pairs have the highest/lowest values respectively.

2.3.4.2 Common to All Model Types

2.3.4.2.1 Model Order Determination (Stability Analysis)

2.3.4.2.1.1 Overview

When comparing state patterns within each model type for robustness and reproducibility, determining the optimal number of states naturally arises. Employing too many states could create an unstable model where the state patterns are no longer consistent across different model initializations or could potentially contain repeated states indicating a lack of parsimony (11,13). Although the model outputs emission probabilities and a model-based log likelihood, these measures are not typically used for assessing neuroimaging-based HMM state patterns (cite 11-13 here). Instead, a few groups examining HMMs in neuroimaging chose their model order by exploring the free energy

for a range of model orders (12,13). Where the free energy levels plateaued indicated the optimal model order for the dataset in question. Stevner et al. 2019 examined the median fraction occupancy (the median value of the average time spent in all states) for model orders 3 to 45, and where this value leveled off signified a model order suitable for their investigation (13). One group chose to pursue 12 states simply because Vidaurre et al. 2017 decided that was best for their investigation despite examining a completely different dataset (57). The overall consensus is that any model order chosen is acceptable as long as there is a valid justification for that choice. These groups utilized the HMM-MAR Matlab toolbox to implement their HMMs and were able to easily calculate the free energy as it was one of the toolbox's included functions (12). However, this toolbox did not provide us with the reproducibility of state patterns and state sequences we aimed for, so we opted to use the hmmlearn library in python instead. Furthermore, we chose not to examine the free energy or median fractional occupancy because previous groups who had utilized the hmmlearn python library had evaluated the robustness and reproducibility of acquired state spatial patterns and we wished to remain consistent with their work (cite 55 here). We therefore employed two different methods to assess whether each HMM subtype's states converged to the same pattern despite the initialization probabilities used while simultaneously deciding model order. For SFC and FFC HMMs, these analyses were performed on the connectivity state highlights rather than on the raw connectivity states.

2.3.4.2.1.2 Method 1: RAICAR-based Stability Analysis

The first method is taken from Chen et al. 2016 and adapted from Yang et al. 2010's Ranking and Averaging Independent Component Analysis by Reproducibility (RAICAR)

method (11,35). This method was developed to assess the optimal number of parcellations in ICA but has since been used to determine the optimal number of states in HMMs fitted to fMRI data (11,35). We used a RAICAR-based method to evaluate the reproducibility of states across different HMM initialization parameters where stability values above a threshold was considered stable. As depicted in Fig. 4, the HMM was run thrice (three "realizations") each with a different starting probability of residing in a certain state so that three sets of state patterns were obtained. State patterns across realizations were matched via the highest Pearson correlations computed and relabeled so that each state assignment universally corresponded to the same spatial pattern. Within each state assignment, the matched state patterns were Pearson correlated between all realizations to obtain $\frac{(\#\ states)!}{2!(\#\ states-2)!}$ R^2 values thereby determining the degree of similarity between the matched patterns. These values were then averaged, sorted from largest to smallest, and plotted against model order. Repeating this for a range of model orders generated a set of bar

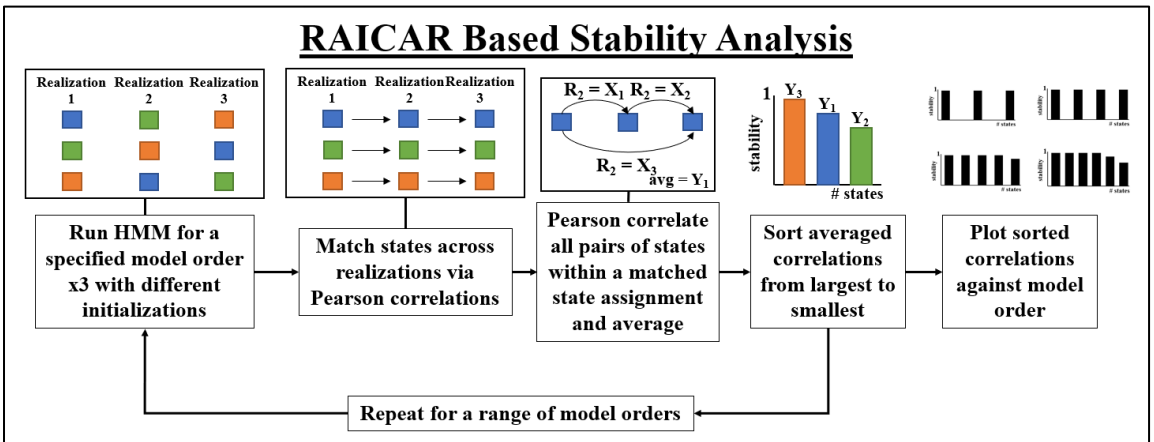

Figure 4: Schematic illustrating the procedure for the RAICAR-based model order determination method. Each of the blocks shown represent a state where matched states are shown in the same color. A basic example of 3 states (model order 3) is shown where states' spatial patterns are first matched across realizations using Pearson correlation values (11,34). 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 model order. This procedure is then repeated for all states within a given model order.

graphs whose heights were compared against a predetermined threshold of stability. Model orders where R^2 values began to dip below a threshold indicated that the model of that order was unstable because the states were no longer matching sufficiently. Thus, the model order was too large by 1 (11,35).

2.3.4.2.1.3 Method 2: ED-based Stability Analysis

The second stability analysis performed (Fig. 5) uses Euclidean distances to assess the reproducibility of states where the smaller the value the more similar the state patterns. In this Euclidean distance-based (ED-based) stability analysis state assignments from two realizations, R_i and R_j, within a certain model order are permuted and their spatial patterns matched via the smallest Euclidean distance such that each state universally corresponds to the same spatial pattern. For example, after permuting the state assignments from R_i and R_j, the Euclidean distance between state 1 from R_i and all states from R_j are computed. The results may show that the smallest Euclidean distance was computed with state 5 from

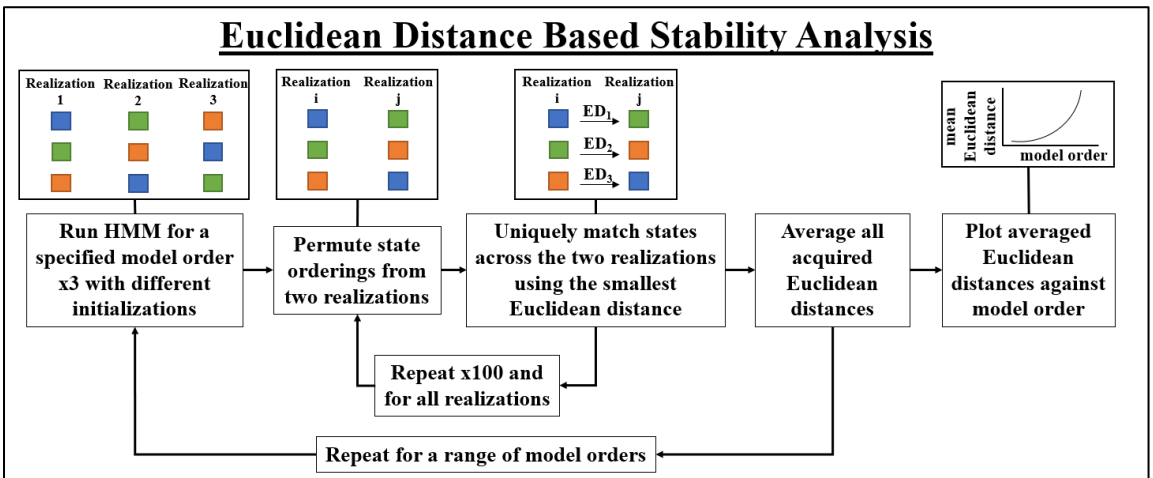

Figure 5: Schematic of the ED-based model order determination method applied to fMRI-based HMMs. Each of the blocks shown represent a state where matched states are shown in the same color. An example with 3 states is shown (11) where state assignments are permuted and matched by determining the smallest possible Euclidean distance between spatial patterns from different realizations. This procedure is repeated 100 times and for states across all pairs of realizations. The mean Euclidean distance from all permutations and all matchings (3 * 100 * model order) are averaged and plotted against model order. This procedure is repeated for a certain range of model orders.

R_j indicating that state 1 from R_i best corresponds to state 5 from R_j. State 5 from R_j is relabeled as state 1 from R_j, then the process is repeated between state 2 from R_i and all states from R_j (except the relabeled state 1 as it has already been matched). This permutation-and-matching procedure is performed 100 times for all pairs of realizations (where matching $R_i \rightarrow R_j$ is not distinguished from matching $R_j \rightarrow R_i$) to ensure that the spatial patterns are paired uniquely without any bias of state assignment. This is repeated for a range of model orders generating a total of 100 * ($\frac{(\#\ states)!}{2!(\#\ states-2)!}$ realizations) * (# states) values for a particular model order which are then averaged to represent its overall stability. This single average is plotted as a function of model order producing a curve where the smallest value succeeded by continuously increasing values for higher model orders indicates the optimal number of states for a dataset. Although this this may seem similar to the RAICAR-based method, the ED-based method is more conservative because stability is assessed by ensuring that the average of hundreds of Euclidean distances is below a prespecified threshold rather than the average of a handful of R^2 values.

2.3.4.2.2 Measures of Interest

As aforementioned, each HMM subtype outputs four measures of interest: mean state patterns, a Viterbi path, covariance matrices, and a transition probability matrix. These outcomes are then used to derive switching frequency and fractional occupancy correlation. Theoretically these measures should be immediately usable for analyses, with the exception of the mean state patterns and covariance matrices outputted from SFC HMM (refer to Section 2.3.4.2.2). The smoothing effect inherent to the connectivity models resulting from the sliding window approach might impact the signification of these

measures. The extent to which these measures can provide insight cannot be confirmed until the HMMs are actually fit to the data, so these measures will be reassessed then.

2.3.5 Types of Analyses

Now that all three HMM types have been defined and their theoretical commonalities and differences delimited, we outline analyses that provide insight into empirically understanding their similarities and differences. Local analyses examine measures specific to each model subtype: state patterns. Thus, local analyses will focus on characterizing states and on evaluating activation or connectivity spatial patterns obtained via information from other HMM subtypes. This is a means to better understand each subtype's ability to capture the dynamics of inputted information. Global analyses examine measures common to all three model types and builds upon local analyses. Comparisons of these measures between model types were made as a means of obtaining insight into differences amongst them.

2.3.5.1 Local Analyses

Local analyses refer to state pattern characterization and state acquisition comparison. Because we examined two different types of fundamental HMMs, we focus on two types of local analyses. The first type of local analysis concentrates on activation state patterns. We acquired activation patterns using the Viterbi paths from all HMMs and compared them to AB HMM's direct output. The second local analysis focuses on connectivity state patterns attained from average connectivity matrices where AB and FFC HMM Viterbi paths (in accordance with methods described in Sections 2.3.5.1.1 and

2.3.5.1.2). This was followed by comparing SFC and FFC HMM state patters with AB HMM derived covariance-based connectivity matrices (as detailed in Section 2.3.3.1.2).

Even after utilizing each HMM type to identify all types of states, it is likely that not recognizing every single state that could possibly exist in a dataset was recognized. There may exist subject-specific states where a single subject might reside in a state that only they experience (13). For data collected during non-rapid eye movement sleep, there existed electroencephalogram-based states associated with various sleep stages potentially leading to different states being active under different conditions (13). Our HMMs were fitted on a group level to find the maximum number of global states that are visited by all subjects. This group-level analysis may ignore subject- or condition-specific states, but, for comparative purposes, we intended to focus on broad, comprehensive states, and can pursue investigation of less populous states in later investigations.

2.3.5.1.1 Activation State Patterns

After determining optimal model order using the methods detailed in Section 2.3.4.2.1, state patterns were analyzed to obtain better insight into the model and to make neuroscientific interpretations. Although AB HMM states are directly outputted from the model, Chen et al. 2016 defined their states using a different method which we also employed to ensure that our state patterns were consistent across various methods of acquisition (11). As seen in Fig. 6, this method extracts and averages the BOLD signal from all time points where the state sequence labeled a state to be active. Repeating this for all states provides a recreation of the mean activation state patterns. We also attempted to recreate the activation states by averaging the BOLD signal where the SFC and FFC

HMM Viterbi paths label a state to be active (Fig. 7). This would determine whether the model types were in fact recognizing different states and ensure that we were not acquiring repeated information from different HMM subtypes. We used both techniques in parallel as a sanity check to confirm that our three model types are distinct and are able to identify changes only in the respective data type (BOLD vs. correlations) inputted. The similarities of these states to each other were assessed via correlations and Euclidean distances.

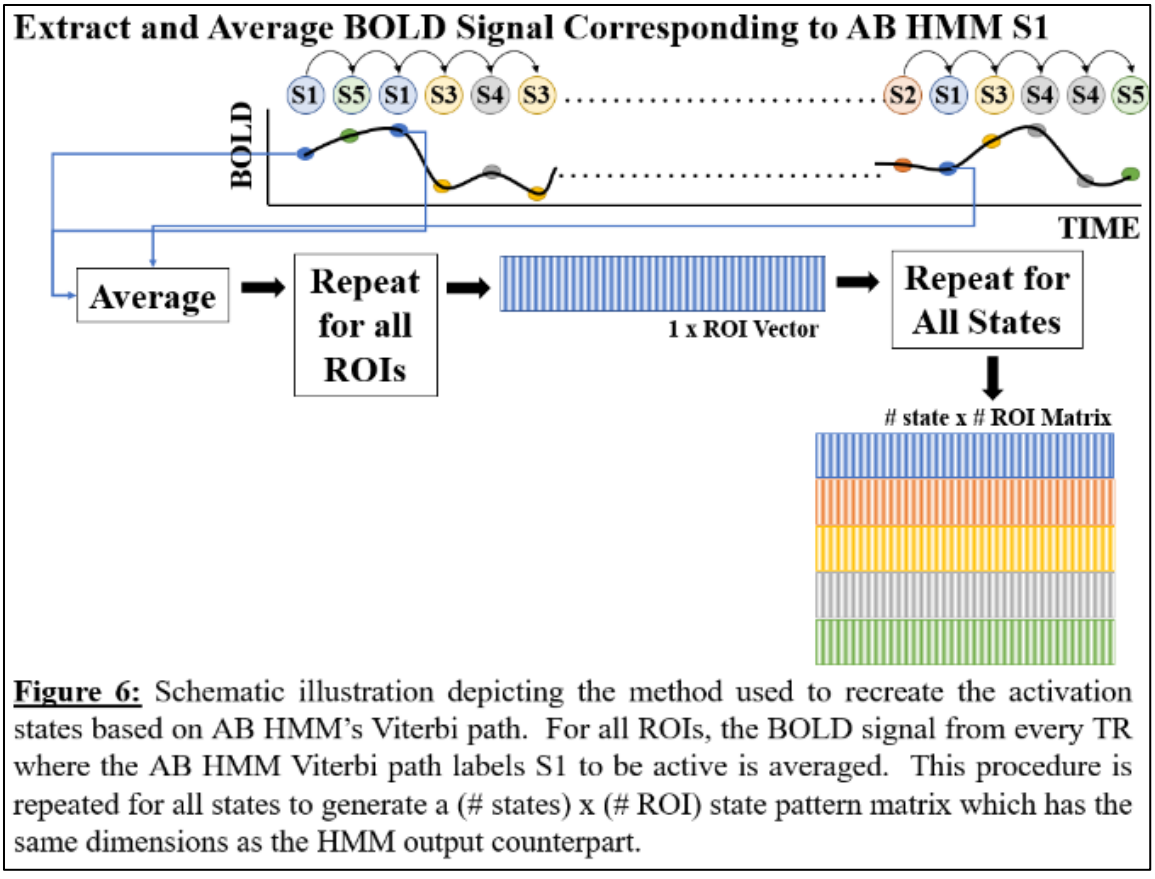

Figure 6: Schematic illustration depicting the method used to recreate the activation states based on AB HMM's Viterbi path. For all ROIs, the BOLD signal from every TR where the AB HMM Viterbi path labels S1 to be active is averaged. This procedure is repeated for all states to generate a (# states) x (# ROI) state pattern matrix which has the same dimensions as the HMM output counterpart.

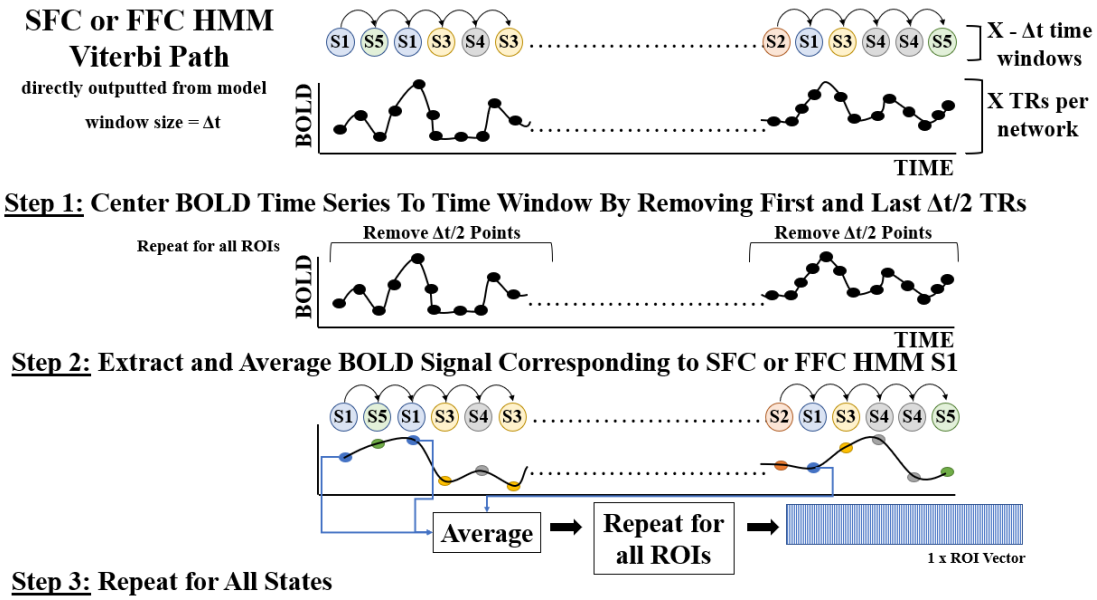

Figure 7: Schematic illustration outlining the method used to recreate the activation states using the Viterbi paths from SFC and FFC HMMs. This is the same method depicted in Figure 6 with the discrepancy in temporal resolution being accounted for by removing some TRs of the BOLD time series to ensure equal length between the BOLD signal and the connectivity-based HMMs' Viterbi paths.

2.3.5.1.2 Connectivity State Patterns

For FFC HMM, the state patterns are directly outputted from the model while SFC HMM acquires state patterns by averaging connectivity matrices where its Viterbi path recognizes a state to be active. As the methods in Fig. 8 depict, the connectivity matrices at every point where the FFC and AB Viterbi paths label a state to be active were also averaged to observe the connectivity patterns arising from this method. The covariance matrices from AB HMM are converted into Pearson correlations as a means of obtaining connectivity states from a BOLD-based HMM. We do not expect any of the connectivity patterns stemming from AB HMM (using either the Viterbi averaging method or the covariance-based correlation method) to be identical to ones identified from SFC and FFC HMMs. The model subtypes should not be interpreting the inputs in the same way since

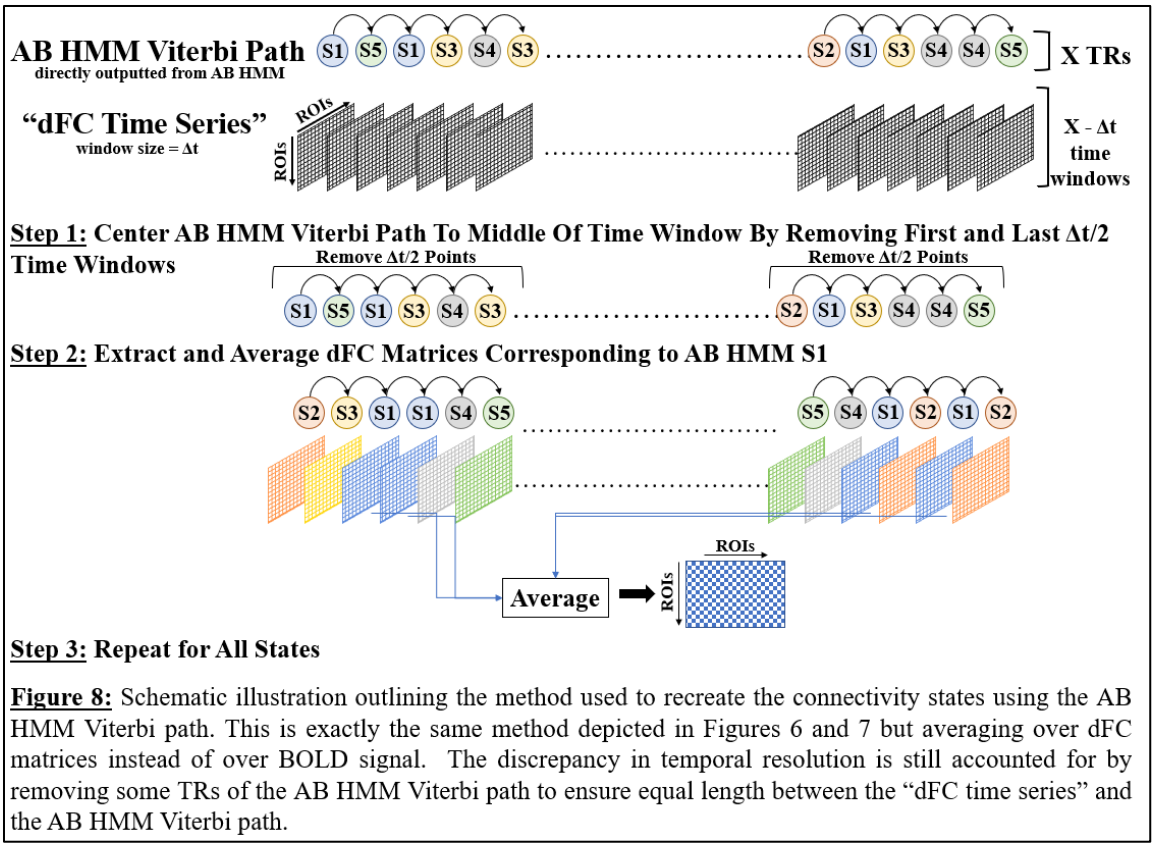

Figure 8: Schematic illustration outlining the method used to recreate the connectivity states using the AB HMM Viterbi path. This is exactly the same method depicted in Figures 6 and 7 but averaging over dFC matrices instead of over BOLD signal. The discrepancy in temporal resolution is still accounted for by removing some TRs of the AB HMM Viterbi path to ensure equal length between the "dFC time series" and the AB HMM Viterbi path.

different data types (intensity vs. connectivity) were used. This set of analyses provide us with an approach allowing direct comparison across all model types such that any discrepancies observed are not due to the difference in method of state spatial pattern extraction, but instead due to the differences in the fitted models themselves indicating that the HMM subtypes employed are behaving uniquely.

2.3.5.2 Global Analyses

Global analyses refer to measures that can be directly obtained or calculated from AB HMM, SFC HMM, and FFC HMM without tailoring the method of acquisition for a specific subtype. We previously introduced four measures of interest that comprise the four global analyses we focused on: Viterbi paths, transition probabilities matrices, switching rates, and fractional occupancy correlations. After acquiring these measures

from each HMM subtype, they were compared to obtain insight into similarities and differences between model type behavior.

2.3.5.2.1 Viterbi Path

The Viterbi path, or hidden state sequence, is directly outputted from all three HMM subtypes and is used to examine switching rate, proportion of time spent in each state, the average duration of a state, and fractional occupancy correlation (12,13). The Viterbi path can be visualized by assigning each state a color and plotting them for every person as a function of TR (AB HMM) or of time window (SFC and FFC HMM) to paint a qualitative picture of which states are active at each time point.

2.3.5.2.2 Transition Probability Matrix

As with the Viterbi path, the transition probability matrix of each HMM subtype is directly outputted. This parameter provides quantitative insight into the relationship between states and can be visualized by observing all transition probabilities on a color scale.

2.3.5.2.3 Switching Rate

Other groups have used switching rate as a distinguishing property between two or more entities: metastates for Vidaurre et al. 2017 and non-rapid eye movement sleep stages for Stevner et al. 2019 (12,13). We use switching rate as a means to further classify temporal differences between our three HMMs of interest. We calculated it between all states within an HMM subtype for comparison across model types to determine whether one is able to capture these dynamics better than the others. This measure is found within each subject by counting the number of times that subjects switched between states and

then dividing the value by the length of the entire time series. Averaging these values across all subjects provides a global depiction of the switching rate between unspecified states, and boxplots can be used to visualize them.

2.3.5.2.4 Fractional Occupancy Correlation

Fractional occupancy correlation is a measure that Vidaurre et al. 2017 has analyzed in detail and is found by correlating the average time all subjects spend in a certain state against the time spent in another state (12). This measure, along with the transition probability matrix, can potentially provide insight into the dynamic relationship between latent brain states. Studying fractional occupancy correlations would aid in obtaining a more quantitative idea of the relationship between time spent in states beyond simply observing the proportion of time spent in a single state via boxplots. Once the fractional occupancy correlations are found, they are visualized in a matrix where all values are placed on a color scale.

Furthermore, Vidaurre et al. 2017 found that metastates (a group of states that make up a network) emerge organically from the plots of fractional occupancy correlations created from intensity-based HMM states (12). If HMM subtypes' model orders are large enough to permit, fractional occupancy correlations can potentially identify metastates and group activation or connectivity states together in a hierarchical fashion. It is possible that the fractional occupancy correlations of the identified states from our three HMMs might provide insight into which model type to analyze when approaching a specific neuroscientific question. While we do not dive deeply into the quantitative aspects of these

analyses, we qualitatively assess the plots of fractional occupancy correlations to hypothesize whether each HMM subtype could potentially observe metastates.

2.3.6 Overarching Questions

Now that we have described in detail the different HMMs of interest, discussed their similarities and differences, and outlined measures of interest we intend to attain, we are left with the following questions about HMM subtype behavior that cannot be answered solely from theory and methodology:

(1) What kinds of questions does each model type answer?

(2) When is it useful to employ each model type?

To obtain insight into these questions, we fit all three HMMs to a large publicly available fMRI dataset to qualitatively and quantitatively compare their outputs.

2.3.7 Dataset and Networks

2.3.7.1 Dataset

The Human Connectome Project (HCP) Unrelated 100 (a subset of the S500 release) dataset is suitable for our investigation because it is readily available and robust dataset (58). The data underwent the HCP minimally preprocessing pipeline: distortion correction, motion correction, alignment to standard space, and surface projection (59). 100 subjects (age = 22–36, gender = 54 female) underwent a 14.4-minute resting state scan (repetition time = 720ms, flip angle = 52°, voxel size = $2mm^3$, echo time = 33ms, field of view = 208mm x 180mm). During early iterations of the HMMs, one subject was found to remain in a single subject-specific state that almost no one else visited for most of the

scan. Removing that subject did not affect the model order chosen (or any other parameter; data not shown) so we conducted all analyses on the remaining 99 subjects.

2.3.7.2 Networks

We narrowed our scope of analysis to pre-defined networks that have previously been associated with resting state. BOLD signal was measured from four networks: 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. Talairach coordinates for DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and converted to Montreal Neurological Institute (MNI) coordinates. MNI coordinates for SN were taken directly from Raichle 2011 (1,60–63). The MNI coordinates used in this investigation can be found in Table 1. Dorsal anterior cingulate cortex and left dorsolateral prefrontal cortex ROIs in FPCN were excluded from analyses because they overlapped or were too closely located to other ROIs. Thus, we only used 29 ROIs when examining the HCP dataset. After labeling each ROI with a $5mm^3$ isotropic marker, the BOLD signal was extracted from each voxel and averaged 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 (64–66).

	ROI	Abbreviation	Full Name	MNI Coordinates	Source
Default Mode Network	1	PCC	Posterior Cingulate Cortex	(2, 54, 16)	Despande et al. 2011
	2	L pIPL	Left Posterior Inferior Parietal Lobule	(-46, -72, 28)	Despande et al. 2011
	3	R pIPL	Right Posterior Inferior Parietal Lobule	(50, -64, 26)	Despande et al. 2011
	4	PFC/vACC	Orbitofrontal Cortex/Ventral Anterior Cingulate Cortex	(4, 30, 26)	Despande et al. 2011
	5	dMPFC BA 8	Dorsomedial Prefrontal Cortex Broadmann Area 8	(-14, 54, 34)	Despande et al. 2011
	6	dMPFC BA 9	Dosomedial Prefrontal Cortex Brodmann Area 9	(22, 58, 26)	Despande et al. 2011
	7	L DLPFC	Dorsolateral Prefrotnal Cortex	(-50, 20, 34)	Despande et al. 2011
	8	L PHG	Parahippocampal Gyrus	(-10, -38, -2)	Despande et al. 2011
	9	L ITC	Inferolateral Temporal Cortex	(-60, -20, -18)	Despande et al. 2011
Fronto-Parietal Control Network	1	L aPFC	Left Anterior Prefrontal Cortex	(-36, 56, 10)	Deshpande et al. 2010
	2	R aPFC	Right Anterior Prefrontal Cortex	(34, 52, 10)	Deshpande et al. 2010
	~~3~~	~~dACC~~	~~Dorsal Anterior Cingulate Cortex~~	~~N/A~~	~~Deshpande et al. 2010~~
	~~4~~	~~L DLPFC~~	~~Left Dorsolateral Prefrontal Cortex~~	~~N/A~~	~~Deshpande et al. 2010~~
	5	R DLPFC	Right Dorsolateral Prefrontal Cortex	(46, 14, 42)	Deshpande et al. 2010
	6	L aINS	Left Anterior Insula	(-30, 20, -2)	Deshpande et al. 2010
	7	R aINS	Right Anterior Insula	(32, 22, -2)	Deshpande et al. 2010
	8	L aIPL	Left Anterior Inferior parietal Lobule	(-52, -50, 46)	Deshpande et al. 2010
	9	R aIPL	Right Anterior Inferior Parietal Lobule	(52, -46, 46)	Deshpande et al. 2010
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 1: 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). Taliarch coordinates for DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and were converted to MNI using (60-63) while MNI coordinates for SN were taken directly from Raichle 2011 (1). The two FPCN coordinates with the strikethrough were excluded from analyses because of their close location to neighboring ROIs.

2.3.7.3 Additional Information

When implementing SFC HMM, we followed the steps outlined in Ou et al. 2014 with the exception of the clustering stage as it was not necessary in our investigation to cluster ROIs stemming from predetermined and predefined networks (28). We examine only 29 ROIs which is few enough that we can avoid clustering. Ou et al. 2014, on the

other hand, examined 358 ROIs acquired from DICCOL and implemented a hierarchical clustering stage to reduce computational cost.

Using this HCP dataset, we characterized the state patterns (Local Analyses) for each model type and for different time windows centered around a 50tp (approximately 36.6 seconds in accordance with information presented in Lurie et al. 2020) window size (14). We also examined the Viterbi path, transition probability matrix, switching frequency, and fractional occupancy correlation (Global Analyses) for each model type.

2.4 Conclusion (Methods Summary)

In this chapter, we introduced, defined, and theoretically analyzed three HMM sub-types (AB HMM, SFC HMM, and FFC HMM) that will be empirically investigated in the next chapter. We first reviewed potential methods of acquiring brain states and concluded that HMMs provide a satisfactory balance between obtaining brain state spatial patterns and temporal dynamics. We crucially noted that the three HMM sub-types of interest to use can be applied to fMRI, electroencephalography, and magnetoencephalography data making them conventional HMMs that can be applied to a variety of datasets. We intend to compare and contrast these three model types in effort to pinpoint advantages, disadvantages, and constraints in each in effort to understand potential questions that they can answer. With this information, we could then develop intelligible judgments about which HMM type is best to pursue under certain conditions in future analyses.

This chapter also aimed to understand information that could be theoretically gleaned from the different HMMs prior to applying them to resting state fMRI data. We discussed how the data in each HMM subtype is arranged and fitted in python to output

mean state patterns, a Viterbi path for every subject, a transition probability matrix, and covariance matrices for each. Although AB and FFC HMM's state patterns are directly outputted from python, they will be verified against the spatial patterns resulting from averaging the BOLD signal and dFC matrices, respectively, from every time point where the corresponding Viterbi path labeled a state to be active. Contrarily, SFC HMM's state patterns are constructed using only this latter process. For all model types, we will use the information in the Viterbi path to examine the switching frequency between each state as well as the fractional occupancy correlation. We also discussed that analyzing states' stability and establishing optimal model order work in conjunction: RAICAR- and ED-based stability analyses pinpoint the maximum number of states possible where the similarity of state patterns between realizations remains above a predetermined threshold.

We also introduced the HCP dataset and resting state networks (DMN, FPCN, DAN, and SN) of interest to us. In the next chapter, we apply all three HMM subtypes to this dataset using the BOLD signal extracted from a total of 29 nodes. We examine the outcomes in hopes of answering to the best of our abilities the overarching questions posed. In other words, the next chapter aims to empirically determine what information can be gleaned from applying AB HMM, SFC HMM, and FFC HMM to a resting state big dataset.

Chapter 3: Application of Hidden Markov Model Methods to an fMRI Dataset

3.1 Abstract

Using the methods and theory outlined in the previous chapter, we now apply three different instantiations of a hidden Markov model to a robust and publicly available neuroimaging dataset. Here we work to obtain empirical insight into all three model types in effort to understand each of their behaviors and differences. We described and interpreted the results of stability, local, and global analyses. As we examined each of these measures individually, we worked to gain experimental knowledge into the two overarching questions posed in the previous chapter. For a functional magnetic resonance imaging dataset, we found that a model measuring fluctuations in activation levels was best at capturing temporal dynamics, and connectivity-based models were better at identifying latent connectivity state spatial patterns. While all model types investigated provided useful information to an investigation, each is better applied under different scenarios which we explored empirically in this chapter.

3.2 Introduction

In Chapter 2 we introduced three hidden Markov model (HMM) subtypes of interest to us: activation based HMM (AB HMM), summed functional connectivity HMM (SFC HMM), and full functional connectivity HMM (FFC HMM). We described how each model is set up to detect changes in activation levels, in a summed representative vector of a dynamic functional connectivity (dFC) values from a sliding window approach, and in all possible correlation values in every dFC matrix, respectively. Each model subtype could be implemented using the hmmlearn library in python and output mean state

patterns, a transition probability matrix, a Viterbi path for every subject, and covariance matrices for each state (55). These four measures of interest were compared in theory across all HMM subtypes. We also discussed additional parameters to analyze in this chapter: switching rate and fractional occupancy correlation.

Now that we have discussed the theoretical aspect of our investigation as well as the corresponding methods and analyses, we next apply each of them to a dataset. In this chapter we fit all three HMM subtypes to the Human Connectome Project (HCP) functional magnetic resonance imaging (fMRI) dataset introduced in Section 2.3.8.1 and contrasted their outputs (58,59). We intended to gain empirical insight into the differing behaviors of each model subtype in effort to obtain insight into the two main questions previously posed: (1) What kinds of information does each model type provide insight into? (2) When is it useful to employ each model type?

In this chapter we first assessed the stability of each model type to determine the optimal model order of each. This was done via the Ranking and Averaging Independent Component Analysis by Reproducibility-based (RAICAR-based) and Euclidean distance-based (ED-based) methods illustrated in Section 2.3.5.2.1 (11,35). Once the best possible number of states was determined, we analyzed the activation states directly outputted from AB HMM as well as the activation spatial patterns resulting from averaging the blood oxygen level dependent (BOLD) signal of every region of interest (ROI) where the AB HMM, SFC HMM, and/or FFC HMMs labeled a state to be active. Then we focused on the connectivity states by examining the spatial patterns acquired for SFC and FFC HMMs when using a time window of 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) in the sliding window analysis to determine the effect that window size choice had on subsequent analyses. We also examined the state patterns resulting from averaging the connectivity matrices where the FFC and AB HMM Viterbi paths labeled a state to be active as well as covariance-based Pearson correlation matrices directly outputted from AB HMM. Subsequently, we evaluated the Viterbi paths, transition probability matrices, switching rates, and fractional occupancy correlations of each model subtype and related them back to our overarching questions. For each measure listed, we discussed how they were acquired and how they were visualized.

3.3 Results and Discussion

3.3.1 Model Order Determination (Stability Analyses)

3.3.1.1 Overview

Using the RAICAR-based and ED-based stability analyses outlined in Section 2.3.4.2.1, the optimal model order for each HMM subtype was determined. When performing the three realizations for AB HMM, SFC HMM, and FFC HMM, the same set of initialization probabilities were used for every model type: two sets were randomly chosen, and one used a uniform distribution with equal probability of starting in a particular state ($^{1}/_{\#\ states}$).

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 (11,35). Our analyses could afford to be more conservative because we explored only 29 ROIs and therefore appointed a threshold of 0.9.

Furthermore, Chen et al. 2016 utilized the same dataset and elected a 9-state model when using 236 ROIs (11). We had no a priori reason to believe that the optimal model order must be the same between AB, SFC, and FFC HMMs, but because we used substantially fewer ROIs then Chen et al. 2016, we did not expect our investigation to appoint more than nine states to any subtype.

3.3.1.2 Model Type 1: AB HMM

RAICAR-based results for AB HMM (Fig. 9a) indicated model order eight to be extremely stable since all stability values were 1; model order nine showed a drop in the stability value of one state to below the predetermined threshold. ED-based results (Fig. 9b) were consistent with this idea: an 8-state model was the largest model order possible while the mean Euclidean distances were close to zero and steadily increased for succeeding higher order models. Results from both methods indicated that AB HMM should be fitted with eight states.

It may appear that a model order as high as 12 could be suitable for AB HMM because its stability values remained above the predetermined threshold. However, model orders 10, 11, and 12 contained a repeated state making them undesirable indicating a lack of parsimony 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 (11). The 10-, 11-, and 12-state models included two occurrences of this activation pattern; although these model orders 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 AB HMM in this investigation as a result of examining

both stability analyses in conjunction and because it met our criterion of choosing fewer than nine distinct states in accordance with Chen et al. 2016's findings (11).

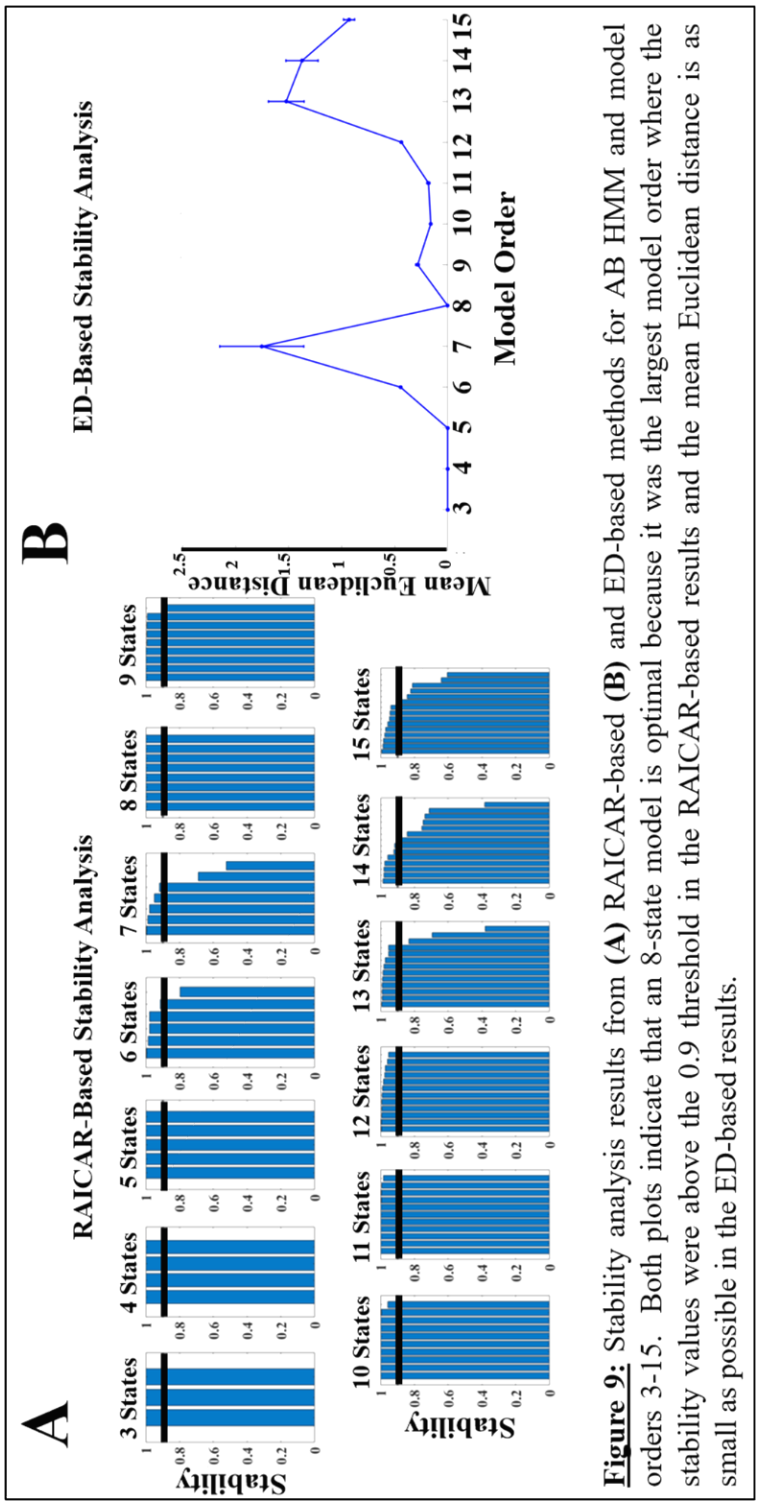

Figure 9: Stability analysis results from **(A)** RAICAR-based **(B)** and ED-based methods for AB HMM and model orders 3-15. Both plots indicate that an 8-state model is optimal because it was the largest model order where the stability values were above the 0.9 threshold in the RAICAR-based results and the mean Euclidean distance is as small as possible in the ED-based results.

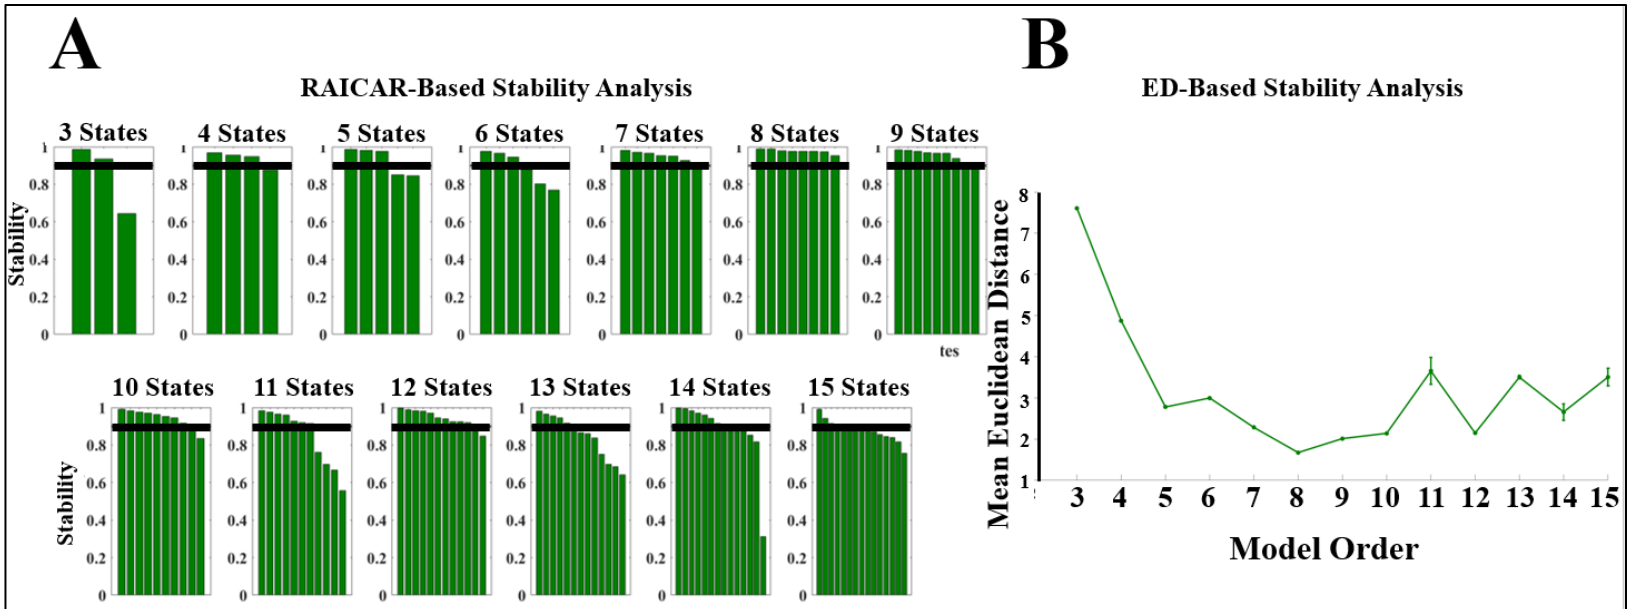

Figure 10: Stability analysis results from **(A)** RAICAR-based **(B)** and ED-based methods for SFC HMM and model orders 3-15. Both plots indicate that an 8-state model was optimal because this was the largest model order before the RAICAR-based stability values fell below the 0.9 threshold, and the Euclidean distances began to continuously increase.

3.3.1.3 Model Type 2: SFC HMM

Fig. 10a and Fig. 10b respectively show the results of the RAICAR-based and ED-based methods performed on the SFC state highlights. As described in Section 2.3.3.2.2, connectivity state highlights were acquired by subtracting the grand mean of all other state patterns from each state to observe each one's defining features. RAICAR-based results indicated that an 8-state model had the highest number of robust and reproducible states since model orders nine and above encompassed stability values below the 0.9 threshold. Similarly, ED-based results showed an 8-state model to be the largest model order possible before the Euclidean distances steadily increased suggesting that model orders nine and above were unstable. Thus, eight states were best to pursue for SFC HMM in this investigation.

3.3.1.4 Model Type 3: FFC HMM

Strictly for the purposes of comparison with SFC HMM, we assigned FFC HMM to have eight states as well. This allowed for a level comparison of local and global analyses across both connectivity-based HMMs. We acknowledge that eight states may not actually be suitable for FFC HMM, and that the same model order between SFC and FFC HMMs may not always occur in a study. Future investigations should perform all stability analyses described in Section 2.3.5.2.1 (RAICAR-based and ED-based methods).

To verify that increasing FFC HMM's model order past eight would decrease stability, FFC HMM robustness and reproducibility were examined for only model orders eight and nine. Fig. 11 shows that both methods indicated an 8-state model was preferable for FFC HMM. Model order nine was unstable as observed by the increase in mean

Euclidean distance and the drop in stability values below 0.9, which were likely not seen in any other model order. This confirmed that an 8-state model for FFC HMM was preferred over a 9-state model.

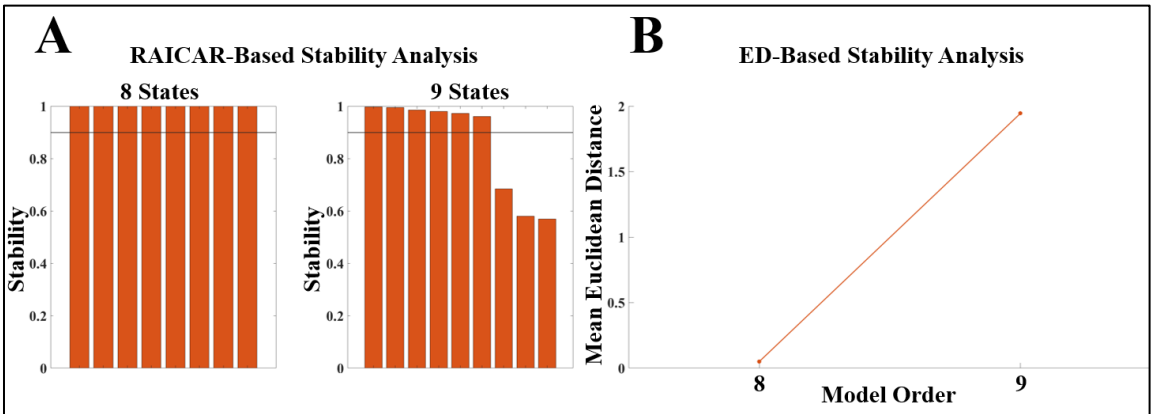

Figure 11: Stability analysis results from **(A)** RAICAR-based **(B)** and ED-based methods for FFC HMM and model orders 3-15. For better comparison with SFC HMM, FFC HMM was assigned a model order of eight and realizations for only model orders eight and nine were acquired. An 8-state model is acceptable for FFC HMM since increasing the model order to nine resulted in a drop in stability below the predetermined threshold in the RAICAR-based results and in a dramatic spike in Euclidean distance in ED-based results.

Because of these results and because eight states were in our predetermined range (less than or equal to nine states) for optimal model order, we pursued eight states for all model types. There is no literature or mathematical reasoning to suggest that all three HMM subtypes should have the same model order within a particular investigation. We coincidentally determined eight states to be best for AB HMM and SFC HMM, but that might not always be the case. This may have occurred because we focused on a relatively small number of ROIs which greatly narrowed our analyses, but a consistent model order for different HMMs should not be assumed.

3.3.2 Local Analyses

Proceeding with eight states, local analyses were conducted for all HMM subtypes. Local analyses refer to procuring and interpretating state patterns. In this section we first discuss activation states and compare state patterns directly outputted from AB HMM with

those from averaging BOLD signal in accordance with all subtypes' Viterbi paths. Then we discuss connectivity states attained directly from SFC and FFC HMMs and evaluate them in relation to ones obtained from alternative methods of acquisition. States from each subtype are distinguished with underscores corresponding to the HMM they stem from, i.e., $S1_{AB}$ corresponds to state one from AB HMM.

3.3.2.1 Activation States

The activation patterns for AB HMM are seen in Fig. 12a. $S1_{AB}$ appears to be a DMN-dominant state since the DMN is activated, and all other networks are deactivated. $S2_{AB}$ shows both DMN and FPCN to be activated. $S3_{AB}$ and $S4_{AB}$ are both attention-dominant networks: $S3_{AB}$ shows DAN and SN to be activated while $S4_{AB}$ shows FPCN, DAN, and SN to be activated. $S5_{AB}$ and $S7_{AB}$ both show all networks to be activated, but with $S5_{AB}$ having activation levels slightly lower in magnitude. $S6_{AB}$ shows DMN, DAN, and SN to be deactivated and FPCN to have minor positive activation levels. $S8_{AB}$ shows

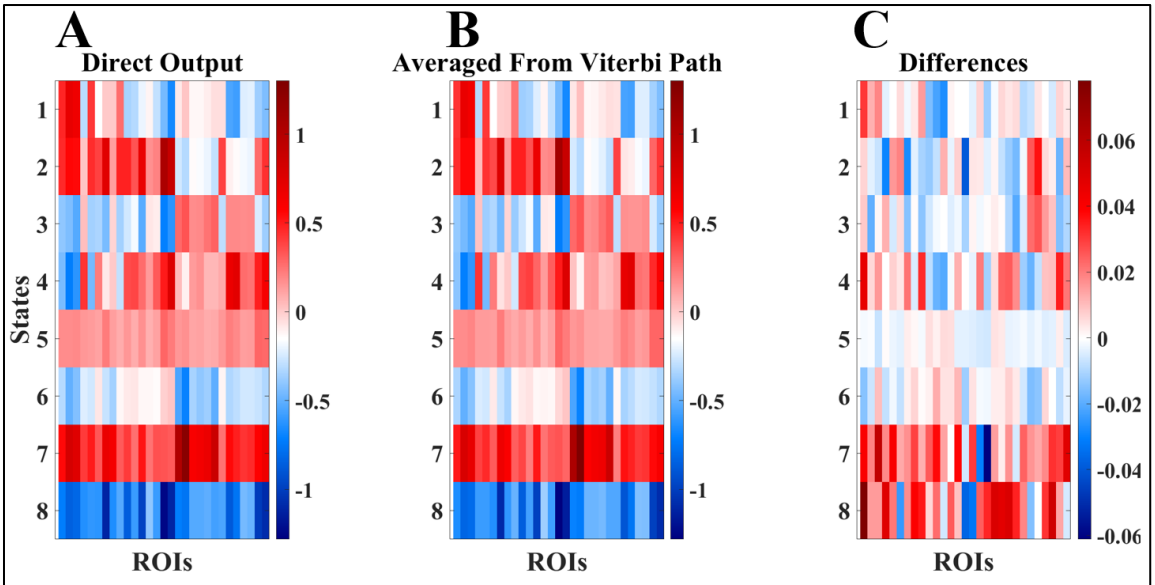

Figure 12: AB HMM state patterns **(A)** directly outputted from the model, **(B)** after averaging BOLD from TRs labeled by the Viterbi path using the process seen in Fig. 6, **(C)** and the difference in magnitude between them.

all networks to be deactivated, perhaps because some other network we are not investigating (i.e., sensorimotor or visual networks) is activated.

Activation patterns acquired using the method outlined in Fig. 6 are seen in Fig. 12b and their difference from AB HMM's direct output are seen in Fig. 12c. Great consistency exists between the activation patterns obtained from two different techniques with little to no differences between them ($R^2 = 0.992$, $p < 0.01$). Either method is acceptable in obtaining activation state patterns when a dataset has a short TR and long duration likely making it sensitive to fast and transient changes and allowing it to capture these dynamics.

Using the methods outlined in Fig. 7, the activation patterns acquired by averaging the BOLD signal from TRs where SFC and FFC's Viterbi paths labeled a state to be active can be seen in Fig. 13b and Fig. 13e respectively and qualitative differences from AB HMM's mean activation output (Fig. 13a and Fig. 13d) can be observed. The SFC and FFC connectivity-based activation plots show very different spatial patterns with much smaller magnitudes compared to the AB HMM direct output. There are even qualitative differences between the SFC- and FFC-based activation patterns. Euclidean distances between each SFC spatial pattern in Fig. 13b and each AB spatial pattern in Fig. 13a are seen in Fig. 13c, and likewise for FFC in Figs. 13d-f. The connectivity-based states had no relation whatsoever to the activation-based states (i.e., $S1_{AB}$ is not expected to have any affiliation to $S1_{SFC}$ or $S1_{FFC}$) motivating the Euclidean distance to be computed between all state pairings. Moreover, the Euclidean distance was computed because Pearson correlations are susceptible to parallelisms and fluctuations in patterns, so, in this case,

using a distance metric was more reasonable in identifying similarities or dissimilarities between vectors. The Euclidean distances show that there was a great disparity between the AB activation states and those derived from the SFC and FFC Viterbi paths. There were two instances in which the Euclidean distance was found to be quite small, indicating similarities between states: between $S3_{AB}$ and the fourth activation pattern acquired from FFC's Viterbi path, and between $S1_{AB}$ and the fourth activation pattern acquired from FFC's Viterbi path. This likely occurred because the distance metric recognized the same ROIs in both types of patterns (ROIs 13 – 16) to have relatively lower activation compared to others (ROIs 1 to 3). Nevertheless, these are only two instances and therefore is not enough to dispute the idea that AB HMM generates activation states that are spatially different from those acquired by averaging according to connectivity-based Viterbi paths (SFC and FFC HMMs).

These results indicate that AB HMM recognized only fluctuations in BOLD signal, and that SFC and FFC HMMs were not sensitive to those same patterns. Each subtype adequately recognized changes only in the data inputted into their respective model type and therefore did not all recognize the same state patterns. Thus, each HMM subtype is distinct in identifying changes in the inputted data, and consequently, in identifying states. AB HMM can be employed to identify activation spatial patterns without fear of those patterns being corrupted by potential correlations in the BOLD signal between ROIs.

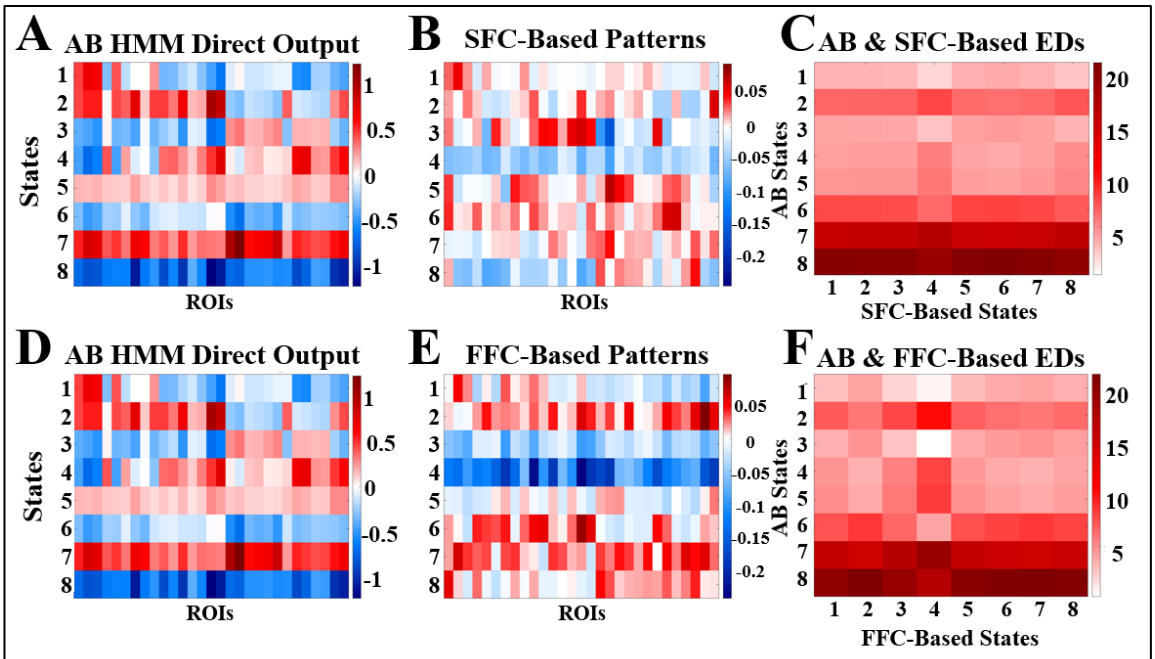

Figure 13: Top Row: Activation patterns **(A)** directly outputted from AB HMM and **(B)** from averaging BOLD signal according to SFC HMM Viterbi path, and **(C)** the Euclidean distances (EDs) between all pairs of these states. Bottom Row: Activation patterns **(D)** directly outputted from AB HMM and **(E)** from averaging BOLD signal according to FFC HMM Viterbi path, and **(F)** the EDs between all pairs of these states. The connectivity-based spatial patterns were acquired using the procedure in Fig. 7.

3.3.2.2 Connectivity States

All analyses were performed on highlights from the SFC and FFC HMM connectivity states. Raw connectivity states can be seen in the top row of Fig. 14a for SFC and 14b for FFC HMMs, and the grand mean excluding a state was calculated for $S1_{SFC}$-$S8_{SFC}$ and for $S1_{FFC}$-$S8_{FFC}$ (middle for of Fig. 14a and 14b, respectively). More specifically, SFC and FFC HMM state highlights were computed using Eq. 7 where X_i is the original raw functional connectivity matrix for state *i*, H_i gives the highlights of X_i, and *j* gives all state assignments for a determined model order excluding the value of *i* (13).

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

The resulting state highlights (H_i) are seen in the last row of Fig. 14a for SFC HMM and of Fig. 14b for SFC HMM. Because the repetition time (TR) was 0.72s, a window length of 50 time points (36 seconds) was selected for subsequent analyses (i.e., Global Analyses) to encompass less than 60s of data as recommended by Lurie et al. 2020 (14). We expand beyond 50 time points later in this section for a comprehensive analysis regarding how window size impacts connectivity state patterns. None of the state ordering amongst the connectivity states were paired, nor did they share any relation to the activation state assignments. The only exception is the covariance-based connectivity states which stemmed from AB HMM and therefore were aligned with the activation states. It is also important to note that because we were examining state highlights 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 highlights shown, negative values are associated with below baseline correlations, not anticorrelations.

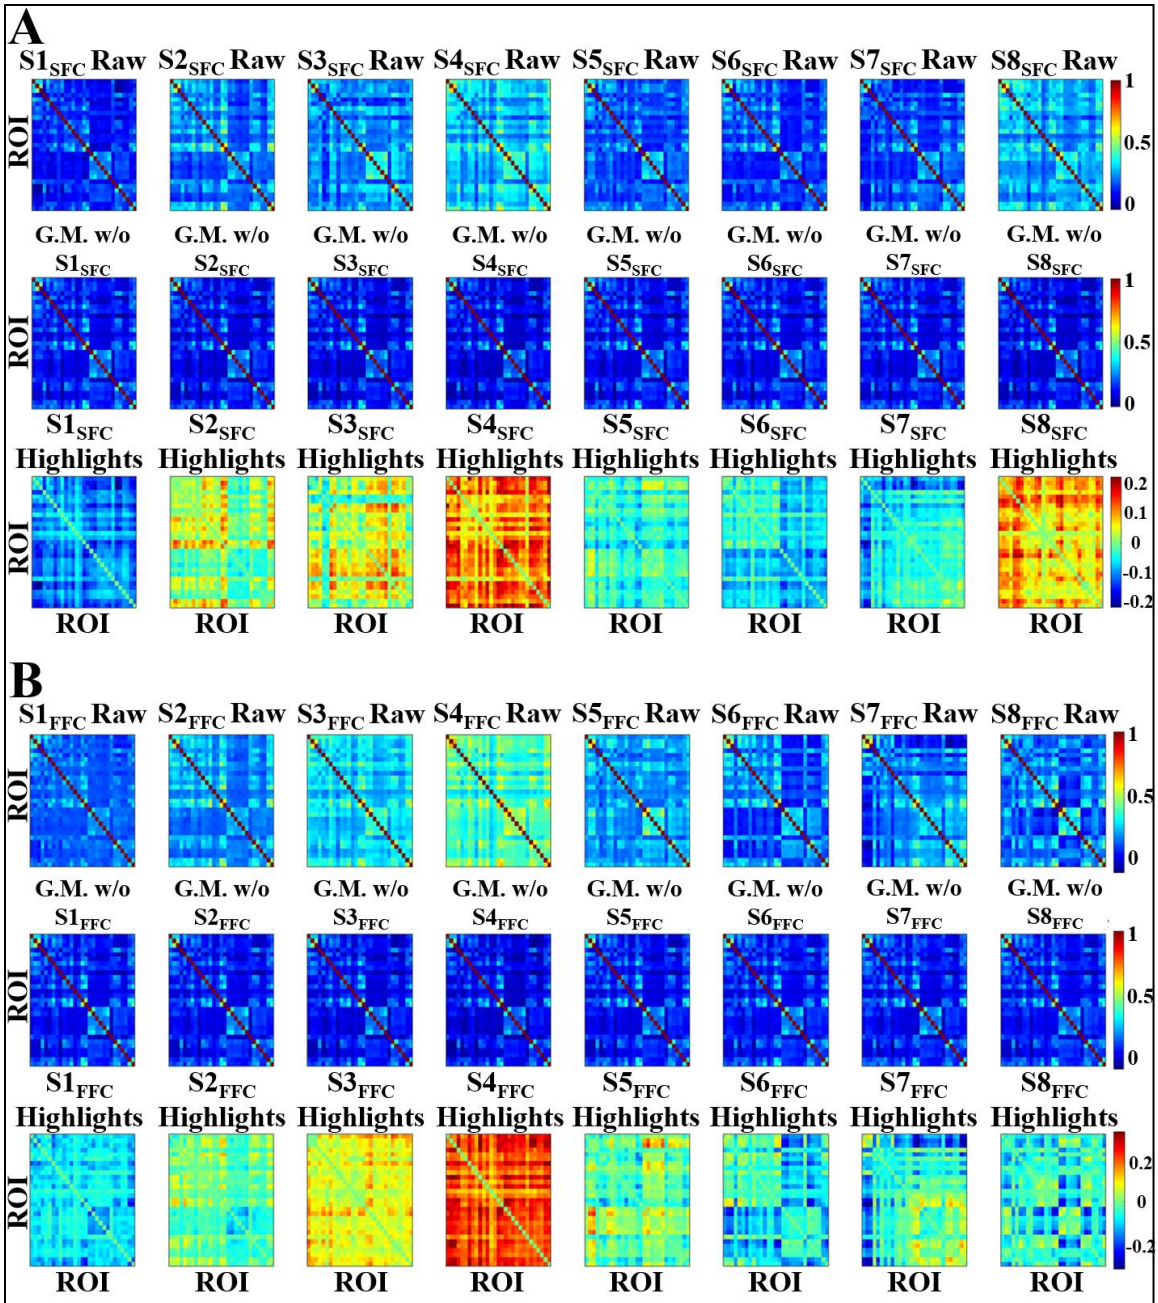

Figure 14: Method of acquiring **(A)** SFC HMM and **(B)** FFC HMM state highlights. Top Row: Raw connectivity states for states 1 through 8. Middle Row: Grand mean (G.M.) for all connectivity states excluding the state indicated. Bottom Row: The information in the middle row is subtracted from that in the top row to obtain the connectivity state highlights observed. The color bars shown correspond to the color axis for each given row.

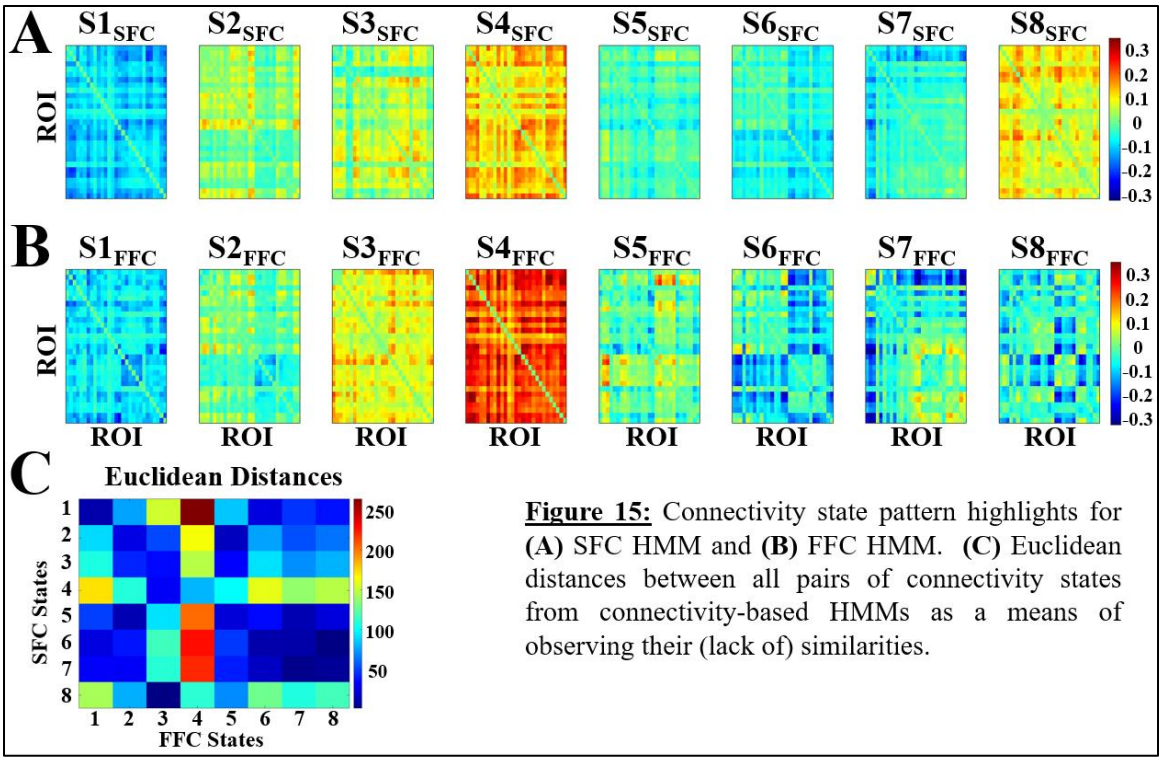

Figure 15: Connectivity state pattern highlights for **(A)** SFC HMM and **(B)** FFC HMM. **(C)** Euclidean distances between all pairs of connectivity states from connectivity-based HMMs as a means of observing their (lack of) similarities.

For SFC HMM state highlights (Fig. 15a), $S1_{SFC}$ showed below baseline correlations amongst all ROIs from all networks, particularly within and between DAN and SN, indicating a disconnect within the attentional system. $S2_{SFC}$ showed slightly elevated correlations within and between all networks expect DAN; DAN appears not to have much interaction with the rest of the system within this state. There were also slightly higher than average correlation values within $S3_{SFC}$ with the exception of a few ROIs (i.e., 5 and 6) in DMN showing below average connectivity with the rest of the system. $S4_{SFC}$ showed the strongest connectivity within and between the attentional networks indicating that the most communication within the system likely occurs in this state. $S4_{SFC}$ also exhibited the opposite behavior to that of $S1_{SFC}$: $S1_{SFC}$ showed the lowest correlation amongst ROIs while $S4_{SFC}$ showed the greatest. $S5_{SFC}$ appeared to be the converse of $S2_{SFC}$ in that there was more of a disconnect amongst DMN, FPCN, and SN. DAN also did not

show distinctive behavior in this state. In $S6_{SFC}$, below baseline connectivity was seen within DAN and between DAN and the other networks. Similarly, SN exhibited the same behavior. These properties indicated a slight disconnect within and between these attentional networks, and no distinctive properties amongst DMN and/or FPCN. $S7_{SFC}$ showed a slight disconnect between DMN and all other networks indicating a disconnect in the resting state networks. $S8_{SFC}$ exhibited the opposite behavior to $S7_{SFC}$ and actually showed elevated connectivity within DMN and between DMN and all other networks. This indicated more communication between resting state and attention networks. Elevated correlations existed amongst the other networks, but the most distinguishing behavior of this state lies with DMN.

For FFC HMM state highlights (Fig. 15b), $S1_{FFC}$ showed no distinguishable state highlights since all connectivity values were around baseline while $S2_{FFC}$ showed slightly higher than baseline connectivity within and between all networks except for DAN. $S3_{FFC}$ showed slightly above average correlations amongst all ROIs from all networks. $S4_{FFC}$ showed all networks to have above average correlations with one another and $S5_{FFC}$ to have slightly above baseline correlations between DAN and all other networks. $S6_{FFC}$ showed below baseline correlations between DAN and all other networks, and SN and all other networks. $S7_{FFC}$ showed DMN to be disconnected from all other networks and SN and DAN to have slightly above average correlation within and between each other. Finally, $S8_{FFC}$ showed below baseline correlations between DAN and all other networks.

Fig. 15c shows a distance matrix where the Euclidean distance was found between all pairs of SFC and FFC state highlights. Vertical state assignments correspond to FFC

states and horizontal state assignments correspond to SFC states. SFC and FFC HMMs could be expected to show some qualitative similarities in their state patterns because they stemmed from the same dynamic functional connectivity sliding time window analysis. Differences between the state patterns arose because SFC and FFC had different model inputs, or perhaps because FFC HMM was fit to 14 times as many values as SFC HMM. Fig. 15c shows a few smaller Euclidean distances indicating that some states across the connectivity-based HMMs exhibited similarities. However, there was no obvious case where a unique correspondence in spatial pattern between SFC and FFC HMMs were observed. A one-to-one match in states would be illustrated with one small Euclidean distance (one blue square) and seven large Euclidean distances (seven red squares) within each row. Because this phenomenon was not observed, the two HMMs outputted states with no exclusive compatibility. State patterns from both model types were quantitatively dissimilar and thus SFC and FFC generated different state patterns; choosing to pursue one method over the other will produce varying results so understanding the underlying meaning of each is critical. Implementing SFC HMM is preferred if the purpose of a study is to examine the general connectedness of an ROI to all other ROIs. If the investigation focuses on specific correlations between all nodes of interest, then choosing FFC HMM is satisfactory. The two connectivity-based HMMs will not produce the same state patterns so model type choice is crucial. Furthermore, it is not always the case that the same model order will be designated across the two connectivity-based HMMs. Varying model orders for SFC and FFC HMMs may produce an even greater discrepancies between their spatial patterns so it is not always assumed that their states will or should be similar.

Fig. 16a shows the connectivity state patterns acquired from averaging the dFC matrices from every time window where FFC HMM Viterbi path labeled a state to be active, and matched closely (Pearson correlation of $R^2 \geq$ ~0.9) with those outputted from model. As with AB HMM, when employing FFC HMM, state spatial patterns can be acquired either straight from the model or by averaging specific time points so long as the dataset is fast and large enough. Because the states converged to the same spatial patterns using two techniques, these results confirm the notion that the connectivity-based model identified changes only in connectivity despite having a plethora of components to fit.

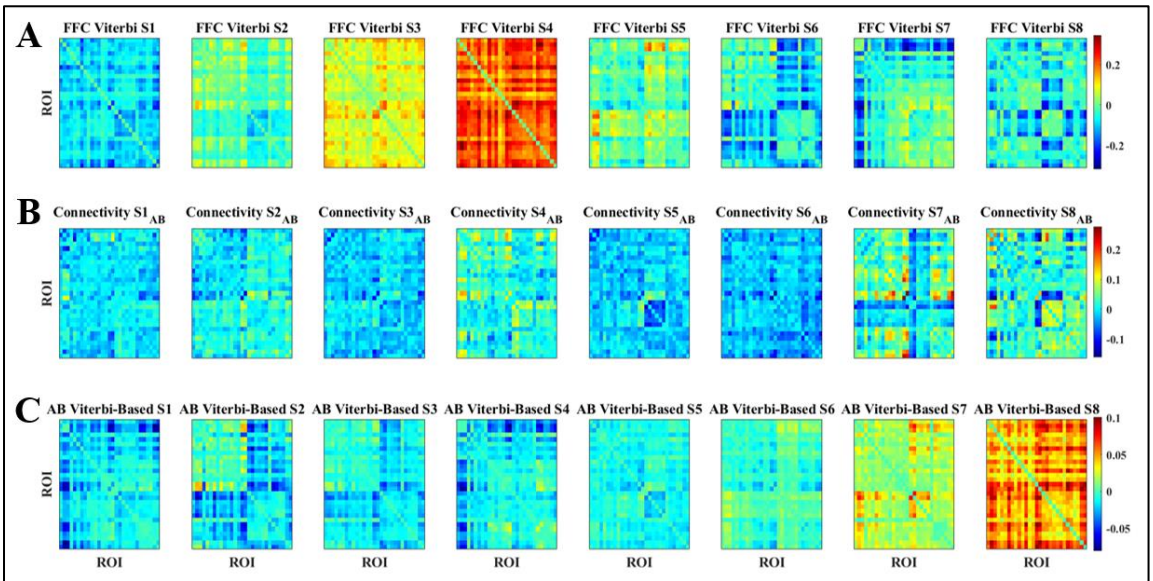

Figure 16: **(A)** Connectivity states acquired by averaging connectivity matrices where the FFC HMM Viterbi path labeled a state to be active. **(B)** Connectivity states acquired from AB HMM where the covariance matrices outputted were convert to Pearson correlations. The state patterns seen in **A** and **B** do not qualitatively match SFC or FFC HMM's state patterns (Fig. 14) and there are not replications of the same states across different model subtypes. We quantitively explore this in Fig. 17. **(C)** Connectivity states acquired by averaging connectivity matrices where the AB HMM Viterbi path labeled an activation state to be active as outlined in Fig. 8. The color bars correspond to the color axes within each given row.

The covariance-based connectivity states (Fig. 16b) do not display distinguishing patterns either within or between the states despite the fact that the state highlights are shown. DAN in $S7_{AB}$ appears to be disconnected from all other networks and is associated with activated DMN, FPCN, DAN, and SN. A few ROIs in within DMN and DAN appear

to have above average connectivity in $S4_{AB}$ and is linked to an attention state where FPCN and DAN show the highest levels of activation. $S8_{AB}$ shows a few ROIs between FPCN and DAN to be more disconnected than in the other states and DAN to have slightly elevated within-network connectivity and is paired with all four networks' deactivation. While these states show considerable qualitative differences from SFC and FFC states, they also show quantitative differences. Fig. 17 displays distance matrices where the Euclidean distance was computed between the AB covariance-based connectivity matrices (horizontal state assignments) and the SFC and FFC connectivity states (vertical state assignments). Because the Euclidean distances showed no state to be identically matched, these states were neither analogous to SFC states nor analogous to FFC states. The HMMs were recognizing differences in connectivity when R^2 (or R^2-related) correlations were fitted but were not able to capture discernable covariance patterns from the z-scored BOLD. Thus, SFC and FFC HMMs provided more information into diverse connectivity

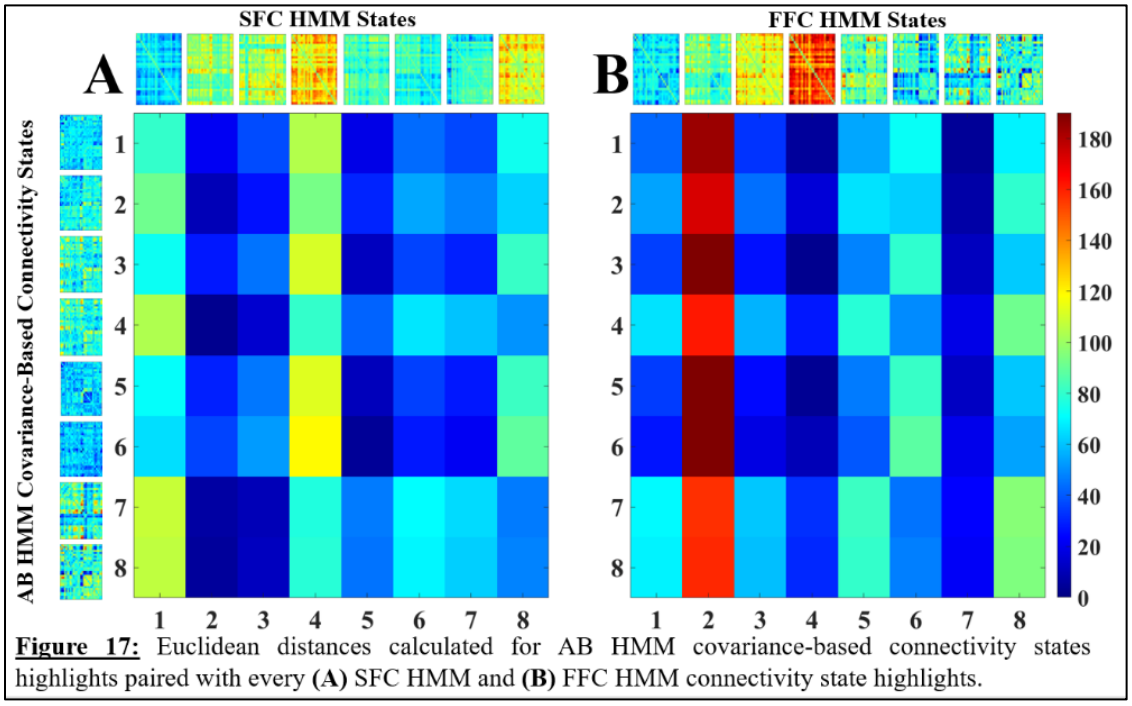

Figure 17: Euclidean distances calculated for AB HMM covariance-based connectivity states highlights paired with every **(A)** SFC HMM and **(B)** FFC HMM connectivity state highlights.

patterns than AB HMM and are best used when a study involves identifying latent connectivity states. However, if an investigation's interest lies in determining connectivity states in conjunction with activation states, then AB HMM is required because it is the only subtype that can link the two state types.

Fig. 16c shows the connectivity state patterns highlights acquired via the method illustrated in Fig. 8, i.e., where the ROI x ROI matrices resulting from the sliding window correlation analysis were averaged wherever the AB HMM Viterbi path labeled an activation state to be active. Although some of these spatial patterns may qualitatively look like those from SFC and/or FFC HMMs, quantitative analyses actually showed them to be dissimilar. The distance matrix in Fig. 18 display Euclidean distances between AB Viterbi-based connectivity states (horizontal state assignments) and the SFC and FFC states (vertical state assignments). Because all values in this plot are consistently high and showed no obvious cases of singularly matched states, SFC and FFC states reflected

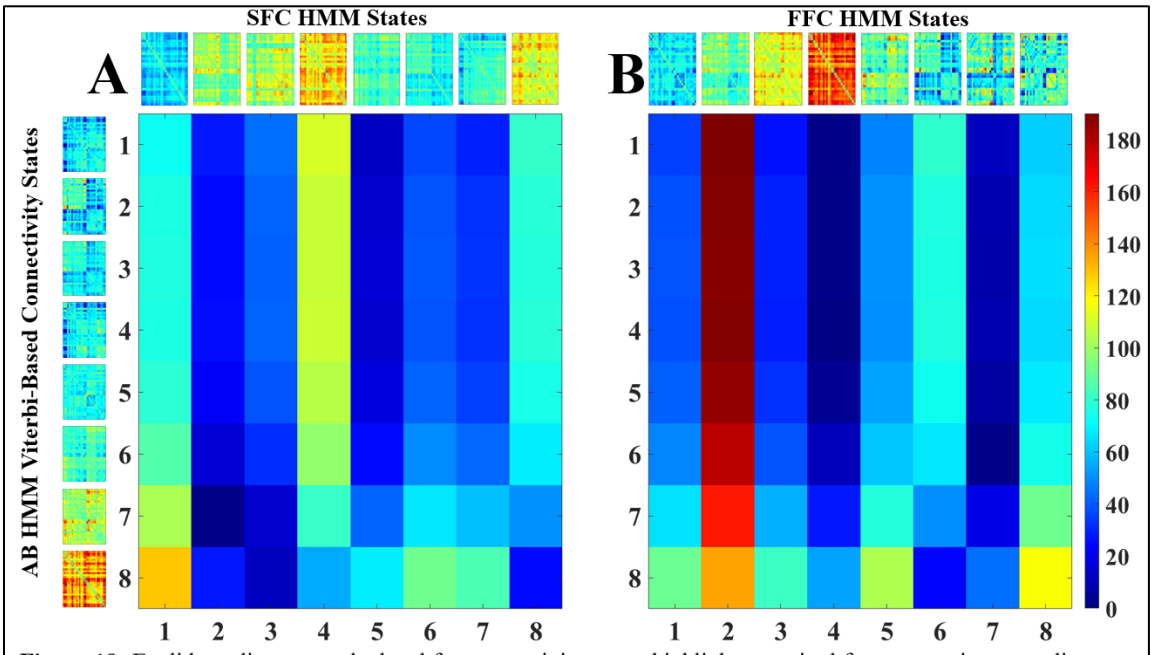

Figure 18: Euclidean distances calculated for connectivity states highlights acquired from averaging according to AB HMM Viterbi path paired with every **(A)** SFC HMM and **(B)** FFC HMM connectivity state highlight.

changes only in connectivity patterns and were robust to fluctuations in BOLD signal. Thus, AB HMM suitably detected activation states while SFC and FFC HMMs satisfactorily ascertained connectivity states and the corresponding changes between them. Each model type was distinct, and the information identified in one was not muddled by the information inputted into others.

To investigate how these patterns changed as a function of temporal resolution, state connectivity spatial patterns were acquired when using different window lengths in the sliding time window analysis, some above 50 time points (tp), and some less. Specifically, SFC and FFC connectivity patterns were found when window sizes of 30 time points (21.6 seconds), 40 time points (28.8 seconds), 50 time points (36 seconds), and 60

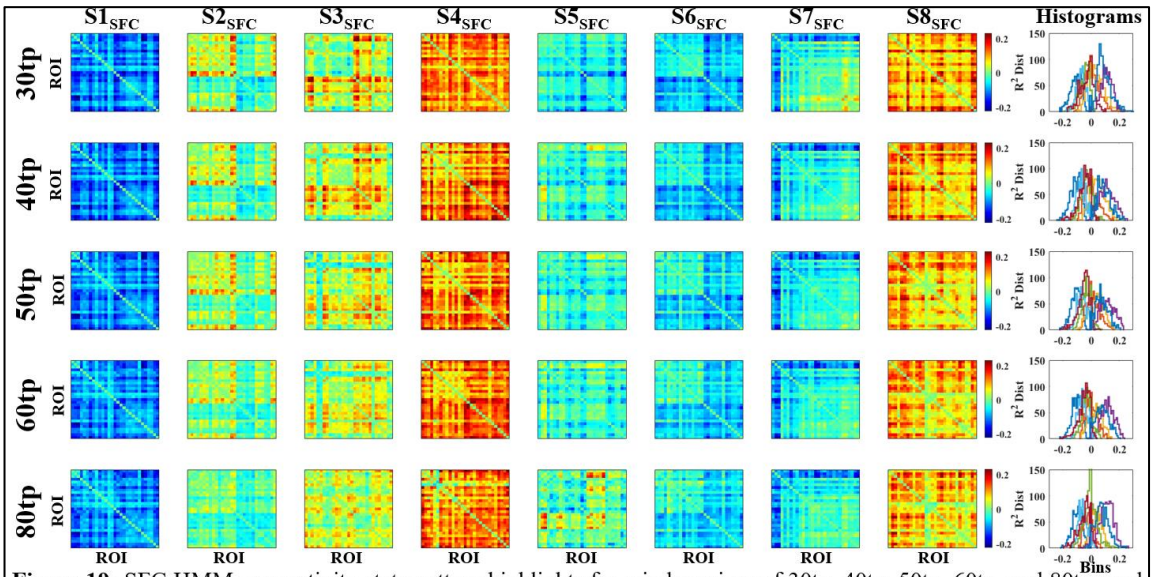

Figure 19: SFC HMM connectivity state pattern highlights for window sizes of 30tp, 40tp, 50tp, 60tp, and 80tp used in the sliding window approach. The histograms display the overall distribution of R^2 values of $S1_{SFC}$-$S8_{SFC}$ for each given window size. The color bars correspond to the color axes within each given row.

time points (43.2 seconds), and 80 time points (57.6 seconds) were employed. Fig. 19 shows the patterns for SFC HMM while Fig. 20 shows those for FFC HMM. The

histograms at the end of each row illustrate the overall distribution of R^2 (connectivity) values in each state for every window size.

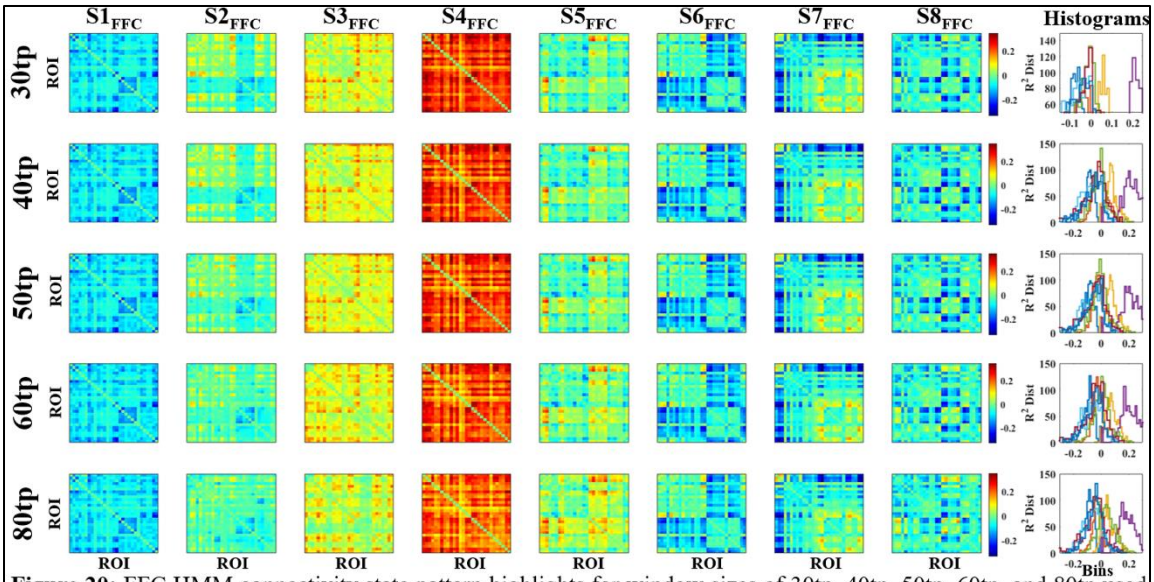

Figure 20: FFC HMM connectivity state pattern highlights for window sizes of 30tp, 40tp, 50tp, 60tp, and 80tp used in the sliding window approach. The histograms display the overall distribution of R^2 values of $S1_{FFC}$-$S8_{FFC}$ for each given window size. The color bars correspond to the color axes within each given row.

SFC state highlights showed little to no variability amongst the different window lengths. All state highlights matched their counterparts across window sizes with Pearson $R^2 \geq 0.5$ except when correlating $S4_{SFC}$ 30 time points (tp) to $S4_{SFC}$ 80tp ($R^2 = 0.3084$), and when correlating $S3_{SFC}$ of any time length to 80tp (R^2 =-0.3587, -0.2588, -0.3062, -0.2563). Likewise, FFC state highlights showed minimal differences across window lengths. All states matched to their analogues with $R^2 \geq$ ~0.5 with 1 exception: matching $S3_{FFC}$ 40tp to $S3_{FFC}$ 80tp ($R^2 = 0.3469$). SFC states potentially showed more variability amongst different time windows than FFC states because of the summing factor. Notably, the poor pattern matchings only occurred when correlating state highlights to the 80tp indicating that this window size (~60 seconds) was too large and that spatial patterns began to diverge from their norm. The histograms (far right columns, Figs. 19 and 20) support this idea

because they became less separated as the window size increased meaning that the state patterns' resemblance to their counterparts from smaller window sizes was reduced. This was particularly noticeable in the 80tp window size where the curves were closest together. These results showed that a window size containing less than approximately 60 seconds of data was preferred for connectivity-based HMMs, consistent with the ideas Lurie et al. 2020 presented (14).

Overall, the connectivity state patterns appeared to be relatively robust to window size indicating that there was some leniency in determining this parameter in sliding window analyses performed on big datasets with a fast TR and many subjects. For an acceptable length, window size is not a limiting factor when deciding whether to pursue SFC or FFC HMM for a particular investigation. We were confident in our choice of 50 time points (~36 seconds) because it was consistent with the suggestion to aim for a window size containing less than 60 seconds of BOLD data, so all global analyses discussed in the next section stemmed from a sliding window approach using 50tp (14).

3.3.3 Global Analyses

Global analyses refer to outputs or computations common to all three model types where methods of performing those analyses did not have to be modified for a specific subtype. We compared the Viterbi path, transition probability matrix, switching rate, and fractional occupancy correlation across AB HMM, SFC HMM, and FFC HMM.

3.3.3.1 Viterbi Path

The Viterbi paths for AB HMM (Fig. 21a), SFC HMM (Fig. 21b), and FFC HMM (Fig. 21c) showed which state was active at every time point for every subject. Each color

represents a different state as illustrated by the color bar. Although the stability analyses appointed eight states to all three subtypes, state assignment across all HMMs were

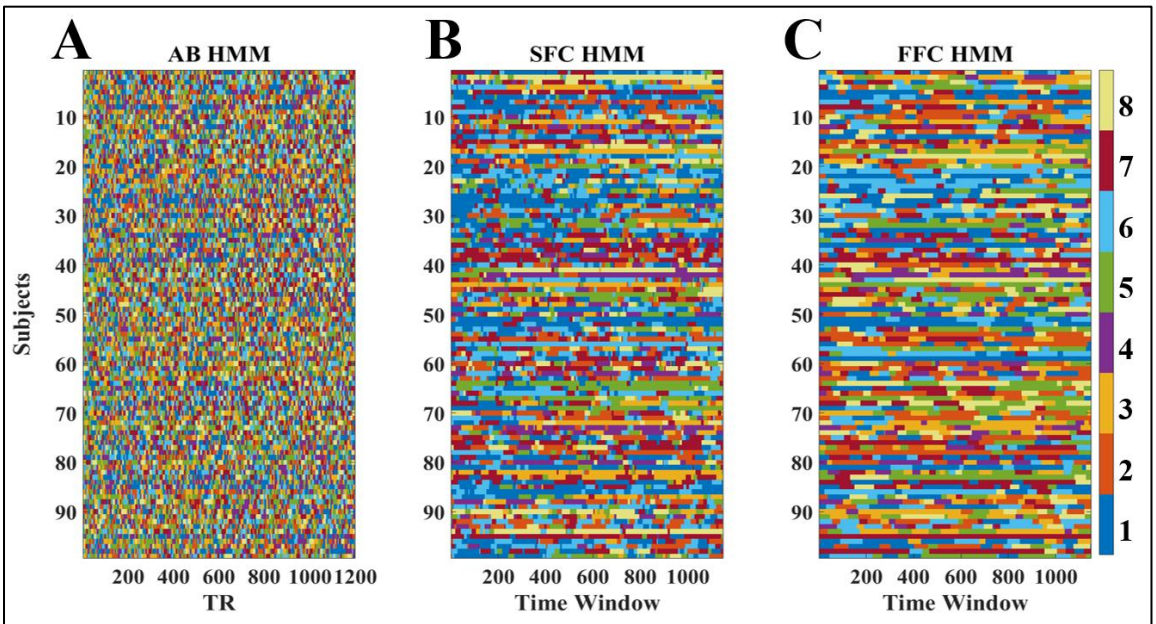

Figure 21: Viterbi Paths for **(A)** AB HMM, **(B)** SFC HMM, and **(C)** FFC HMM.

arbitrary. That is, $S1_{SFC}$ was unique to SFC HMM and had no relation to $S1_{AB}$ or $S1_{FFC}$. These plots show that, compared to AB HMM, the connectivity-based HMMs were smoothed over time. Temporal discrepancy between the activation and connectivity HMMs potentially contributed to this in that SFC's and FFC's x-axes employ time windows rather than TRs, and therefore possess Δt (50) fewer time points. Autocorrelation was another contribution since the time window shifted only one time point at a time in the sliding window analysis. Although the temporal smearing is blatantly obvious, the Viterbi path was reproducible across different initializations ($R^2 \geq \sim 0.84$). For all realizations, the models recognized the same connectivity states to be prevalent during the same time windows and the same switches between states even with the temporal smearing. Neuroscientific conclusions could not be made directly from these plots, and instead

additional analyses, such as switching rate and fractional occupancy correlation, were needed. It is possible that qualitative differences in this measure can be observed across different conditions.

The FFC HMM Viterbi path showed increased temporal smoothing compared to that of SFC HMM and therefore contained fewer transitions between states. Although we do not know exactly why this happened, we suspect that FFC HMM fitting so many more components than SFC HMM could have played a role. That is, compared to SFC HMM, FFC HMM fitted a more holistic picture of the information in a sliding window analysis and therefore from fitting a profuse amount of connectivity values per time window (406). It nevertheless was stable because the Viterbi paths converged to the same solution for every realization indicating that 406 components per time window was not an irrational amount of data to fit. It was, however, a much larger amount than SFC's 29 components per time window which may have caused FFC HMM to take longer in recognizing switches from one state to another: FFC HMM worked 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 model order as SFC HMM. It is possible that because FFC HMM had more variables that it could identify more states and therefore more possible spatial patterns. As previously mentioned, eight states were chosen for a comprehensive comparison with SFC HMM. However, we did not comprehensively dissect these ideas, so they remain hypotheses and require additional analyses.

Overall, AB HMM provided the most insight into fMRI temporal dynamics, so if an investigation focuses on time-related measures (i.e., switching), then AB HMM is the logical choice. Furthermore, caution must be taken when fitting connectivity-related values because temporal smoothing is inherent to a windowed analysis and consequently reduced state transitions in the hidden state sequence which can affect subsequent analyses. Caution is also needed when deciding how many specific nodal connections FFC HMM should fit which may require preliminary tests.

3.3.3.2 Transition Probability Matrix

Like the Viterbi paths, transition probability matrices were directly outputted from each model subtype (Fig. 22a) and were of great interest because they quantitatively described the likelihood of transitioning out of State X and into State Y. As with the Viterbi path, neuroscientific interpretations of the transition probabilities could not be directly compared across all HMMs because the states were not the same. However, this measure could provide quantitative insight into the relationship between states within each model type, something that could not be observed from the Viterbi path alone. The probability of remaining in the same state overpowered transitions between different states, so the diagonals have been removed for better visualization of the off-diagonal elements.

The AB HMM transition probability matrix (Fig. 22a) showed notably high transitions from $S3_{AB}$ into $S5_{AB}$, from $S6_{AB}$ into $S5_{AB}$, and from $S8_{AB}$ into $S3_{AB}$ and $S6_{AB}$. The SFC HMM transition probability matrix (Fig. 22a) showed notably high transitions between $S1_{SFC}$ and $S6_{SFC}$, from $S2_{SFC}$ into $S6_{SFC}$, and from $S7_{SFC}$ into $S1_{SFC}$. There were no transitions between $S1_{SFC}$, $S2_{SFC}$, $S3_{SFC}$, $S4_{SFC}$, and $S8_{SFC}$. Likewise, no subjects

transitioned between $S6_{SFC}$ into $S4_{SFC}$. The FFM HMM transition probability matrix (Fig. 22a) showed notably high transitions from $S2_{FFC}$ into $S6_{FFC}$, from $S3_{FFC}$ into $S8_{FFC}$, from $S6_{FFC}$ into $S8_{FFC}$, and from $S6_{FFC}$ into $S7_{FFC}$. There were no transitions from $S2_{FFC}$ into $S8_{FFC}$, from $S5_{FFC}$ into $S2_{FFC}$, or from $S8_{FFC}$ into $S2_{FFC}$. The plots reflect the temporal resolution discrepancy between the connectivity-based HMMs and AB HMM because the magnitude of transitions in SFC and FFC HMMs were smaller compared to those of AB HMM. FFC HMM's transition probabilities had smaller magnitudes than those of SFC HMM thereby reflecting its exaggerated temporal smoothing. To compare overall behavior across model subtype, the relative values rather than the actual magnitudes should be examined. Neuroscientific interpretations about the actual transition probabilities magnitudes can be examined within a model type and between conditions for a dataset.

Along with the autocorrelation and smoothing elements, SFC HMM's off-diagonal transition probabilities may have exhibited transition probabilities close to zero because this model type used summed node-wise correlations as input data. Because some R^2 values within a time window might have been positive and others might have been negative, summing over them averaged out connectivity behavior to zero, which likely exacerbated the temporal smoothing effect and hindered the HMM from recognizing potential connectivity dynamics. FFC HMM's off-diagonal transition probabilities might have had even smaller values because this model type fit many components at once, so it was quite computationally demanding to identify transitions between 406 different components per time window. With these ideas in mind, each transition probability matrix nevertheless provided useful insight into their respective model subtype's dynamics.

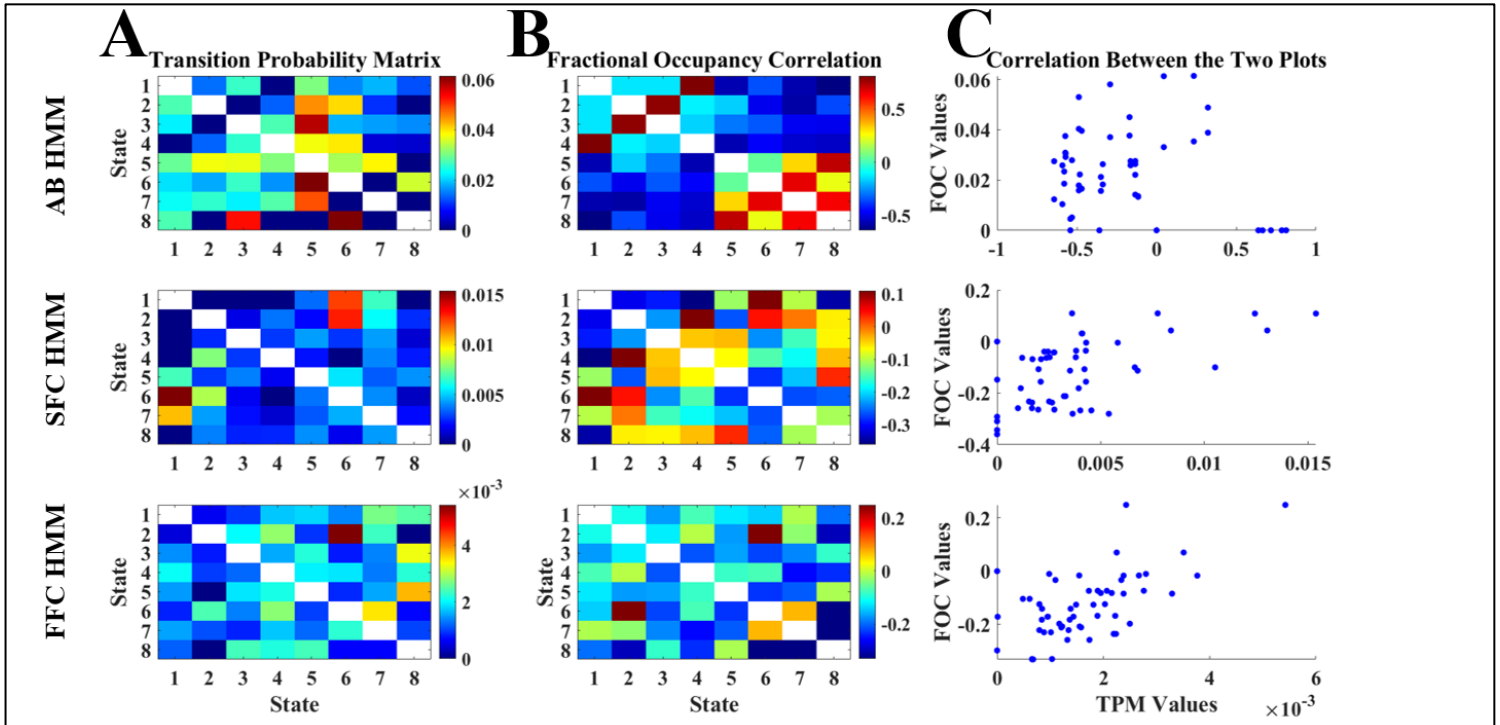

Figure 22: **(A)** Transition Probability Matrix (TPM), **(B)** Fractional Occupancy Correlation (FOC), and **(C)** correlation between these two measures of interest for AB HMM (top row), SFC HMM (middle row), and FFC HMM (bottom row). The values on the diagonal have been removed for easier visualization of the off-diagonal elements.

3.3.3.3 Switching Rate

We initially aimed to examine switching rate as a means of comparing whether one HMM subtype could better capture switching dynamics over the others. However, the temporal smoothing factor inherent to the connectivity-based models made this very difficult, and thus we could not compare this measure between model types. Nevertheless, switching rate can still be compared within a model type or within an investigation of fMRI data collected between different conditions. Because the fMRI dataset examined in this chapter contained only pure resting state only one condition occurred, this measure will be revisited in Chapter 4 where we focused on a dataset containing BOLD signal acquired under different conditions. Although it is possible to compute the switching rate within the connectivity-based HMMs, the results would be greatly smoothed across time and therefore would be overwhelmed with noise so the results might not show anything meaningful. As a result, this measure is best explored when an AB HMM is employed due to its superiority in interpreting temporal dynamics of an fMRI dataset.

3.3.3.4 Fractional Occupancy Correlation

Studying the fractional occupancy correlation provided additional quantitative insight into the relationship between states, and, as Vidaurre et al. 2017 suggested, metastates may organically arise (12). The AB HMM fractional occupancy correlation plot (Fig. 22b) shows that this might be the case because two subgroups were observed: one containing the attention and resting state dominant networks ($S1_{AB}$-$S4_{AB}$) and the other containing the states where either all networks were activated or deactivated ($S5_{AB}$-$S8_{AB}$).

This plot shows that subjects tended to reside in and transition within each subgroup (either in $S1_{AB}$-$S4_{AB}$ or in $S5_{AB}$-$S8_{AB}$) more often than between them.

Fig. 22b shows that a potential metastate may have occurred comprised of $S1_{SFC}$, $S2_{FFC}$, $S4_{SFC}$, and $S6_{SFC}$. The high fractional occupancy correlation values exhibited between these particular states (the dark red blocks in the plot) may serve as the diagonal for a metastate, similar to the ones seen in the AB HMM counterpart. The FFC HMM fractional occupancy plots (Fig. 22b) demonstrated a relatively strong correlation in residence between $S2_{FFC}$ and $S6_{FFC}$ and between $S6_{FFC}$ and $S7_{FFC}$, and a relatively strong anticorrelation in residence among $S8_{FFC}$ and $S6_{FFC}$ and $S7_{FFC}$. Subgroups were not immediately noticeable in this plot likely because the temporal smoothing was overpowering any dynamics in the data. Nonetheless, further investigation into confirming whether metastates arise from the connectivity-based HMMs is needed but is outside the scope of our investigation. We aimed to compare HMM subtypes and to propose which subtype to apply under certain circumstances, not to make neuroscientific observations. Overall, AB HMM adequately recognized fMRI temporal dynamics and therefore could naturally identify metastates while SFC HMM had the potential to detect them, but additional information was needed.

A potential relationship between fractional occupancy correlation and transition probabilities is seen in Fig. 22c. These analyses could provide additional insight into the interaction between states within a model type and could provide clearer neuroscientific interpretations of the acquired states and their corresponding behavior. However, they will not be interpreted in this dissertation because it is not the focus of this study. The plots

were included simply as a means of demonstrating possible supplementary analyses that could be carried out utilizing the information extracted from HMM subtypes.

3.4 Conclusion

Overall, our empirical analyses revealed that all three HMM subtypes provided different yet useful perspectives into an fMRI dataset, and model choice for a study depends on the question or concept investigated. Each HMM subtype was distinct and identified changes only in the data type (activation or connectivity) inputted. AB HMM identified changes only in activation levels, SFC identified changes only in the general connection between a node and all other nodes, and FFC identified changes only in specific connections between a reasonable number of nodes of interest. Spatial patterns for AB and FFC HMMs could be obtained either from the model's direct output or from averaging BOLD signal or connectivity matrices from their respective Viterbi paths. The latter technique was inherent to obtaining SFC states. Furthermore, AB HMM covariance-based connectivity states identified weak connectivity patterns, so SFC and FFC were more robust to establishing latent connectivity states. However, it cannot be assumed that both connectivity-based HMMs will produce the same spatial patterns. Therefore, model type choice in an investigation is important and depends on whether the goal is to examine general vs. specific nodal connections. If an investigation's interest lies in analyzing connectivity states within activation states, then AB HMM is favored and the covariance matrices must be converted into Pearson correlations.

For a large and fast dataset like the one investigated in this chapter, there was a bit more flexibility when choosing the window size for a sliding window analysis. Selecting

a size containing less than 60 seconds of data was optimal (consistent with Lurie et al. 2020) because the spatial patterns began to deviate from those observed when using longer windows (14). SFC and FFC HMMs' Viterbi paths were smoothed across time compared to AB HMM's, likely because they arose from a windowed analysis. Consequently, AB HMM is the ideal candidate for analyses that focus on the temporal dynamics. FFC's Viterbi path was more smoothed than SFC's likely because it was fit to 14 times as many components making it more difficult to identify changes in many nodal correlations at once and because its model order was constrained to match that of SFC HMM. Similarly, all transition probability matrices presented provided insight into their own model type's dynamics. Because of the largely varying scale of values, the relative changes in transition probabilities must be examined when comparing this measure across HMM subtypes. Differences in the magnitudes themselves can be analyzed within each model type.

Switching rate provided no meaningful interpretations between model types due to the substantial smoothing factor but could still be used to compare different conditions within a dataset and within a model type. Because this measure relies on the temporal resolution of the HMM, it is best analyzed within AB HMM. Metastates naturally arose from AB HMM fractional occupancy correlations likely because of its superior ability in capturing temporal dynamics. SFC could potentially identify metastates more clearly than FFC HMM, but additional information outside the scope of this investigation is needed. The theoretical and empirical perspectives outlined in this chapter facilitates future investigators to make informed decisions about which HMM subtype is best to pursue in their neuroimaging analyses and investigations.

Chapter 4: LC Dataset Findings for an Activation-Based Hidden Markov Model

4.1 Abstract

The locus coeruleus (LC) is a small subcortical structure situated at the top of the brainstem and is the principal source of norepinephrine in the brain; therefore, it is responsible for controlling cognitive functions related to arousal such as stress, the sleep-wake cycle, and attention (8,29–32). Degradation of this structure has been associated with aging and with neurodegenerative diseases such as Parkinson's and Alzheimer's diseases, so to understand the role that LC dysfunction plays in these conditions, its normal function must first be understood (33,34). Although an affiliation between LC and attention has been widely established, the corresponding underlying neural mechanisms driving this relationship are not well understood (29,31). In this chapter we employed a hidden Markov model (HMM) to derive latent brain states that centered around attention networks and fit it to a modified pseudo resting state paradigm where LC activity was noninvasively up-regulated via a handgrip task (11–13,36,37,67–71). Functional magnetic resonance imaging, pupillometry, and neuromelanin data were collected to understand how HMM-related measures varied as a function of changes in LC activity, and whether they held a relationship with LC structure. Specifically, we compared state transition probabilities, network variability, mean state duration, and state switching rates between an active squeeze and a sham session. Our findings suggest that observing a consistent relationship between LC engagement and latent brain state dynamics is unlikely when using the measures employed in this investigation. This could occur because our subject pool consisted of healthy young adults where LC structure is unlikely to be compromised, our

handgrip task may not have been strong enough to evoke a prominent LC response, or the measures employed in this investigation are not sensitive enough to detect it (72–74).

4.2 Introduction

The locus coeruleus (LC) (Fig. 23) circuit is the main source of norepinephrine in the brain; it projects to the entire brain and is deeply involved in cognitive functions related to arousal including attention, stress, and the sleep-wake cycle (7,8,29–32). In addition, LC dysfunction is hypothesized to occur in prodromal stages of Alzheimer's and Parkinson's diseases (33,34). LC degradation is also prevalent in normal aging since cell loss within this structure is thought to impair memory and cause cognitive reserve depreciation (75). For example, older adults with compromised LC structure have exhibited significant correlations between poor memory performance and thinner LC axons, while their younger counterparts demonstrated no relationship between memory performance and LC microstructure (9,10,75). Because acquired LC blood oxygen level dependent (BOLD) signal can be noisy due to LC's small stature, and because accurate measures of LC BOLD in animals have been invasively acquired, pupillometry data is typically collected as a noninvasive proxy measure for LC activity (76–78). Using

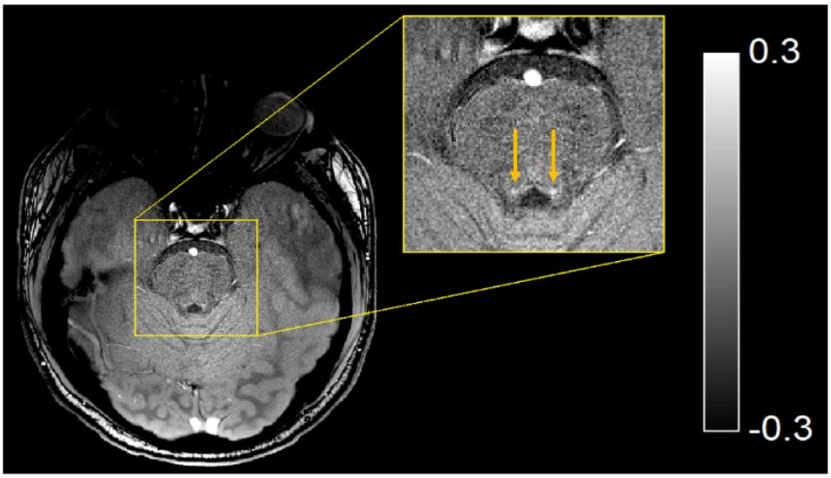

Figure 23: The locus coeruleus (LC) is a 2mm in diameter nucleus located in the brainstem as illustrated by the two dots on either side of the region anterior to the fourth ventricle (7,8).

pupil size as a noninvasive surrogate for LC activity is a method supported in literature reporting strong correlations between pupillometry data and LC BOLD activity (localized from both neuromelanin-sensitive structure imaging and a LC atlas published by Keren et al. 2009) in functional magnetic resonance imaging (fMRI) studies in humans (76,78). After regressing out physiological noise, it has been shown that continuous measures of pupil diameter throughout both resting state and task stimuli are able to index tonic variations in LC BOLD activity, and are less liable to trial-by-trial noise than pupil dilation locked into task-related events (76). Electrophysiological studies in monkeys have similarly shown a reliable relationship between LC activity and changes in pupil diameter due either to spontaneous fluctuations, or to external stimuli (77). Thus, there is strong reason to believe that pupillometry data can act as an accurate and robust surrogate for LC BOLD activity.

In normal cognition, the relationship between arousal and task performance during a focused attention task is explained by the Yerkes-Dodson curve (29). Moderate LC firing rates correspond to optimal task performance while low and high LC firing rates are associated with inadequate task performance because subjects are inattentive and distracted, respectively (29). Despite observed correlations between LC engagement and attention, the underlying mechanisms driving changes in network dynamics within this relationship are not well understood (29–31). Characterizing this dynamic affiliation provides foundational insight into LC by facilitating the development of computational models that represent the dynamic spatiotemporal influence of LC over arousal.

Chapters 2 and 3 demonstrated that hidden Markov models (HMMs) sufficiently quantify the spatiotemporal behavior of latent states. Typically, HMMs are used for weather prediction, in computational biology, and in finance, however, HMMs are becoming more popular in neuroimaging (26,43,44). Neuroimaging-based HMMs identify latent brain states which quantify network or nodal interaction as well as the probability of transitioning between those hidden states (11–13,27,28,46–50,52–54). As a result, an HMM can be applied to fMRI data to identify the spatiotemporal behavior of latent brain states as a function of arousal or LC up-regulation and characterize LC's dynamic underlying relationship with arousal. To oversee how HMM-derived state behavior changes as a function of arousal, we compared this behavior before and after LC activity had been up-regulated. LC activity can be noninvasively up-regulated via a handgrip task which has previously shown to induced sympathetic arousal and increases norepinephrine activity; therefore, LC activity should increase as a result of squeezing (36,37,67–71,79).

In this chapter we aimed to ascertain how LC activity up-regulation affected brain states derived from an HMM in effort to better understand LC's relationship with attention. We modified a squeezing task to create a pseudo-resting state paradigm that increases LC activity so we can analyze changes in brain state behavior as a function of LC up-regulation (37,80,81). Pupillometry data were collected to indirectly measure changes in LC activity as a function of squeezing as well as the relationship between pupil dilation and latent brain state behavior (76,77). Thus, it was used as a noninvasive proxy measure of quantifying LC's relationship with latent brain states. Neuromelanin data were collected to determine the relationship between latent brain states or pupil dilation with LC structure (7,9).

4.3 Methods

In this section we first explain the dataset used in this investigation which has been optimized to analyze subcortical structures such as the LC as well as the experimental paradigm employed to up-regulate LC activity. We also introduce four attention-related networks and describe the HMM that will be applied to the LC dataset. Then we describe how local analyses focusing on state pattern characterization will be performed followed by global analyses focusing on HMM outputs and subsequent calculations. We end this section by describing analyses performed on pupillometry data.

4.3.1 Datasets and Networks

4.3.1.1 Experimental Paradigm and fMRI Data

The experimental paradigm is illustrated in Fig. 24 and shows all subjects first underwent a five-minute pure resting state episode prior to any squeeze. Following this RS0 block, subjects underwent a 12.8-minute experiment where they alternated between resting state and bringing their dominant hand to their chest to squeeze a squeeze-ball at

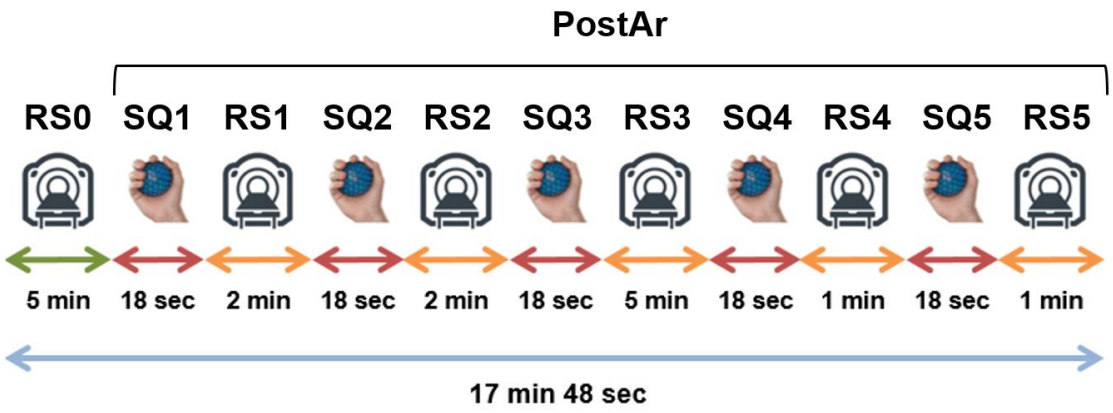

Figure 24: Experimental paradigm consisting of a pure resting state block and longer post-squeeze resting state blocks to allow for dynamic analysis (36). Subjects underwent both an active squeeze where they lifted their arm to their chest and squeezed a squeeze-ball at maximum grip strength for the length of the entire block, and a sham session where they refrained from squeezing. Blocks SQ1 through RS5 are referred to as post-arousal, or PostAr and block RS0 is baseline.

maximum grip strength (SQ1-RS5) (37,80). All five squeeze blocks lasted 18 seconds while the interspersed five resting state blocks had durations of two-, two-, five-, one-, and one-minute, respectively. Blocks SQ1 through RS5 occurred after the arousing stressor has been introduced and are thus deemed post-arousal (PostAr) blocks. In order to create a within-subject experimental design, all subjects underwent two sessions: one where they executed the squeeze (active session), and one where they still brought their arm up to their chest but were instructed simply to touch the ball and not to squeeze it (sham session). Subjects with an odd identification number squeezed during their first session and held the ball for their second session, while even numbered subjects experienced the opposite. Following this resting state paradigm, subjects underwent an auditory oddball detection task, the details of which are not relevant to this dissertation (82).

Thirty-one participants (18 females, mean age 25 years ± 4 years) enrolled in this study at the University of California, Riverside Center for Advanced Neuroimaging (81). All subjects gave written informed consent to participate, received monetary compensation for their participation, and the study's procedures were approved by the University of California, Riverside Institutional Review Board. Magnetic resonance imaging (MRI) data were called on a Siemens 3T Prisma MRI scanner (Prisma, Siemens Healthineers, Malvern, PA) with a 64 channel receive-only head coil. fMRI data were collected using a 2D echo planar imaging sequence (echo time (TE) = 32 ms, repetition time (TR) = 2000 ms, flip angle = 77°, and voxel size = 2x2x3mm^3, slices=52) while pupillometry data were collected concurrently with a TrackPixx system (VPixx, Montreal, Canada). Anatomic images from a MP-RAGE sequence (TE/TE/inversion time = 3.02/2600/800 ms, flip angle =8°, voxel

size = 0.8x0.8x0.8mm^3) were used for registration from subject space to common space. One subject was excluded due to a history of attention deficit hyperactive disorder and consumption of related medication resulting in all functional data to be analyzed with N = 30 (80,82). This subject was also excluded from all other types of data analyzed in this dissertation.

The functional data underwent a standard preprocessing pipeline in the functional magnetic resonance imaging of the brain software library (FSL): slice time correction, motion correction, susceptibility distortion correction, and spatial smoothing using a kernel gaussian smoothing factor of $\frac{Full\ Width\ Half\ Maximum}{2.3598} = \frac{2}{2.3598} = 0.8475$ (83,84). Finally, all data were transformed from individual subject space to Montreal Neurological Institute (MNI) standard space using the following procedure in FSL (83,84). First, the T_1-weighted image was skull stripped using the brain extraction tool. Next, brain extracted T_1-weighted images were aligned with the MNI brain extracted image using an affine transformation. Finally, a nonlinear transformation (FNIRT) was used to generate a transformation from individual T_1-weighted images to T_1-weighted MNI common space (83,84).

4.3.1.2 Neuromelanin Data

Neuromelanin MRI (NM-MRI) data were used to compute LC magnetization transfer contrast (MTC) which quantifies LC neuronal density and therefore quantifies the amount of neuromelanin in a subjects' LC (7,9). Data were acquired using a magnetization-prepared 2D gradient recalled echo (GRE) sequence: TE/TR = 3.10/354 ms, 416 × 512 imaging matrix, 162 × 200 mm^2 (0.39 × 0.39 × 3 mm^3) field of view, 15 slices, flip angle = 40°, four measurements, MTC preparation pulse (300°, 1.2 kHz off-resonance,

10 ms duration), and 470 Hz/pixel receiver bandwidth with a scan time of 10 minutes and 12 seconds (8,85). The four measurements were saved individually for offline registration and averaging. NM-MRI slices were prescribed perpendicular to the dorsal edge of the brainstem in the T_1-weighted image. Two subjects chose not to participate in the neuromelanin scans, so all NM-MRI data analyzed in this chapter were done with N = 28.

To process the NM-MRI data in FSL, images from the four GRE measurements were registered to the first image using a linear transformation tool in FLIRT and averaged (83,84). A transformation between this averaged NM-MRI image and T_1-weighted image was derived using a rigid body transform with boundary-based registration cost function in FLIRT. Prior to the rigid body transformation, the T_1-weighted image was parceled into grey matter, white matter, and cerebral spinal fluid regions. The quality of each registration between T_1-weighted and NM-MRI images was assessed by overlaying the white matter boundary from the T_1-weighted image on the NM-MRI image. No significant deviation was observed in all subjects. Contrast from the magnetization transfer preparation pulse, denoted MTC, was then estimated using Eq. 8.

$$MTC = {(I - I_{ref})}/{I_{ref}} \qquad \text{Eq. 8}$$

I denotes the intensity of a voxel in the NM-MRI image and I_{ref} refers to the mean intensity of a reference region in the NM-MRI image. To ensure consistent placement of reference region in NM-MRI images across subjects, a reference region was drawn in the pons in MNI T_1-weighted common space and then transformed to individual NM-MRI images. A LC atlas in MNI space was used in this study to localize the region around LC for MTC

measurement (9). The LC atlas was transformed to NM-MRI space using the aforementioned transformations and using a threshold level of 0.5. After binarizing, each subject's mean MTC was measured in the LC ROI resulting in a total of N = 28 values.

4.3.1.3 Pupillometry Data

The pupillometry data were collected using a sampling rate of 2kHz, preprocessed using the ET-remove artifacts toolbox (github.com/EmotionCognitionLab/ET-remove-artifacts), and downsampled to match the temporal resolution of the fMRI data (37). To measure pupil dilations relative to baseline, this dataset was normalized by dividing by subject-specific means of RS0 providing us with percent signal changes. Three subjects' data were improperly collected resulting in N = 27 for most pupillometry-related calculations. These three subjects are different from the ones who decided not to participate in the neuromelanin scans.

4.3.1.4 Attention-Related Networks

The Human Connectivity Project (HCP) dataset analyzed in Chapter 3 focused on four networks previously associated with resting state: default mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), and salience network (SN). The same networks are examined here because we implemented a pseudo-resting state dataset, and because they are all related to attention in some manner. DMN (a resting state network) and DAN (an attention network) were selected because squeezing ought to invoke a transition from the resting state into a task-positive state (86); FPCN because it is linked to DAN and regulates perceptual attention (87); and SN because it determines which stimuli in our environment are most deserving of attention (37,88).

Talariach coordinates for regions of interest (ROIs) within DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and converted to MNI coordinates while SN MNI coordinates were taken directly from Raichle 2011 (1,60–62). Two ROIs from FPCN (dorsal anterior cingulate cortex and left dorsolateral prefrontal cortex) were excluded due to their close location to other ROIs.

LC was localized using the probabilistic atlas described in Langley et al. 2020 (9,10). Briefly, a reference region was drawn in the pons and its mean (μ_{ref}) and standard deviation (σ_{ref}) were calculated. Next, the LC atlas was transformed from MNI space to native NM-MRI space using the transforms described in Section XXXX using a threshold level of 0.01. After binarizing and dilating the transformed LC atlas, voxels in the dilated LC region of interest with intensities greater than $I > \mu_{ref} + (4*\sigma_{ref})$ were considered part of LC.

Table 2 shows the MNI coordinates for all networks and ROIs discussed and were used to center a $5mm^3$ isotopic marker. It displays the same information as Table 1 but also includes the LC (whose dilated voxels were split into rostral and caudal regions) (64–66). BOLD signal from each voxel within an ROI were extracted and averaged to represent the overall signal for an ROI. This was repeated for 31 total ROIs: 9 from DMN, 7 from FPCN, 6 from DAN, 7 from SN, and 2 from LC. The MNI anatomical coordinates for the four attention networks and LC are found in Table 2 and were used to center a $5mm^3$ isotopic marker (64–66). Although the LC is only a single ROI split into rostral and caudal, it is an entity distinguishable from the large-scale networks. So, it will henceforth be referred to as a network for the purposes of this dissertation.

	ROI	Abbreviation	Full Name	MNI Coordinates	Source
Default Mode Network	1	PCC	Posterior Cingulate Cortex	(2, 54, 16)	Despande et al. 2011
	2	L pIPL	Left Posterior Inferior Parietal Lobule	(-46, -72, 28)	Despande et al. 2011
	3	R pIPL	Right Posterior Inferior Parietal Lobule	(50, -64, 26)	Despande et al. 2011
	4	PFC/vACC	Orbitofrontal Cortex/Ventral Anterior Cingulate Cortex	(4, 30, 26)	Despande et al. 2011
	5	dMPFC BA 8	Dorsomedial Prefrontal Cortex Broadmann Area 8	(-14, 54, 34)	Despande et al. 2011
	6	dMPFC BA 9	Dosomedial Prefrontal Cortex Brodmann Area 9	(22, 58, 26)	Despande et al. 2011
	7	L DLPFC	Dorsolateral Prefrotnal Cortex	(-50, 20, 34)	Despande et al. 2011
	8	L PHG	Parahippocampal Gyrus	(-10, -38, -2)	Despande et al. 2011
	9	L ITC	Inferolateral Temporal Cortex	(-60, -20, -18)	Despande et al. 2011
Fronto-Parietal Control Network	1	L aPFC	Left Anterior Prefrontal Cortex	(-36, 56, 10)	Deshpande et al. 2010
	2	R aPFC	Right Anterior Prefrontal Cortex	(34, 52, 10)	Deshpande et al. 2010
	~~3~~	~~dACC~~	~~Dorsal Anterior Cingulate Cortex~~	~~N/A~~	~~Deshpande et al. 2010~~
	~~4~~	~~L DLPFC~~	~~Left Dorsolateral Prefrontal Cortex~~	~~N/A~~	~~Deshpande et al. 2010~~
	5	R DLPFC	Right Dorsolateral Prefrontal Cortex	(46, 14, 42)	Deshpande et al. 2010
	6	L aINS	Left Anterior Insula	(-30, 20, -2)	Deshpande et al. 2010
	7	R aINS	Right Anterior Insula	(32, 22, -2)	Deshpande et al. 2010
	8	L aIPL	Left Anterior Inferior parietal Lobule	(-52, -50, 46)	Deshpande et al. 2010
	9	R aIPL	Right Anterior Inferior Parietal Lobule	(52, -46, 46)	Deshpande et al. 2010
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
Locus Coeruleus	1	R LC	Rostral Locus Coeruleus	Probabilistic Atlas	Langley et al. 2020
	2	C LC	Caudal Locus Coeruleus	Probabilistic Atlas	Langley et al. 2020

Table 2: List of MNI coordinates used for ROIs in the default mode network (DMN), fronto-parietal control network (FPCN), dorsal attention network (DAN), salience network (SN), and locus coeruleus (LC). Taliarch coordinates for DMN, FPCN, and DAN were taken from Deshpande et al. 2011 and were converted to MNI using (60-63). MNI coordinates for SN and LC were taken directly from Raichle 2011 and Langley et al. 2017 respectively (1,9,10). The two FPCN coordinates with the strikethrough were excluded from analyses because of their close location to neighboring ROIs. The information for ROIs within DMN, FPCN, DAN, and SN are reproduced here, and the LC has been included.

4.3.2 Hidden Markov Model

In this investigation we aimed to focus on the temporal dynamics of the described dataset and on comparing measures of interest across two conditions (active vs. sham). According to the discussion in Section 3.4, of the three HMMs described, the activation

based HMM (AB HMM) was the ideal choice to examine these characteristics and for the goals of this study. Furthermore, preliminary tests in applying the connectivity-based hidden Markov models to this dataset yielded that the inherent smoothing factor in a windowed analysis was overwhelming and obliterated any useful dynamics. This likely occurred because the LC dataset was much smaller and slower than the HCP Unrelated 100 dataset. Because the AB HMM is the only HMM subject employed in this investigation, it will henceforth be referred to simply as HMM.

To implement the HMM, the same methods described in Section 2.3.4.1 in Chapter 2 were used (11–13). The data for RS0 and SQ1-RS5 were z-scored separately for a total of four z-scorings performed per subject: RS0 for active and sham sessions, and SQ1-RS5 (PostAr) for active and sham sessions. The normalized BOLD signal from the predefined ROIs were concatenated across all subjects then inputted into an HMM from the hmmlearn library in python (55). Outputs of interest include the Viterbi path (hidden state sequence), covariance matrices, mean state patterns, and a transition probability matrix. Connectivity states directly corresponding to each activation state were acquired by converting the outputted covariance values into Pearson correlation values using the Matlab command corrcov.m (http://www.mathworks.com/) (12,13).

4.3.3 Model Order Determination (Stability Analyses)

As described in Chapter 2, model order determination organically arises when assessing the robustness and reproducibility of state patterns. The Ranking and Averaging Independent Component Analysis by Reproducibility based (RAICAR-based) and Euclidean distance based (ED-based) methods discussed in Section 2.3.5.2.1 in Chapter 2

were employed to assess the stability of different model orders applied to this LC dataset and therefore to determine the optimal model order (11,35). Both methods required the HMM to be implemented multiple times each with a different initialization probability (the probability of starting within a certain state) where each implementation was termed a "realization". For the RAICAR-based method, the state patterns were matched across realizations via Pearson correlations and reordered so that each state possessed the same spatial pattern (with R^2 above a prespecified threshold) across all realizations (11,35). Then the spatial patterns for each state were correlated across all pairs of realizations. For example, with three realizations, there would be three total pairings because realization 1 could be paired with realization 2, realization 2 with realization 3, and realization 1 with realization 3. These R^2 values were averaged for all within-state correlations, sorted from highest to lowest values, and plotted against model order (Fig. 4). For the ED-based method, state spatial patterns were permuted and matched via the smallest Euclidean distance 100 times for all pairs of realizations (Fig. 5). The Euclidean distances from all permutations and matchings within a model order were averaged and plotted against model order.

4.3.4 Local Analyses

As mentioned in Section 3.3.2, local analyses refer to state pattern characterization. For the previous investigation we compared the acquired state patterns across the three different HMMs, however, in this investigation we used only one subtype (AB HMM) and therefore focused only on making neuroscientific interpretations about the states.

Despite examining both activation and connectivity state patterns, not all possible latent brain states in this dataset were captured (13). It is entirely possible that there were states visited in one condition (either active or sham) that were not visited in the other since a squeezing task was performed in one session. However, we fit the model to both conditions together in effort to identify globally prevalent states. Ideally, future analyses would fit an HMM to active and sham sessions separately which cannot be done with this dataset because it is slow and short (less than half the length with more than twice the TR of the HCP Unrelated 100 dataset).

4.3.4.1 Activation State Patterns

The HMM directly outputted mean state patterns. In accordance with Chen et al. 2016's method of state pattern acquisition, we also examined spatial patterns acquired by averaging the BOLD signal from TRs where the Viterbi path labeled a state to be active (Fig. 6) to determine how they compared with those from the model output (11).

In addition to examining comprehensive state patterns, we analyzed activation state patterns specific to the active and sham conditions as a means of determining whether a certain state pattern was dominant in one session and therefore driving the overall spatial patterns. These spatial configurations were acquired by averaging the BOLD signal from each subject's Viterbi path for active and sham sessions separately.

4.3.4.2 Connectivity State Patterns

Because we employed what Chapters 2 and 3 referred to as AB HMM, we were able to acquire connectivity states that directly corresponded to each activation state. These

connectivity-based states were found by mathematically converting the covariance matrices directly outputted by the model into Pearson correlations (12,13).

4.3.5 Global Analyses

Global analyses refer to measures of interest that can be acquired from any HMM applied to neuroimaging data. As with Chapters 2 and 3, the global analyses we investigated from the LC data were the Viterbi path, transition probability matrix, switching rate, fractional occupancy, and average state duration. These measures were calculated for RS0 and blocks SQ1-RS5 because this portion of the scan occurred after the handgrip task (both active and sham). They were then compared across sessions as a means of analyzing how they changed as a function of active squeeze as well as from RS0 as a means of determining how they changed from baseline.

4.3.5.1 Viterbi Path

The Viterbi path—a direct HMM output—was used to qualitatively assess differences between active and sham sessions and to obtain qualitative insight into the temporal dynamics of the LC dataset. Since the active and sham sessions were fitted together, the model provided the hidden state sequence for concatenated active and sham conditions. We therefore separated the first and second halves of the outputted Viterbi path to examine the active and sham state sequences separately. Each state was assigned a color and plotted as a function of TR for every subject.

4.3.5.2 Fano Factor

The Fano factor is a measure of the noise-to-signal ratio of a time series and should provide insight into any changes in the stability of each network within each state as a

function of active squeeze (89–91). The Fano factor was computed using the raw BOLD signal for every TR where the Viterbi path labeled a state to be active using the equation seen in Eq. 9.

$$Fano\ Factor = \frac{variance}{mean} = \frac{\sigma^2}{\mu} \quad \text{Eq. 9}$$

Using Eq. 9, the Fano factor was calculated for the entire session (all resting state and squeeze blocks together) and across all subjects because the state spatial patterns directly outputted from the HMM were fitted on a group level for the entire scan duration. Active and sham sessions were kept separate to obtain an idea of how the noise-to-signal ratio of each network and state varied as a function of the task. The Fano factor was calculated using the raw BOLD signal because it is defined as the mean divided by the variance (Eq. 9) (89–91). The mean of the z-scored BOLD signal fitted by the HMM had been centered to zero and using it would have provided an inaccurate representation of the noise-to-signal ratio. The results were plotted using boxplots for easier visualization of this measure for a network within every state and across active vs. sham. Statistical tests were performed to determine whether the Fano factor of states and networks changed as a function of the handgrip task. Furthermore, the Fano factor was correlated with LC MTC and plotted on a bar graph to determine whether a relationship between LC neuromelanin content and noisiness of networks and states existed.

4.3.5.3 State Visits

State visit analyses refer to any computation relying on information found using the hidden state sequence. The transition probability matrix is excluded from this definition

because it was directly outputted from the HMM and is explored in the next section. The differences of these state visit related measures between active and sham sessions per subject were calculated to highlight individual changes. RS0 provided baseline measures since it was a pure resting state block occurring prior to any arousal inducer. To observe the temporal progression of these measures relative to baseline, fractional occupancy, average duration, and switching rate were calculated for RS0 and for blocks SQ1 – RS5, the blocks after the squeeze prompt was administered. Because the squeeze blocks were points where arousal should have been induced via the handgrip task, blocks SQ1-RS5 will henceforth be referred to as post-arousal (PostAr). The difference in measures between RS0 and PostAr were computed to understand how they changed relative to baseline. The difference in these measures were subsequently computed across condition to understand how they changed as a function of active squeezing. To determine whether LC structure was predictive of active squeeze inducing changes in fractional occupancy and average state duration, these measures were Spearman correlated with LC MTC because we had no a priori reason to believe the two measures should be linearly related. After performing this for each state and across all subjects the correlation values were plotted on a bar graph for visualization.

4.3.5.3.1 Fractional Occupancy

The proportion of visits to a state was calculated by first counting the total number of TRs each subject spent in a certain state for each block, RS0 and PostAr. That value was then divided by the total number of TRs in each block (150 for RS0 and 375 for PostAr) to obtain the fractional occupancy of each state within either RS0 or PostAr. This measure

examined differences in the overall time spent in a state between the conditions of active and sham. That is, we calculated the fractional occupancy of a certain state in RS0 and PostAr, and for active and sham sessions to observe the temporal progression of fractional occupancy as a function of the task relative to baseline.

4.3.5.3.2 Average State Duration

Average state duration is defined as the mean time spent in a state once a subject entered that state. The average duration in a state was computed for active and sham sessions to determine whether active squeeze impacted this measure. This was done by calculating the number of consecutive TRs spent in a state once the subjects enter it, averaging it, then dividing by the total number of TRs within a block. The average duration was computed for RS0 and PostAr, and the difference between them found within each session to assess these measures relative to baseline. Subsequent statistical tests were conducted to determine whether average state duration was statistically and significantly different across state, condition, block, or any combination of them.

4.3.5.3.3 Switching Rate

The switching rate between sessions were examined as a means of determining whether performing the handgrip task affected nonspecific state transitions. This measure was calculated by counting the number of state switches each subject underwent and dividing it by the length of the block. This procedure was performed for RS0 and PostAr and the difference in switching rate between the two blocks was found to acquire switching rate within each session relative to baseline. A paired t-test was performed to determine if the mean switching rate was significantly different across conditions.

The switching rate was also examined to obtain insight into the model's ability to recognize attentional dynamics. We expected there to be a difference in switching rate between conditions because Stevner et al. 2019 found a relationship between switching rate and non-rapid eye movement sleep stage (13). They found that subjects switched between HMM-derived state more during the awake stages. Because squeezing a squeeze-ball has been shown to increase arousal causing the subjects to become more alert, we hypothesized that the switch rate would be greater during the active session.

4.3.5.4 Transition Probability Matrix

The transition probability matrix directly outputted from the HMM gave a general overview of transitions for both conditions as a whole and included baseline (RS0) effects since it was fitted to concatenated sessions. Thus, it showed transition probabilities on a global level whereas we were interested in changes in transition probabilities once the handgrip task was introduced, i.e., during PostAr. To calculate this, we concatenated Viterbi paths for active and sham sessions during the PostAr block and identified the number of times a subject transitioned out of a certain state and into another, then divided that value by the total number of transitions in the block (374). To examine transition probabilities specific to each session, we again calculated the transition probability matrix for PostAr during active and sham sessions separately. To remove the effects of baseline, the transition probability matrix was again calculated for RS0 and subtracted from each individual session's PostAr transition probability matrix. The mean of these values across all subjects was found to obtain a global average of transition probabilities specific to each session. All transition probabilities were placed on a color scale for visualization.

To determine whether LC neuromelanin content facilitated switching between particular states, state-specific transitions within each session were correlated with LC MTC. That is, once subject-specific transition probabilities were acquired and prior to obtaining a global average, each square in the matrix was Pearson correlated with LC MTC. This matrix was also visualized via color patterns and observed on a color scale.

4.3.6 Pupillometry Analyses

Pupillometry analyses paralleled some of the global analyses. Specifically, we computed the mean pupil dilation during overall switching rate (transitions amongst any state) as a function of condition to determine whether LC engagement could be observed between active and sham sessions. The hemodynamic response function caused a lag in pupil response after stimulus onset, and therefore was taken into consideration (76). Pupil dilation during overall switching rate was calculated by identifying a switch in subjects' state sequences, then calculating the difference between the normalized pupil size two TRs before the switch and the first TR after the switch. This calculation was contingent on the subject remaining in the same state for two TRs before or after the identified switch to ensure that they settled into a stable state. After calculating these differences, the average within a subject was taken and plotted in a boxplot.

Pupil dilations were also computed for state-specific transitions in a similar manner. If a subject remained in the same state two TRs before and after the switch, the difference in pupil size two TRs before a switch and one TR after a switch was found, and the states assignment during these switches was also recorded. These analyses were performed to

understand the relationship between active squeezing and switching between HMM-derived brain states. The results were plotted in a matrix and visualized on a color axis.

Once the mean subject-specific changes in pupil size during state specific transitions were found for both conditions, the differences in these values were computed to identify distinct pupil dilation patterns as a function of active squeezing. Each of these differences across session were correlated with LC MTC to determine whether LC neuromelanin content impacted pupil dilation during switches between HMM-derived latent brain states. Because some subjects' data were missing due to improper data collection or the aforementioned exclusion criterion, only subjects whose data were accounted for in both the pupillometry and neuromelanin datasets were correlated (at most N = 25). The resulting matrix of correlations was also visualized on a color axis.

4.4 Results and Discussion

4.4.1 Model Order Determination (Stability Analyses)

Fig. 25 shows the results from performing the RAICAR-based and ED-based stability analyses for model orders 3-15 using three realizations. Both plots indicate that five states were best for this investigation. This was the maximum model order where the stability values for all states remained above the predetermined threshold, and where the Euclidean distance remained as low as possible (zero) before dramatically increasing. Chen et al. 2016 and Yang et al. 2010 both used the RAICAR-based method, explored 236 ROIs and 162 independent components respectively, and employed a stability threshold of 0.8 (11,35). As discussed in Section 3.3.1 we utilized a 0.9 threshold because we examined substantially fewer ROIs (31) so our analyses and interpretations could afford to be more

stringent. However, whether we used a threshold of 0.9 or 0.8 at least one stability value for model orders six and above fell below 0.8; so, a 5-state model was selected via both standards.

4.4.2 Local Analyses

Both the activation patterns and the connectivity patterns of the five states identified in the HMM are described and interpreted in this section.

4.4.2.1 Activation State Patterns

Fig. 26a shows the activation state patterns directly outputted from the HMM. State 1 (S1) represents a DMN-dominant state because DMN showed the highest level of activation while State 2 (S2) corresponds to an attention-dominant state since DAN and SN showed the highest levels of activation. State 3 (S3) shows all networks investigated to be activated, and State 4 (S4) is the squeeze state because it was prevalent during the squeeze blocks of the paradigm as indicated by the Viterbi path in Fig. 29. Because this state was even prevalent during the "squeeze" blocks of the sham session where the

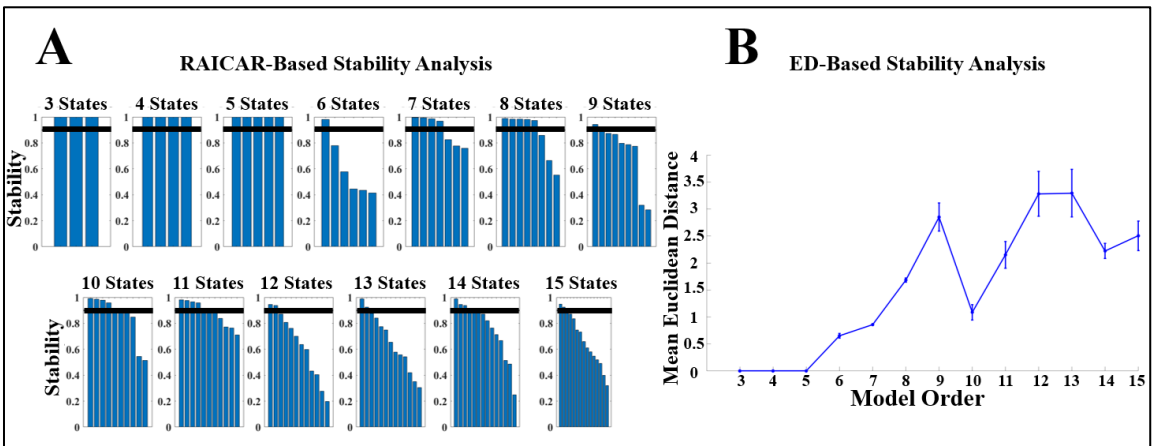

Figure 25: Stability analysis results from **(A)** RAICAR-based **(B)** and ED-based methods for AB HMM and model orders 3-15. Both plots indicate that a 5-state model was optimal because this is the largest model order where the stability values were above the 0.9 threshold in the RAICAR-based results and the mean Euclidean distance was as small as possible in the ED-based results.

subjects lifted their arm to their chest but refrained from squeezing, S4 will be referred to as the arousal state. S4 also shows a dramatic increase in the activation level of some ROIs in the attention networks, DAN and SN. Finally, State 5 (S5) shows all networks to be deactivated, likely because another network not examined in this investigation was activated (i.e., visual or networks). Furthermore, many of these states showed qualitative similarities to one another. For example, S1 and S4 visually look like they have similar patterns, but with opposite signs. DMN showed high levels of activation in S1 but was deactivated in S4. Conversely ROIs 25 and 26 in SN were deactivated in S1 but were activated in S4.

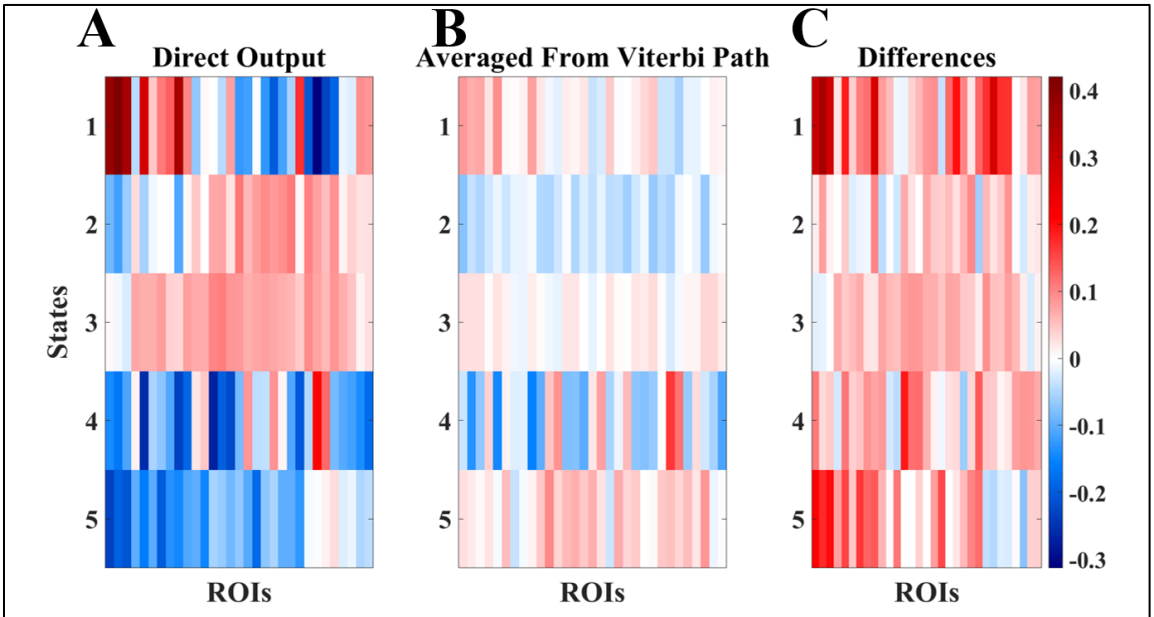

Figure 26: Activation state patterns **(A)** directly outputted from the model, **(B)** after averaging BOLD from TRs labeled by the Viterbi path using the process seen in Fig. 6, and **(C)** the difference in magnitude between them.

Fig. 26b shows the activation state patterns acquired from averaging the BOLD signal where the state sequence labeled a state to be active (as laid out in Fig. 6) and Fig. 26c shows the difference between these state patterns and those directly outputted. Although the states created from averaging the BOLD were different from the model's

direct output, in this case we did not expect them to be similar because the probability of assigning a state to a time point was not always close to 1. It is also possible that the states were subjected to noise resulting from a global average of hundreds of TRs for 30 subjects. When Viterbi-averaging the activation patterns, we assumed that the probability of residing in a certain state at any time point was absolute, or 1, but that was not always the case as sometimes the probability of a state being active within a certain TR was less than 0.8 or even less than 0.7, and we did not take into account these probabilities. Thus, when we recreated the states, we performed a simple, global average rather than a weighted average which likely accounts for most of the differences seen between Fig. 26a and Fig. 26b. This may also be why the recreated HCP states in Fig. 12 were almost exactly the same as the model output—more often than not the probability of residing in a state was between 0.9 and 1. This likely occurred because the HCP dataset was much faster, much longer, and

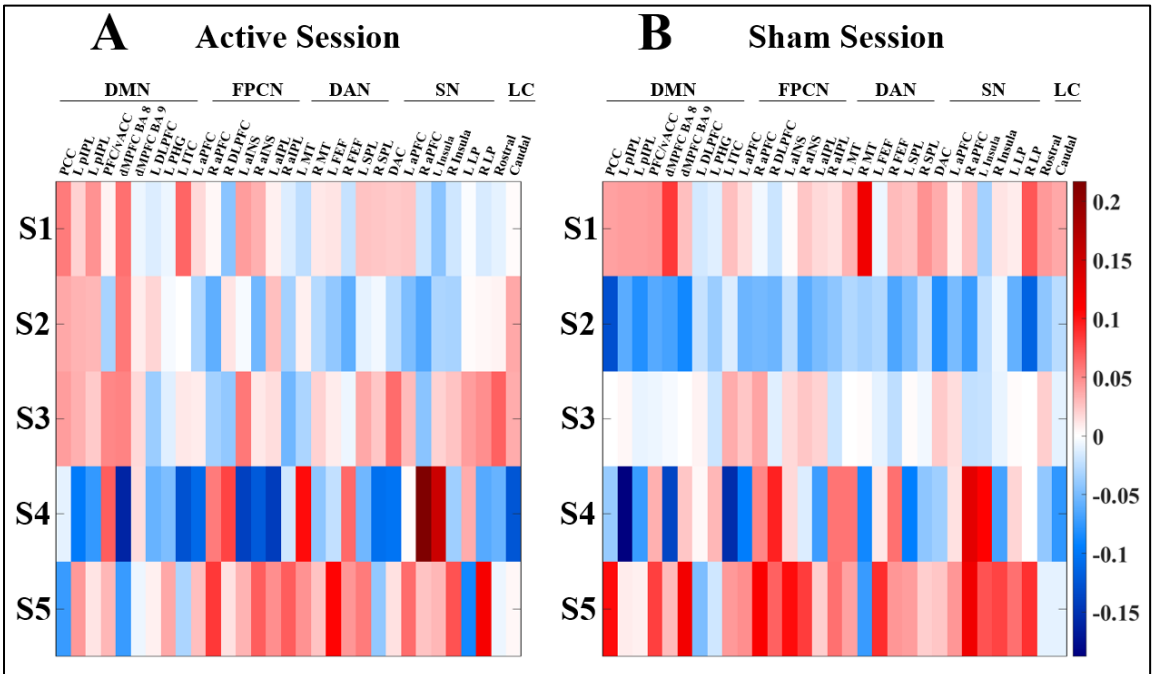

Figure 27: Activation state patterns for **(A)** active and **(B)** sham sessions after averaging BOLD from TRs labeled by the Viterbi path using the process seen in Fig. 6.

did not include a squeezing task so it was likely more sensitive to changes in BOLD fluctuations. All in all, we trusted the state patterns outputted directly from the model because Chapter 2 and 3 showed us that we could trust them (and they can potentially be the same as the averaged version) and so we moved forward with Global Analyses using the spatial patterns from the direct output (Fig. 26a).

To facilitate direct comparison with the active- and sham-specific patterns, the activation state patterns using the Viterbi averaging method (Fig. 6) was applied to active and sham sessions separately (Fig. 27a and Fig. 27b, respectively). Qualitatively, one state pattern was not dominant in one condition over the other indicating that the BOLD data in one session (active or sham) was not driving the overall, global state patterns. The differences stated above likely account for the spatial discrepancy between these plots and the model's direct output as well: the spatial patterns were subject to noise from averaging across many TRs and we did not perform a weighted average.

4.4.2.2 Connectivity State Patterns

Fig. 28a shows the raw connectivity states corresponding to the activation states after converting covariances to Pearson correlations (12,13). State patterns was subtracted from the grand mean of all other connectivity states (Fig. 28b) to highlight unique features of each pattern (Fig. 28c). For Fig. 28c, it is important to note that because the highlights are displayed, the values in the matrices are not actually R^2 values but are values relative to baseline. Thus, S1 shows a state where DMN activity was higher compared to all other but was functionally disconnected from all other networks. SN shows below baseline correlations to DMN, FPCN, and DAN as well. In S2 DAN and SN showed the highest

activation values and were functionally disconnected from all other networks. The three attentional networks (FPCN, DAN, and SN) appeared to have slightly below baseline within-network connectivity. S3 pairs the whole-brain activation (all networks examined were activated) with connectivity values around baseline. S4 links the arousal state with a connectivity state where most correlations were around baseline, but some were sporadically above or below. As discussed in Section 4.4.2.1, S4 corresponds to the arousal squeeze state so these correlation patterns correspond to a few of DAN and SN ROIs' above baseline activation levels. Whole-brain deactivation in S5 corresponded to the strongest correlations amongst all networks of interest. LC showed no distinguishing characteristics in any of these states. Although this structure did not have a substantial effect on state patterns and interpretations, we expected its effect to be prevalent when calculating measures of interest described in the Global Analyses methods section (Section 4.3.5). That is, although strong LC activity was not observed in the state patterns, we expected it to influence the differences in measures of interest calculated between the active and sham sessions explored further in the Global Analyses results and discussion section (Section 4.4.3).

Interestingly, these states showed that networks with higher activation corresponded to below baseline correlations, or functional disconnectedness. These connectivity state showed that even though some networks were deactivated in some states, that does not necessarily mean that nothing was happening. The networks were still communicating, but they may have been communicating more with themselves than with other networks. This was prevalent in S1, S2, and S4. Although DAN and SN were

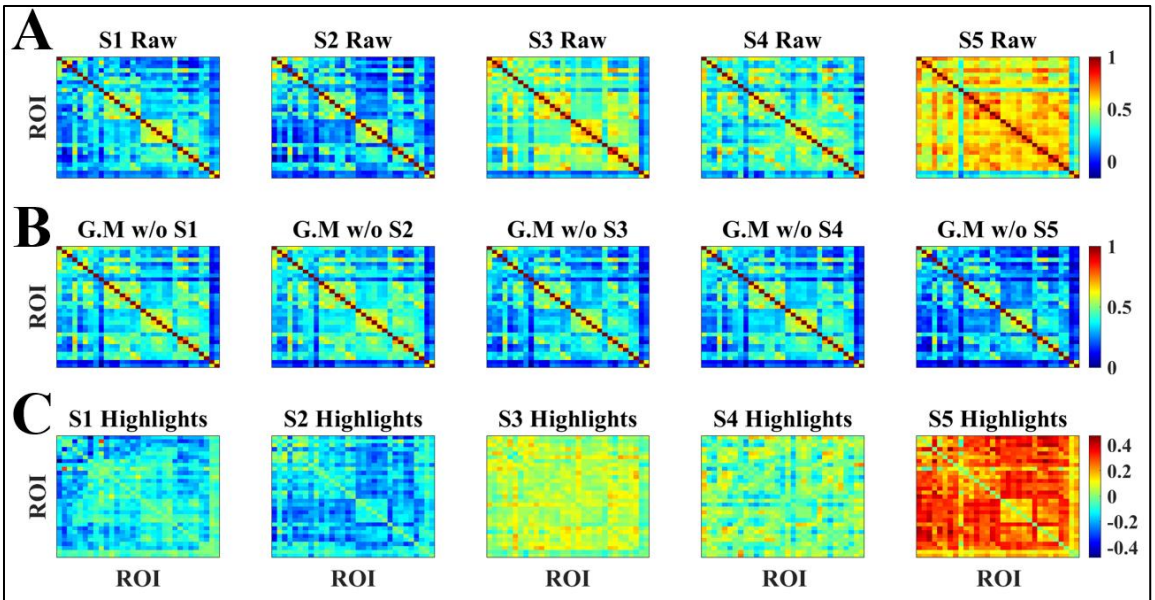

Figure 28: Connectivity states derived from outputted covariance matrices. **(A)** Raw connectivity states for states 1 through 5. **(B)** Grand mean (G.M.) for all connectivity states excluding the state indicated. **(C)** The information in the middle row is subtracted from that in the top row to obtain the connectivity state highlights observed. The color bars shown correspond to the color axis for each given row.

deactivated in S1, they were still communicating with themselves (i.e., the six DAN and seven SN ROIs were still relatively functionally connected to one another) since the R^2 values in Fig. 28c were around baseline rather than below it, and because the networks were still clearly visible around the diagonal of the matrix. The same concept was seen in S2: DMN was deactivated but its within-network functional connectivity remained higher than the functional connectivity between DAN or SN and all other networks. In S4, although all networks showed deactivation, their BOLD signal was still fluctuating as indicated by the relatively strong correlations between all four networks.

4.4.3 Global Analyses

Now that we have made sense of both the activation and connectivity states, we examine the dynamics of the data in more depth by exploring global analyses: analyses that can be performed when an activation-based or connectivity-based HMM was applied to

neuroimaging data. We explore the hidden state sequence (Viterbi path), transition probability matrix, switching, and fractional occupancy acquired when applying an HMM to the LC dataset.

4.4.3.1 Viterbi Path

The Viterbi paths for active and sham sessions are seen in Fig. 29. These plots qualitatively showed that even though the dataset was small and slow the HMM was still able to recognize dynamics in the data. Remarkably, the squeeze blocks were obvious as almost all subjects visit S4 (the orange state) in both sessions for almost the entire length of the squeeze blocks at time points 151-159, 220-228, 289-297, 448-456, and 487-495 – visible as vertical orange bars running through almost all subjects. Although we can make qualitative neuroscientific interpretations of the states and their differences in active vs. sham sessions, quantitative analyses are presented below.

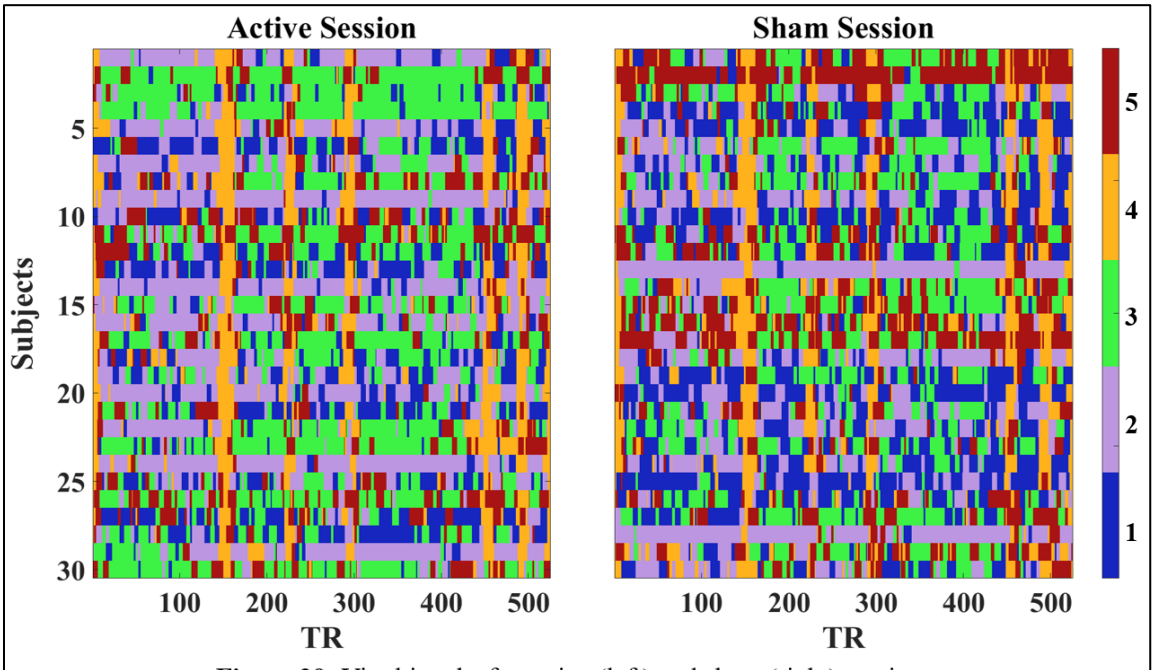

Figure 29: Viterbi paths for active (left) and sham (right) sessions.

4.4.3.2 Fano Factor

The Fano factor (Eq. 9) was calculated for every network within each state to assess their spatial patterns' susceptibility to noise as LC activity was up-regulated (Fig. 30). Two outlier subjects were removed (resulting in N = 28) because their Fano factor for FPCN in S2 was greater than 20 which greatly skewed the results and prevented adequate analysis of the other states' and networks' noisiness. A 2 (condition: squeeze vs control) x 5 (state) x 5 (network: DMN, FPCN, DAN, SN, LC) repeated measures analysis of variance (ANOVA) was conducted to determine whether the mean Fano factor was statistically different across these factors. That is, the ANOVA determined whether the mean Fano factor across all observations (subjects) as a function of condition regardless of state or network was statistically different from mean Fano factor as a function of state regardless of conditions or network, or from mean Fano factor as a function of networks regardless of condition or state. It also tested whether the level of one factor depended on that of any other, or, interactions. Because the Fano factor is a ratio of the noise-to-signal of a time series, the data were not normally distributed and needed to be transformed prior to executing the ANOVA. To do this, the data were $\log_{10}$ transformed.

The omnibus ANOVA showed significant main effects of state ($F(4,108) = 12.858$, $p < 0.01$) and network ($F(4, 108) = 39.920$, $p < 0.01$), but not in condition ($F(1,27) = 0.535$, $p = 0.471$), or in any interaction (condition*state: $F(4,108) = 0.336$, $p = 0.826$; condition*network: $F(4,108) = 1.051$, $p < 0.378$; state*network: $F(16,432) = 1.345$, $p = 0.256$; condition*state*network: $F(16,432) = 0.264$, $p = 0.884$) when using the Greenhouse-Geisser correction for sphericity violations (according to Mauchly's test of

sphericity). Because Fano factor did not statistically change as a function of condition, nor did this factor interact with the others, the data were collapsed across condition to simplify analysis of state and network. Consequently, a step-down 5 (state) x 5 (network) repeated measures ANOVA was performed. This tested whether the mean Fano factor of the networks regardless of states was significantly different from the mean Fano factor of all states regardless of network, as well as the interaction between them despite condition. Since the dataset was relatively short, rather than averaging the values across conditions, the raw active and sham BOLD signals were concatenated and the Fano factor recalculated (Fig. 31a). This alternative method provides a holistic, global sense of the dataset and was consistent with the data format fitted by the HMM to obtain the states' spatial patterns. As before, the recalculated Fano factors were $\log_{10}$ transformed prior to implementing the 5 x 5 repeated measures ANOVA.

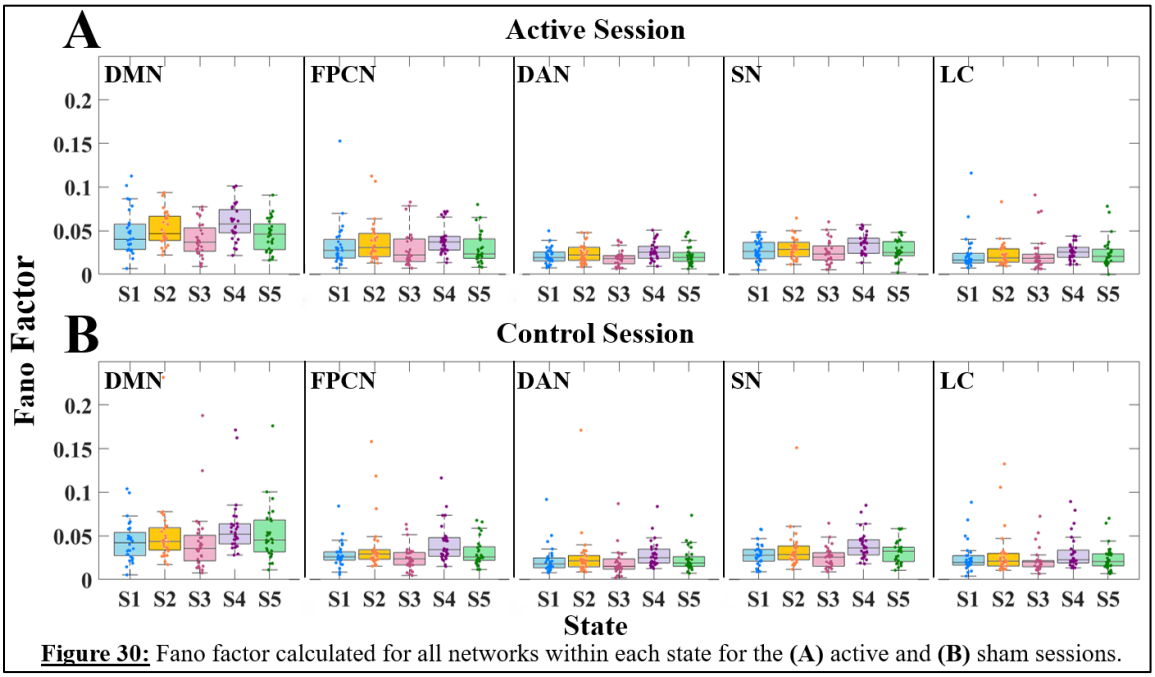

Figure 30: Fano factor calculated for all networks within each state for the **(A)** active and **(B)** sham sessions.

This step-down 5 x 5 ANOVA showed significant main effects of state ($F(4,108) = 5.623$, $p < 0.01$) and network ($F(4,108) = 8.885$, $p < 0.01$), but not in an interaction ($F(16,432) = 1.684$, $p = 0.159$) between them when using the Greenhouse-Geisser correction for sphericity violations. A lack in significant interaction prevents us from performing any more step-down ANOVAs, but the significant main effect of state allows us to look at deviation contrasts for both state and network. Deviation contrasts for state determined whether the Fano factor of each state was significantly different from the global Fano factor of all states excluding S5. The same interpretation was applied when examining deviation contrasts for the factor of networks: deviation contrasts decided whether each network's Fano factor was significantly different from the global Fano factor of DMN, FPCN, DAN, and SN, excluding LC.

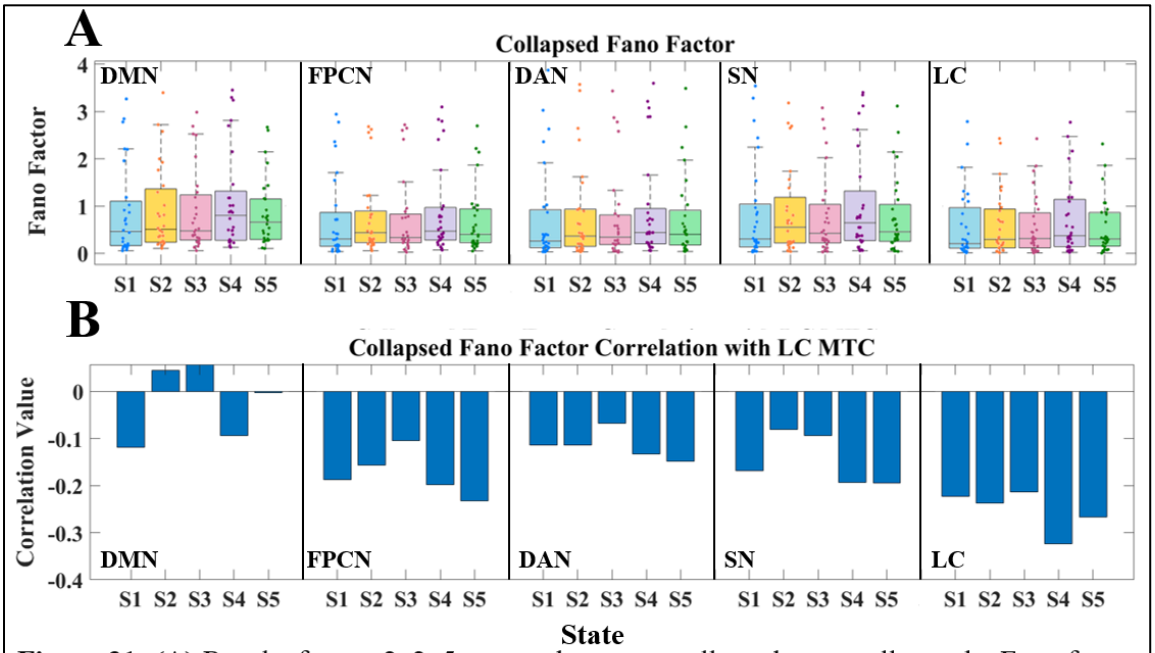

Figure 31: **(A)** Results from a 2x2x5 repeated measures allowed us to collapse the Fano factor across condition. The Fano factor was recalculated when concatenating active and sham conditions, and the values plotted. **(B)** The $\log_{10}$ Fano factor independent of condition was Spearman correlated with LC MTC, but no ρ value reached statistical significance. "*" corresponds to significant correlation values while "†" corresponds to trending values.

The deviation contrasts showed that on average and regardless of condition and state, DMN had a significantly different Fano factor than the global level ($F(1,27) = 15.684$, $p < 0.01$). In terms of state, S1 (DMN-dominant state) and S4 (arousal state) had significantly different Fano factors than the global Fano factor ($F(1,27) = 4.668$, $p = 0.04$) and $F(1,27) = 37.313$, $p < 0.01$), respectively).

These results, taken from the main effects of network and state, potentially indicate that DMN-related behavior was disrupted the most of the four networks. This could have occurred because the dataset analyzed had a pseudo resting state paradigm, so the squeeze (or the prompt to squeeze) disrupted resting state and therefore disrupted the core resting state network, DMN. S4 exhibiting significantly more noise than the global noise likely alludes to the fact that it was the shortest state in duration. Although S4 was exhibited in both sessions, active squeezing occurred in the active session while simply holding the ball occurred in the sham session. This minor discrepancy in the task potentially explains the relatively large amount of noise in the spatial pattern of S4 (the "squeeze", or arousal, state). However, because the omnibus ANOVA showed no significance as a function of condition, this is merely a qualitative interpretation. Nevertheless, the subjects were doing something different during these blocks (bringing their arm up to their chest and either squeezing or touching the ball) than all other blocks (resting state) within a session.

LC MTC values (acquired using the method outlined in Sections 4.3.1.2 and 4.3.1.3) was Spearman correlated with the $\log_{10}$ transformed Fano factor calculated for concatenated active and sham session for every network within every state (Fig. 31b) to determine whether LC neuromelanin content affected the noisiness of any network or state.

Stars (*) indicate raw significant R^2 values while daggers (†) correspond to R^2 values nearing significance. Although one correlation produced significant p-values, no significant relationship between LC structure and the noise-to-signal ratio of any network within any state was observed. This lack of affiliation may have occurred because we examined a younger population where neuronal loss was expected to be minimal, if occurring at all (72–74).

4.4.3.3 State Visits

In this section we analyze all calculations related to visiting a state: proportion of time spent in a state (fractional occupancy), the average time spent in a state once a subject entered it (average state duration), as well as the frequency of transitioning between states (switching rate). These measures were also Spearman correlated with LC MTC.

4.4.3.3.1 Fractional Occupancy

The proportion of time spent in a state within RS0 and PostAr were calculated and plotted (Fig. 32) to observe any changes (relative to baseline) in the overall amount of time spent in a certain state as a function of active squeezing. A 2 (condition: squeeze vs control) x 2 (block: RS0 vs PostAr) x 5 (state) repeated measures ANOVA was conducted to test whether the mean fractional occupancy of a state across all observations (subjects) of one element (state, block, or condition) when holding all other elements equal was significantly different from the mean fractional occupancy as a function of another element when holding all else equal. It also tested for any interactions between these factors: whether the fractional occupancy as a function of one factor depended on another. A Lilliefors test showed that the data were not normally distributed and needed to be transformed prior to

running the ANOVA. Although the fractional occupancy is a proportion, the data could not simply be $\log_{10}$ transformed because they contained zeros. This would result in values of infinity thereby making it impossible to perform any statistical test. Instead, following (92), the data were square root arcsine transformed where zeros were replaced with $^1/_{4n}$ and ones with $100 - {}^1/_{4n}$ where n = length of each block (150 for RS0 and 375 for PostAr). Then the square root of the arcsine of each data point were found (92).

The omnibus ANOVA showed main effects of block ($F(1,29) = 7.669$, $p = 0.010$), and state ($F(4,116) = 3.539$, $p = 0.039$), and an interaction between the two ($F(4,116) = 17.088$, $p < 0.01$) when using Greenhouse-Geisser corrections for sphericity violations. Condition (active vs. sham) was not found to be significant as a main effect ($F(1,29) =$

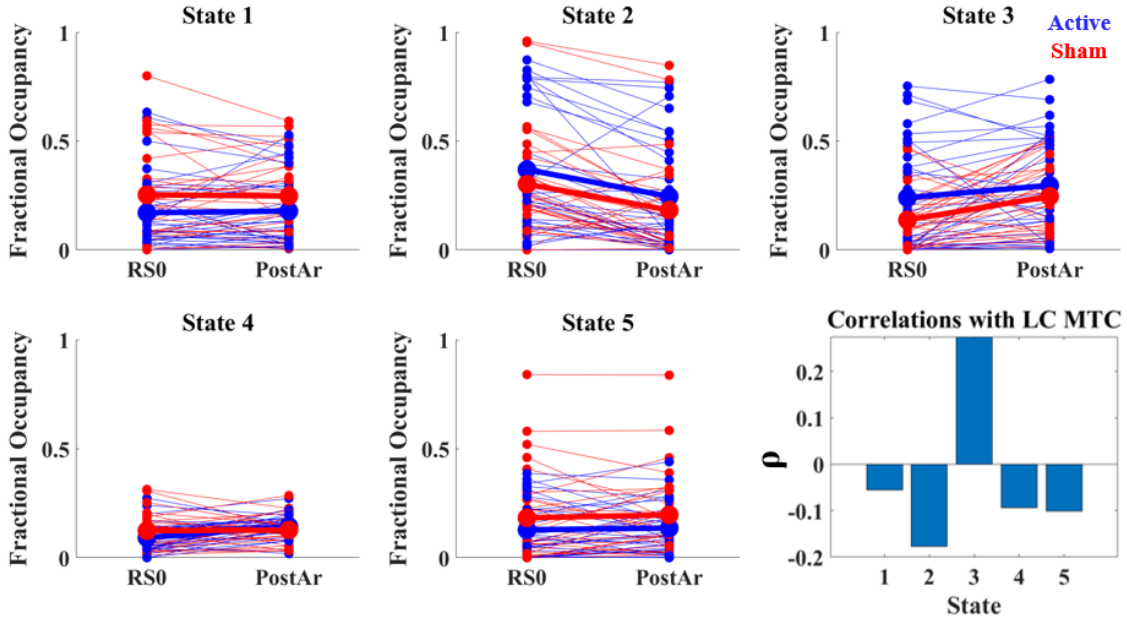

Figure 32: The fractional occupancy of states 1 through 5 for baseline (RS0) and all blocks post-arousal (PostAr) for the active (blue) and sham (red) sessions. The larger dots and thicker connecting line indicate the mean of the corresponding measure. The subplot in the lower right corner shows the correlation between LC MTC and the difference in fractional occupancy (FO) across block and condition $[(FO_{Active\_PostAr} - FO_{Active\_RS0}) - (FO_{Sham\_PostAr} - FO_{Sham\_RS0})]$ after undergoing the transformation. Thus, it shows the Spearman correlation between LC MTC and the fractional occupancy of a state as a function of LC activity up-regulation relative to baseline. No ρ values reached significance or neared significance.

2.002, p – 0.168) or in any interaction (condition*block: $F(1,29) = 1.057$, $p = 0.312$; condition*state $F(4,116) = 1.876$, $p = 0.147$); condition*block*state: $F(4,116) = 0.844$, $p = 0.425$) meaning the average time spent in any state in any block (RS0 or PostAr) in the active session was not significantly different from its sham session counterpart. This allowed us to collapse the data across condition and to examine whether the mean fractional occupancy as a function of only state was significantly different from the mean fractional occupancy as a function of only block regardless of condition. This was tested using a step-down 2 (block: RS0 vs PostAr) x 5 (state) repeated measures ANOVA. Because the dataset was relatively short, the active and sham Viterbi paths were concatenated for every subject and the fractional occupancy recalculated rather than averaging these values across condition (Fig. 33). The data again were square root arcsine transformed prior to implementing the ANOVA.

The step-down 2 x 5 repeated measures ANOVA yielded significant main effects of block ($F(1,29) = 5.257$, $p = 0.029$) and state ($F(4,116) = 4.137$, $p = 0.023$) as well as an interaction between the two ($F(4,116) = 15.128$, $p < 0.01$) when using the Greenhouse-Geisser correction for sphericity violations. Thus, the mean fractional occupancy of each state depended on both residing within a specific state and residing within a specific block. The interaction between both factors allowed us to perform additional step-down tests to assess statistical significance between certain states and across block. Specifically, we performed two step-down one-way ANOVAs with least square difference contrasts to assess pairwise fractional occupancy comparisons of specific states within either block

(92). We also performed paired t-tests to determine if the mean fractional occupancy of a particular state was different across block (Fig. 33).

The step-down paired t-tests (Fig. 33a) showed that subjects spent more time in S3 (p = 1.1644e-04, t = -4.4504) and in S4 (p = 0.0038, t = -3.1422) during PostAr than RS0

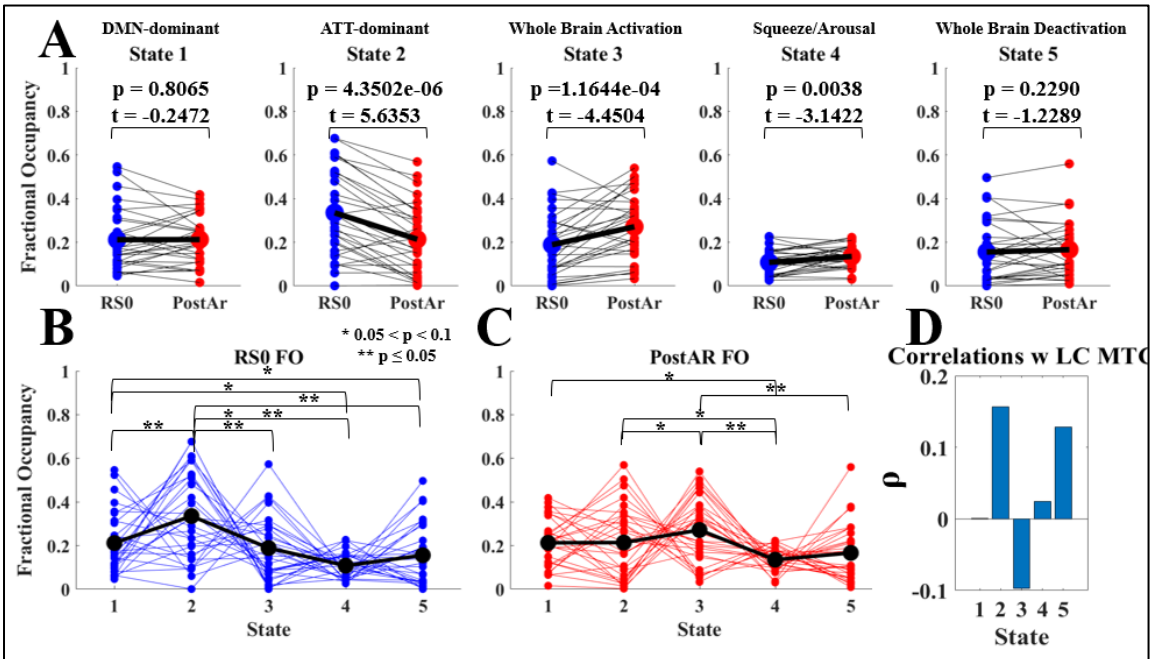

Figure 33: Results from a 2x2x5 repeated measures ANOVA allowed us to collapse fractional occupancy measures across condition. **(A)** The active and sham sessions were concatenated and the fractional occupancy for each state was recalculated for baseline (RS0; blue) and for post-arousal (PostAr; red). For easier visualization of one-way ANOVA results, the fractional occupancies were replotted so that all values during **(B)** RS0 and **(C)** PostAr are plotted together. The larger dots and thicker connecting line indicate the mean of the corresponding measures. **(D)** The Spearman correlation between LC MTC and the difference in fractional occupancy (FO) across block [$FO_{ConcatenatedCondition\_PostAr} - FO_{ConcatenatedCondition\_RS0}$] are shown. Thus, it shows the degree to which LC MTC and fractional occupancy during different blocks (RS0 or PostAr) were correlated regardless of condition, but none reached significance or neared significance.

regardless of whether they actively squeezed or simply held the ball. We expected subjects to reside in the arousal state (S4) more once they were presented with the prompt to squeeze since actual squeeze blocks occurred in PostAr and only sporadically occurred during RS0. The higher fractional occupancy of the whole-brain activation state (S3) during PostAr than in RS0 is consistent with the idea that the handgrip task is meant to increase arousal (37). These t-tests also showed that subjects occupied S2 (p = 4.3502e-06, t = 5.6353)

more during RS0 than in PostAr. Because previous studies have shown that the handgrip task increases arousal causing subjects' salient attention selectivity to increase, our results may initially appear counterintuitive since they show less time spent in the DAN/SN-dominant state after performing the handgrip task. However, while it is possible that the overall proportion of time spent in S2 did not increase as a function of the handgrip task (regardless of whether the subjects actively squeezed), it is possible that the time spent in S2 once subjects entered it, or the probability of transitioning into S2 during PostAr may still have increased. These measures may still indicate that subjects' attention processing increased post-handgrip and are explored in the next few sections.

The first step-down one-way ANOVA showed that, on average, there was a statistically significant difference in the mean fractional occupancy of all five states during RS0 ($F(4,145) = 8.906$, $p < 0.01$). Fig. 33b shows the fractional occupancy during RS0 of S1 through S5. Least square difference contrasts showed that compared to S4, subjects spent more time in S1 ($\mu_{diff} = 0.12389$, $\sigma_{SE} = 0.04384$, $p < 0.01$) and in S2 ($\mu_{diff} = 0.23893$, $\sigma_{SE} = 0.04384$, $p < 0.01$). This is to be expected because S4 corresponds to the squeeze blocks which occurred only during PostAr, so S4 occurred sporadically during RS0. Subjects also spent more time in S1 compared to S5 ($\mu_{diff} = 0.08756$, $\sigma_{SE} = 0.04384$, $p = 0.048$) which was to be expected because subject underwent pure resting state during RS0. Thus, they spent more time in a state where the foremost resting state network was activated rather than in one where that network (along with all others investigated) were deactivated (93). Furthermore, these statistical tests showed that subjects also spent more time in S2 than in S1 ($\mu_{diff} = 0.11504$, $\sigma_{SE} = 0.04384$, $p = 0.010$), in S3 ($\mu_{diff} = 0.15868$, $\sigma_{SE} = 0.04384$,

$p < 0.01$), or in S5 ($\mu_{diff} = 0.20260$, $\sigma_{SE} = 0.04384$, $p < 0.01$). Although we expected subjects to spend less time in an attention-dominant state (S2) during pure resting state, this may have occurred due to an increase in arousal from adjusting to the MRI atmosphere (37,38). This also could account for why the fractional occupancy of S1 was not significantly greater than that of S3.

The second step-down one-way ANOVA (Fig. 33c) showed that there was a main effect in the mean fractional occupancy amongst all five states during PostAr ($F(4,145) = 3.964$, $p < 0.01$). Least square difference contrasts showed that subjects occupied S1 ($\mu_{diff} = 0.08449$, $\sigma_{SE} = 0.04384$, $p = 0.035$) and S3 ($\mu_{diff} = 0.14086$, $\sigma_{SE} = 0.04384$, $p < 0.01$) more so than S4. This could have occurred because S4 was associated with the squeeze blocks and therefore only occurred for 45 total TRs out of the entire 375. It is possible that no statistical difference between the time spent in S4 and S5 was observed because the effects of the handgrip task (regardless of active squeezing) induced arousal and consequently prevented subjects' attentional networks from being deactivated for too long. This was likely also attributed to subjects spending more time in S3 than in S5 ($\mu_{diff} = 0.12201$, $\sigma_{SE} = 0.04384$, $p < 0.01$): the periodic squeeze blocks likely induced whole-brain activation more often than whole-brain deactivation as a result of increasing arousal overall. Curiously, we did not observe subjects to spend more time in S2 than in S4, a phenomenon that was expected since the handgrip task was designed to increase attentional processing. As aforementioned, although we did not observe subjects to reside in S2 more than other states once the arousal blocks were introduced, it is possible that an increase in attentional processing related measures was observed via average state duration and/or

switching into S2 from S4. These statistical tests also showed that subjects spent more time in S3 than in S2 ($\mu_{diff} = 0.07915$, $\sigma_{SE} = 0.04384$, $p = 0.047$) which might have occurred because previous studies have shown that the handgrip task induces fronto-parietal activation, and our fronto-parietal control network (FPCN) is activated in S3, but not in S2. Furthermore, it is possible that although some states did not exhibit significant differences in their fractional occupancies across session or across blocks, the trajectory through state space was different. No additional information about LC's dynamic relationship with attention could be gleaned from examining the proportion of time spent in a state since the factor of condition, did not reach significance in the omnibus ANOVA.

To determine whether LC neuromelanin content had any relation to the fractional occupancy of a state, the two values were correlated across all subjects. To do this, the proportion of time spent in each state for PostAr and baseline (RS0) were subtracted within each session. These results were then subtracted across conditions and Spearman correlated with LC MTC. That is, we performed the following calculation where fractional occupancy is abbreviated as FO: FO Correlated = $(FO_{Active\_PostAr} - FO_{Active\_RS0}) - (FO_{Sham\_PostAr} - FO_{Sham\_RS0})$. The correlations seen in Fig. 32 show the degree to which LC MTC was associated with the fractional occupancy of all five states as a function of active squeeze and relative to both sessions' baseline. Because no significant differences in fractional occupancy across condition reached significance in the omnibus ANOVA, a similar procedure was used for the Spearman correlation with the collapsed values in that the difference in factional occupancy was calculated between PostAr and baseline regardless of session: FO Correlated = $FO_{ConcatenatedCondition\_PostAr} - FO_{ConcatenatedCondition\_RS0}$.

These results (Fig. 33d) indicate how correlated LC neuromelanin content was with the fractional occupancy of all five states relative to baseline and regardless of whether active squeezing occurred. However, no value appeared to be significant or nearly significant indicating that we could not establish a relationship between LC neuromelanin content and fractional occupancy of an HMM-derived brain state.

4.4.3.3.2 Average State Duration

The average duration of a state between sessions and relative to baseline (RS0) was examined to determine whether the handgrip task affected the average time subjects remained in a state once entering it (Fig. 34). A 2 (condition: squeeze vs control) x 2 (block: RS0 vs PostAr) x 5 (state) repeated measures ANOVA was performed. This was to determine whether the mean duration of a state regardless of block or condition was significantly different from the mean duration of a state across condition regardless of block or state assignment, or from the mean duration of a state across RS0 or PostAr regardless of condition or state assignment. The ANOVA also tested for interactions between these factors: whether the duration of a state was dependent on any combination of state assignment, condition, or block. As with fractional occupancy, the average state duration data needed to be transformed prior to conducting the ANOVA. Again, a $\log_{10}$ transformation would not suffice because the data contained 0s and $\log_{10}(0) = \infty$. The square root arsine transformation was again performed using the aforementioned imputation and transformation (92).

The 2 x 2 x 5 repeated measures ANOVA showed significant main effects of state ($F(4,112) = 5.449$, $p < 0.01$) and block ($F(1,28) = 354.983$, $p < 0.01$) as well as an

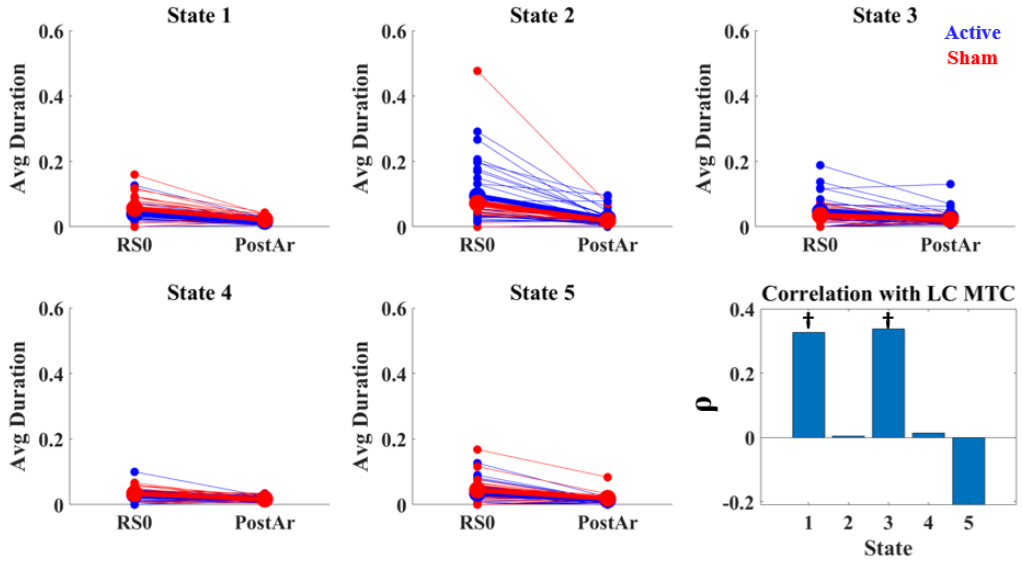

Figure 34: The average state duration of states 1 through 5 for baseline (RS0) and all blocks post-arousal (PostAr) for the active (blue) and sham (red) sessions. The larger dots and thicker connecting line indicate the mean of the corresponding measures. The subplot in the lower right corner shows the Spearman correlation between LC MTC and the difference in average duration (AD) across block and condition $[(AD_{Active\_PostAr} - AD_{Active\_RS0}) - (AD_{Sham\_PostAr} - AD_{Sham\_RS0})]$. Thus, it shows the correlation between LC MTC and the average duration of a state across condition and relative to baseline. "*" corresponds to significant ρ values while "†" corresponds to trending values.

interaction between them ($F(4,112) = 12.975$, $p < 0.01$) but no significance with condition (condition: $F(1,28) = 2.355$, $p = 0.136$; condition*state = $F(4,112) = 2.042$, $p = 0.127$; condition*block = $F(1,28) = 0.364$, $p = 0.551$; condition*state*block: $F(4,112) = 0.527$, $p = 0.640$). These results indicate that there was no statistical difference in average duration as a function of condition for any state within either block, nor did duration during active or sham session impact the other factors. For easier visualization and analysis of the state and block factors, the data were collapsed across condition. Because the dataset was relatively short, rather than averaging the state duration values across conditions, the active and sham Viterbi paths were concatenated, then the average duration recalculated (Fig. 35). The results were again square root arcsine transformed prior to conducting the ANOVA. A step-down 2 (block) x 5 (state) repeated measures ANOVA was implemented to

determine whether there was a statistically significant difference between the mean duration of a specific state within either block and the mean duration of any of the states across RS0 and PostAr.

The step-down repeated measures ANOVA showed a significant main effect of state ($F(4,112) = 4.635$, $p = 0.013$) and an interaction between state and block ($F(4,112) = 6.403$, $p < 0.01$), but not a main effect of block ($F(1,28) = 1.932$, $p = 0.176$). The significant interaction allowed us to perform additional step-down statistical tests. We conducted two one-way ANOVAs with least square difference contrasts testing for significant differences in mean state duration across specific states within each block. Paired t-tests were also executed to determine the significance of changes in mean duration relative to baseline.

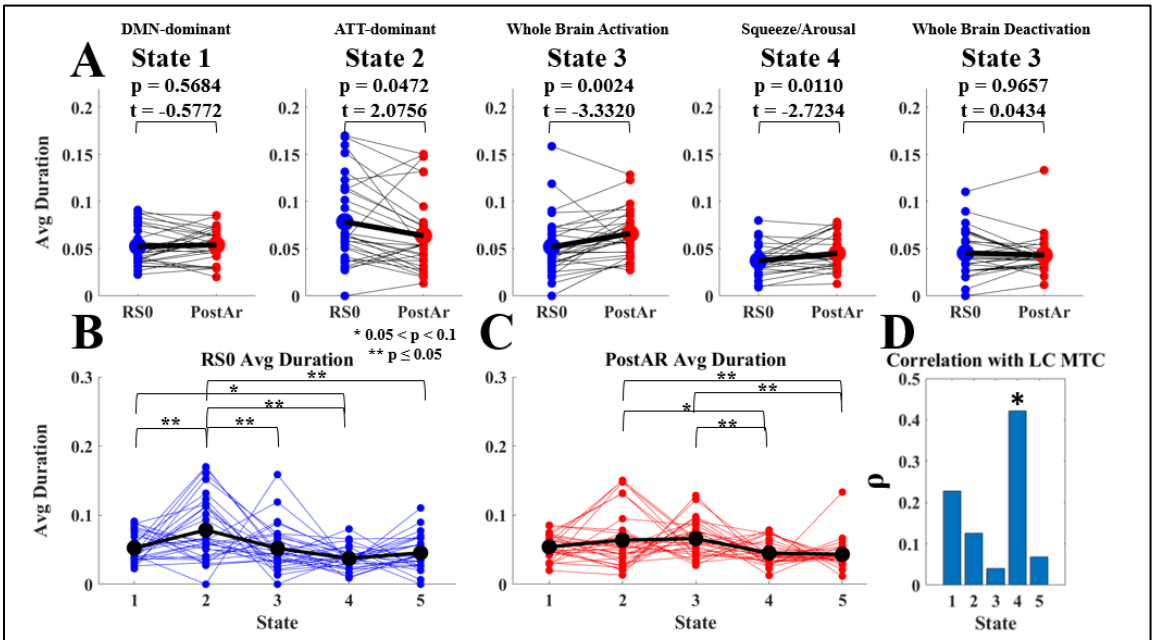

Figure 35: Results from a 2x2x5 repeated measures ANOVA allowed us to collapse average duration across condition. **(A)** The active and sham sessions were concatenated and the mean duration for each state was recalculated for baseline (RS0; blue) and for post-arousal (PostAr; red). For easier visualization of one-way ANOVA results, all state duration values during **(B)** RS0 and **(C)** PostAr were plotted together. The larger dots and thicker connecting line indicate the mean of the corresponding measures. **(D)** The subplot in the lower right corner shows the Spearman correlation between LC MTC and the difference in fractional occupancy (AD) across block $[AD_{ConcatenatedCondition\_PostAr} - AD_{ConcatenatedCondition\_RS0}]$. Thus, it shows the degree which LC MTC was correlated with the average duration of a state relative to baseline and regardless of condition. "*" corresponds to significant ρ values while "†" corresponds to trending values.

The step-down paired t-test (Fig. 35a) showed that once subjects entered S3 or S4, they remained in that state longer during PostAr than RS0 (S3: $p = 0.0024$, $t = -3.3320$; S4: $p = 0.0110$, $t = -2.7234$). This makes sense regarding S4 since this state was prevalent during the squeeze blocks occurring only during PostAr and therefore occurred sporadically during RS0. Regarding the whole-brain activation state (S3), it is likely that subjects entered a state where all attention networks were activated in response to the arousing handgrip task, regardless of whether they actively squeezed. Subjects also exhibited a longer average duration of S2 during RS0 than PostAr ($p = 0.0472$, $t = -2.0756$), consistent with results from the previous section showing fractional occupancy of S2 to be higher in RS0 than in PostAr. Because both fractional occupancy and state duration of S2 were reduced in PostAr compared to RS0, this could be indicative of norepinephrine (NE) depletion. Previous animal literature has shown NE levels to decrease with the onset of an arousing stressor (37,94,95). Our results could be evident of this phenomenon because the arousing handgrip task was expected to increase subjects' salient attentional processing via increased residence in S2 (37), but instead it significantly decreased. However, invasive histological tests are needed to confirm this neuroendocrinological event (37,96).

The first step-down one-way ANOVA showed a main effect for mean state duration within RS0 ($F(4,140) = 6.035$, $p < 0.01$). The plots from Fig. 35a were adjusted so that the average duration of all states during RS0 were included in one plot, which is illustrated in Fig. 35b. Least square difference contrasts showed subjects to remain in S1 ($\mu_{diff} = 0.03663$, $\sigma_{SE} = 0.01687$, $p = 0.032$) longer than S4 once entered, likely because S4 occurred during the squeeze blocks in PostAr and were relatively short. We expected subjects to

remain in a DMN-dominated state more so than an arousal state during RS0 since this block occurred prior to the handgrip task and therefore was pure resting state. These contrasts also showed subjects to remain in S2 longer than S1 ($\mu_{diff} = 0.04131$, $\sigma_{SE} = 0.01687$, $p = 0.0.16$), S3 ($\mu_{diff} = 0.04926$, $\sigma_{SE} = 0.01687$, $p < 0.01$), S4 ($\mu_{diff} = 0.04926$, $\sigma_{SE} = 0.01687$, $p < 0.01$), or S5 ($\mu_{diff} = 0.06237$, $\sigma_{SE} = 0.01687$, $p < 0.01$) which may have occurred because of the scanner environment. It is possible that adjusting to the MRI atmosphere may have generated some stress triggering subjects' LC to engage via increases in sympathetic arousal (37,38). This may have caused subjects not only to reside in S2 more so than all other states, but also to persistently remain in it.

The second step-down one-way ANOVA tested for differences in mean state duration within PostAr and showed a general significance for all five states ($F(4,140) = 4.580$, $p < 0.01$). Fig. 35c shows the average duration of states one through five during PostAr along with the results from the ANOVA included. Least square difference contrasts demonstrated that subjects remained in S2 ($\mu_{diff} = 0.03247$, $\sigma_{SE} = 0.01371$, $p = 0.019$) and in S3 ($\mu_{diff} = 0.04376$, $\sigma_{SE} = 0.01371$, $p < 0.01$) longer than in S4. This likely is explained by S4 occurring during the shortest block (only occurring for nine TRs at a time for five times) so we expected subjects to have the least duration in this state. We also observed that, compared to S5, subjects had a longer duration in S2 ($\mu_{diff} = 0.03726$, $\sigma_{SE} = 0.01371$, $p < 0.01$) and in S3 ($\mu_{diff} = 0.04855$, $\sigma_{SE} = 0.01371$, $p < 0.01$). The longer duration of S2 than S5 is likely accounted for because the handgrip (regardless active and sham) prompted an increase in arousal, so subjects tended to persist more in the attention-dominant state (S2) (even if they did not spend more time in this state overall) than in one where all

networks were deactivated as a result of the handgrip task regardless of condition. While this was likely for the active session where the subjects actively perform the handgrip task, it is also possibly explained in the sham session. Although the subjects did not squeeze, they were still given the prompt and still brought their arm to their chest which might have triggered a small change in LC activity as evidenced by previous work (81). Paralleling results from the fractional occupancy analysis, the duration of S3 may have been longer than S5 because the intermittent squeeze blocks induced arousal increasing the likelihood of all attention networks investigated being activated rather than deactivated. This may also explain why S1 and S5 were not sustained longer than S4: the handgrip task prevented subjects from remaining in states where a resting state network was activated (S1) or where all attention-related networks were deactivated (S5). Moreover, it is possible that the subjects persisted in the states for insignificantly different durations across block and condition, but the state path they traversed varied. No additional information about how LC modulated attention could be gleaned from examining the average duration of a state since the factor of condition did not reach significance in the omnibus ANOVA.

LC MTC was Spearman correlated with the average duration of all five states before (Fig. 34d) and after (Fig. 35d) collapsing the data across condition. The ρ values in Fig. 34 were acquired by correlating LC MTC with the difference in average duration (AD) of a state from baseline and across block, i.e., Correlated AD = $(AD_{Active\_PostAR} - AD_{ActiveRS0}) - (AD_{Sham\_PostAr} - AD_{Sham\_RS0})$. That is, these correlations show the degree to which LC MTC was correlated with the average duration of all five states as a function of condition and relative to baseline. The plot displays that two trending correlations are seen.

This means that the amount of LC neuromelanin content affects the total time spent in S1 (DMN-dominant) and almost affects that of S3 (whole-brain activation). However, the values shown have not taken into account any corrections of family-wise error rates.

Because the omnibus ANOVA determined average duration of any state to be insignificantly different across active and sham sessions, ρ was recalculated for the collapsed mean duration values. The ρ values in Fig. 35d were acquired by Spearman correlating LC MTC with the difference in average duration across baseline, i.e., Correlated AD = $AD_{ConcatenatedCondition\_PostAr} - AD_{ConcatenatedCondition\_RS0}$. Thus, these ρ values show the correlation between LC neuromelanin content and average duration relative to baseline. It appears as if there is a significant correlation between LC MTC and the average duration of S4 as a function of the handgrip task. This would indicate that the amount of neuromelanin in a subjects' LC impacts how long a subject persists in the arousal state. However, the values shown have not taken into account any corrections of family-wise error rates, so it is possible that there is actually no statistical significance between average duration of S4 regardless of condition and LC MTC.

4.4.3.3.3 Switching Rate

While neither the overall time in nor the average duration of each state showed significant differences between active and sham sessions, it is possible that the rate of transitioning between states might change across condition. Fig. 36 shows the overall switching rate for both sessions regardless of which state subjects were transitioning between and relative to baseline. However, a paired t-test showed that this switching rate was not significant ($p = 0.2197$, $t = 1.4537$). These results indicate that actively squeezing

did not induce an increase or a decrease in the rate at which subjects switched between any state. While the transition between any state was not significantly different between sessions, it is possible that the handgrip task affected state-specific transitions which is explored further in the next section via transition probability matrices.

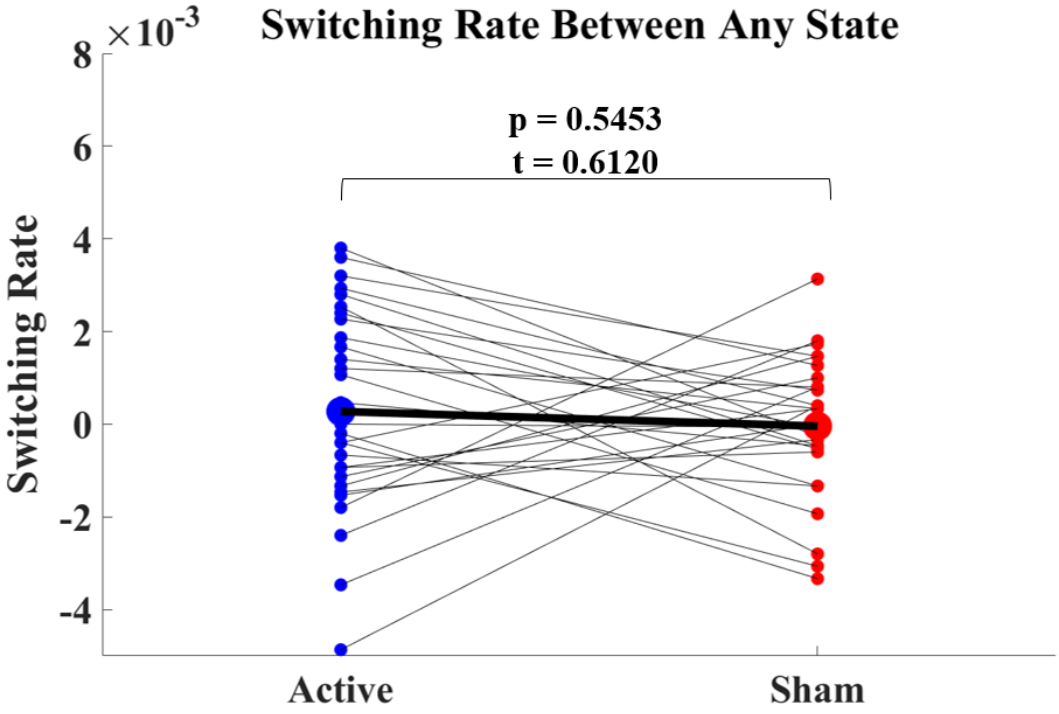

Figure 36: Switching rate for active and sham sessions during PostAr relative to RS0 and disregarding which states subjects transitioned between. A paired t-test was performed to determine if the difference between them was significant.

4.4.3.4 Transition Probability Matrix

Fig. 37 shows all calculated transition probability matrices (TPMs) with the diagonals removed for easier visualization of the off-diagonal elements. The TPM directly outputted from the HMM is not shown as we wanted to examine transition probabilities during PostAr only. Fig. 37a shows the TPM for concatenated active and sham PostAr blocks which was quantitatively similar to the HMM's direct output (Spearman's ρ =

0.8957). When assessing these global transition probabilities calculated by hand, we analyzed values on par with those seen in the direct output. Any differences between the two may have occurred because Fig. 37a was calculated using a count method (counting the number of times subjects switched) and therefore was acquired using a binary analysis. It is possible that HMM's direct output was acquired using a nonbinary analysis where transition probabilities may have been weighted in a similar manner as the spatial activation patterns the model outputted may have been weighted by the posterior probability. Fig. 37a shows transition probabilities on a global level whereas as we were interested in changes in transition probabilities as a result of active handgrip, and therefore interested in the transition probabilities of active and sham sessions individually.

Fig. 37b shows the transition probability matrix for the active session during PostAr since this is the segment of the experimental paradigm where the active squeezing occurs. Fig. 37c shows the sham session counterpart. In addition to the qualitative similarities between the $TPM_{Active\&Sham}$, TPM_{Active}, and TPM_{Sham}, they also exhibit quantitative similarities. The active session TPM (Fig. 37b) had a ρ value of 0.9906 when Spearman correlated with the HMM TPM (Fig. 37a) while the sham session (Fig. 37c) had a ρ value of 0.9070. This indicated that Fig. 37b and Fig. 37c were close enough to Fig. 37a that they could be compared to the global TPM but were not exact replicas since differences occurred between active and sham which were then analyzed and interpreted. These high correlation values also indicated that the transitions between states in either session were not driving the transitions seen on a global level. The transition probability matrices for both sessions showed that subjects had a higher probability of transitioning from S2 into

S4 and from S4 into S2 in the active session over the sham. Conversely, subjects were more likely to transition from S1 into S3, from S3 into S1, from S5 into S4, and from S5 into S4 in the sham session compared to the active. However, because these TPMs were calculated for only PostAr, we could not explore how the transition probabilities changed as a function of the handgrip task since this matrix disregarded RS0, or baseline values. In this case, baseline was important to consider because one subject (or even one session in one subject) may intrinsically have had a higher switching rate than another subject or session, so to understand how the handgrip task affected this measure, we removed baseline effects.

To do this baseline removal, we computed the TPM for PostAr and for RS0 for each subject and for each session, then found the difference between them. Fig. 37d and Fig. 37e show these results for active and sham sessions, respectively. Since these TPMs show the probabilities relative to baseline, we termed them *relative transition probability matrices* (RTPMs). That is, $RTPM_{Active} = TPM_{Active\_PostAr} - TPM_{Active\_RS0}$ and $RTPM_{Sham} = TPM_{Sham\_PostAr} - TPM_{Sham\_RS0}$. To understand how these transition probabilities relative to baseline changed as a function of active squeeze we found the difference between these RTPMs (Fig. 37f). These subtractions were done first on an individual level prior to obtaining the global average. However, these RTPMs could not be directly compared to the HMM TPM direct output since they now displayed changes in transition probabilities rather than raw transition probabilities. Because the magnitudes of the RTPMs were small, we explored only the difference between active and sham and focused on cases where larger differences between the sessions occurred.

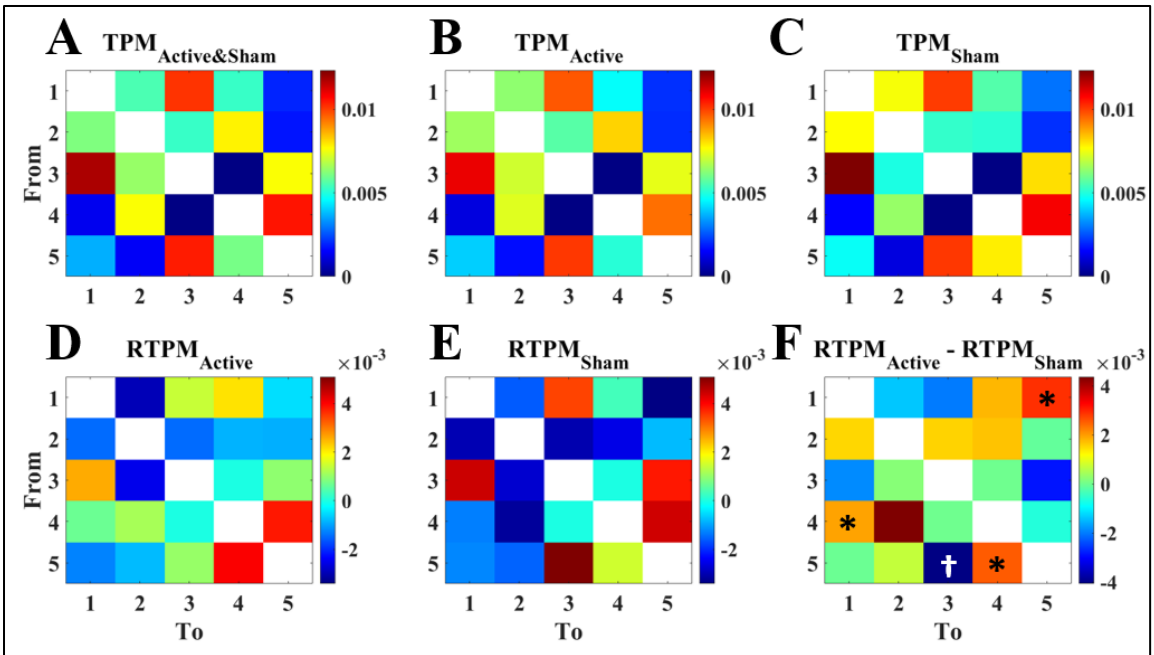

Figure 37: Transition probability matrices (TPM) **(A)** for both active and sham sessions as a global measure, **(B)** for just active and **(C)** for just sham sessions for all subjects during PostAr. To determine how the transition probabilities changed relative to baseline, the TPM for PostAr and RS0 were computed for each session and subtracted (RS0 not shown). These relative transition probabilities (RTPM) are seen for **(D)** active and **(E)** sham session as well as **(F)** the difference between them. Stars (*) show significant p-values and daggers (†) show trending p-values after a Wilcoxon sign rank test was conducted for state-specific transition probabilities.

To determine whether the RTPMs were different from one another overall, the Euclidean distance was found between each individual subject's $RTPM_{Active}$ and $RTPM_{Sham}$. Because this was a nonparametric test applied to a distance metric, there was no need for transformations. The differences in Euclidean distances across all 30 subjects were compared against zero using a Wilcoxon sign rank test and a p-value of 1.7344e-06 resulted. That is, $ED_i = ED(RTPM_{Active,i}, RTPM_{Sham,i})$ where $i \in \{subjects\}$ and a Wilcoxon sign rank test compared the vector ED_i against 0. This indicated that the active squeeze transition probabilities relative to baseline were significantly different from their sham counterpart. However, these analyses did not account for the transitions between specific states. To test whether the difference in transition probabilities of state-specific transitions

relative to baseline were significantly different from zero, a Wilcoxon sign rank test was conducted for each distribution of all subjects' off-diagonal transition probabilities in $RTPM_{Active} - RTPM_{Sham}$ (Fig. 37f). That is, $D_i = RPM_{Active,i} - RTPM_{Sham,i}$ where $i \in \{subjects\}$ and D_i was then compared to 0 using a Wilcoxon sign rank test. The stars (*) and daggers (†) in Fig. 37f show correlations that were, respectively, statistically significant and trending.

This analysis revealed a significantly higher transition probability as subjects switched out of S4 and into S1 during the active session and relative to baseline ($p = 0.0467$; Fig. 37f). Although this may seem counterintuitive as we expected the active handgrip to elicit an LC response and increase arousal, this observed phenomenon may be consistent with the idea that subjects returned to resting state following the handgrip task. One explanation for this phenomenon to occur in the active squeeze session more likely over the sham is NE depletion. It is possible that the arousing stressor (squeezing) depleted the LC's supply of NE causing subjects to return to the DMN-dominant state rather than switching into an attention-dominant state (i.e., S2) thereby prompting an attention "reset" (97–99). The significant level used was uncorrected as no family-wise error rates had been considered. It is possible that the significant transition probability would not survive any corrective measures consequently rendering the value insignificant.

The results in Fig. 37f showed that subjects also exhibited a higher probability relative to baseline of switching from S5 into S4 once the active squeeze was performed ($p = 0.0226$). Although this switch was barely detectable in active PostAr alone, removing the effects of baseline showed that this phenomenon was exacerbated once the active

squeezing blocks were introduced. This same transition from S5 into S4 was not overtly observed in PostAr or PostAr relative to RS0 during the sham session meaning that the likelihood of this transition increased as a function of squeezing. This indicates that subjects were more likely to reside in S5 right before the active squeeze blocks.

Fig. 37f also showed that subjects were more likely to transition into the attention-dominant state (S2) after performing the active squeeze blocks in the active session (S4). Although there was no distinguishable difference in this transition probability between TPM_{Active} (Fig. 37b) and $RTPM_{Active}$ (Fig. 37d), this probability actually decreased between TPM_{Sham} (Fig. 37c) and $RTPM_{Sham}$ (Fig. 37e). This shows that the higher probability observed in $RTPM_{Active} - RTPM_{Sham}$ (Fig. 37f) was due not to this probability remaining the same, but because the active squeeze prevented a decrease in this transition probability, although insignificantly. As a result, we cannot definitively conclude that actively squeezing a squeeze-ball invoked arousal and increased attentiveness as a result of subjects transitioning into an attention-dominant state (S2) in response to active squeezing (36,37,67–71,80). Thus, we cannot conclude that the handgrip task increased salient attentional processing because actively squeezing did not increase the likelihood of subjects entering the DAN/SN-dominant state (S2). Although the difference in this value across condition was large, it was likely insignificant because the variance was large. This high between-subject variability could indicate that subjects reacted differently when transitioning into the DAN/SN-dominant state (S2) after the handgrip task which is potential grounds for future research. Although Mather et al. 2020 found that handgrip increased phasic arousal in both younger and older adults, our results show that we could

not reproduce their findings using our modified handgrip task in younger subjects (37). A similar observation can be made with the probability of transition out of S1 and into S5 ($p = 0.0168$). Importantly, this seemingly large and significant transition probabilities during the active session compared to sham and relative to baseline did not occur because the transition probabilities increased as a result of active squeezing. Instead, the active squeeze prevented a reduction in the transition probability. Subjects transitioning out of the DMN-dominant state and into the whole-brain deactivation state could indicate that a processing system not analyzed in this study is at work in response to the active squeeze. The idea that DMN was activated prior to this could indicate that the unrelated activated system could be associated with resting state in some way.

Negative values in Fig. 37f delineate instances where the transition probability was higher during the sham session over the active. Thus, it appears that the transition probability between S3 and S5 (regardless of direction) relative to baseline was greater during the sham session compared to the active. The transition out of S5 and into S3 showed a trending p-value ($p = 0.0937$). Relative to baseline, the transition probability of S3 $\leftrightarrow$ S5 exhibited minimal changes during PostAr in the active session but increased during the sham session. Thus, the negative values seen in $RTPM_{Active} - RTPM_{Sham}$ occurred because the active squeeze impeded an increase in this transition probability relative to baseline, rather than the active squeeze having no or a reduction effect on this probability. The increase in transition probability between S3 and S5 during the sham session could have occurred because subjects still lifted their arm to their chest which may have marginally up-regulated LC activity slightly inducing arousal (80). The high

probability of transitioning between S3 and S5 during the sham condition is consistent with one seen in previous literature where pure resting state (with no periodic arousal-inducing tasks) was used (11). This indicates that our results were consistent with those seen in literature where subjects were likely to transition between whole-brain activation states and whole-brain deactivation state in the absence of a task (11). It is possible that the transition S5→S3 was significant but S3→S5 was not because the slight changes in LC activity triggered caused subject's likelihood of transitioning out of a whole-brain deactivation state and into a whole-brain activation state where attentional networks were activated to increase rather than the other way around (80).

The interpretations presented above are qualitative as we visually compared the differences in values between relative transition probabilities across conditions based on their color differences. Although the Wilcoxon sign rank test showed a few significant transition probabilities (S1↔S4 and S5→S4) and one almost significant switch (S5→S3), family-wise error rate corrections had not been taken into consideration.

Interestingly, no subject transitioned from S3 into S4 or back in either session. One possible explanation for this observed phenomenon is that LC could be allocating attention. Previous studies have shown that an increase in LC activity may cause a rapid shift in allocation of attention in response to the sudden onset of a stimulus (97,98). This may be why subjects did not transition from the arousal state (S4) into the whole-brain activation state (S3) where all attention-related networks were concurrently activated, and the LC induced some phenomenon that could not be observed in this investigation. Instead transitions from S4 into S1, S2, or S5 were seen where DMN, FPCN, DAN, and SN were

not synchronously activated—only a few (two at most as seen in S2) were activated at the same time. The LC allocated attention to DAN and SN since the subjects were observed to transition out of S4 was into S2. Attention was also allocated to a resting state network (DMN in S1) since some subjects returned to resting state following the squeeze blocks. Some attention could have been allocated to a network not examined (S5), albeit more so in the sham session than in the active, but never to DMN, FPCN, DAN, and SN at the same time. Because this idea of attentional allocation is associated with network resetting, our results may provide preliminary evidence that HMM-derived states can quantify the LC network reset hypothesis (97,98). However, additional tests focusing on attentional shift during these specific transitions are needed.

The correlations between LC MTC and the difference in transition probabilities as a function of the handgrip task relative to baseline are illustrated in Fig. 38. That is, the LC MTC values were Spearman correlated with $RTPM_{Active} - RTPM_{Sham}$ across all subjects prior to obtaining the global average. Stars (*) indicate raw significant ρ values while daggers (†) corresponding to ρ values nearing significance. A significant correlation is seen between LC MTC and the transition S5→S4 ($\rho = 0.0150$) indicating that LC neuromelanin content may affect the probability of subjects transitioning out of the whole-brain deactivation state and into the arousal state as a function of the handgrip task relative to baseline. Although this raw correlation appeared significant or trending, it likely would not survive family wise error corrections.

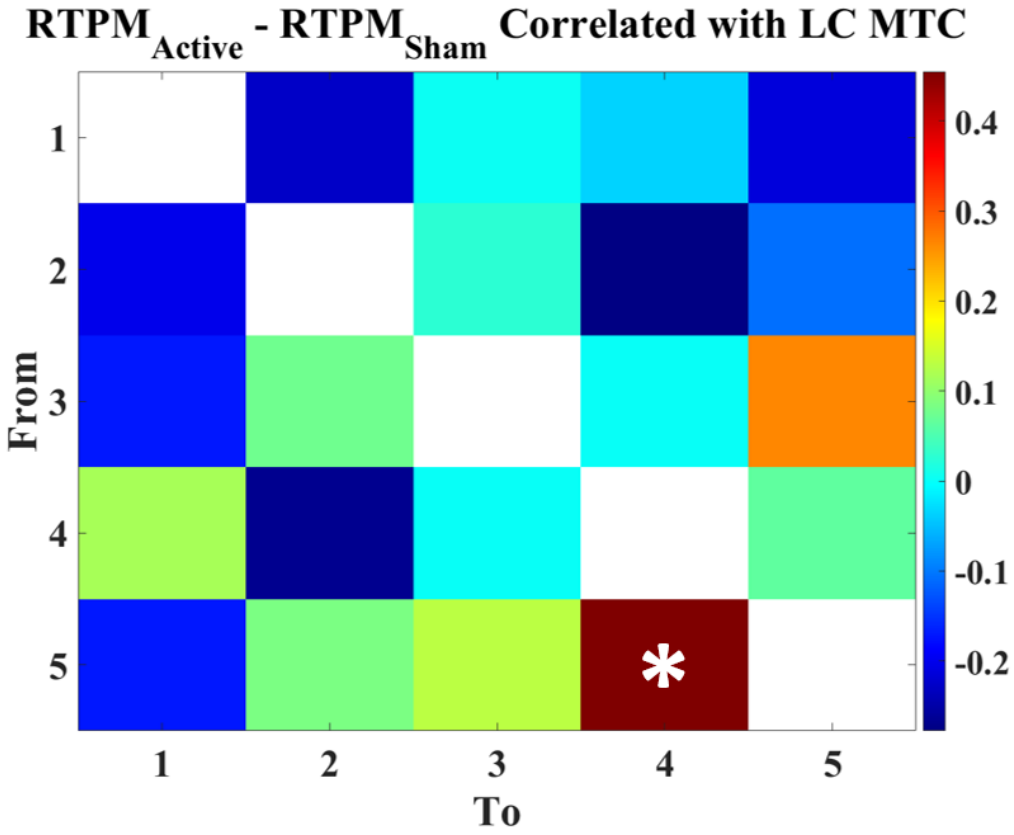

Figure 38: ρ values from Spearman correlating the difference in transition probabilities across condition relative to baseline ($RTPM_{Active}$ – $RTPM_{Sham}$; Fig. 37f) with LC MTC values. "*" corresponds to significant correlation values while "†" corresponds to trending values.

4.4.4 Pupillometry Analyses

The mean pupil dilation when transitioning between any state was first examined, taking care to meet the criterion outlined in Section 4.3.6 (Fig. 39). However, a paired t-test showed no distinguishable differences between the active and sham sessions ($p = 0.4569$, $t = -0.7553$). This might have occurred because we ignored which states subjects transitioned between when computing dilation thereby potentially averaging over important, state-specific dynamics.

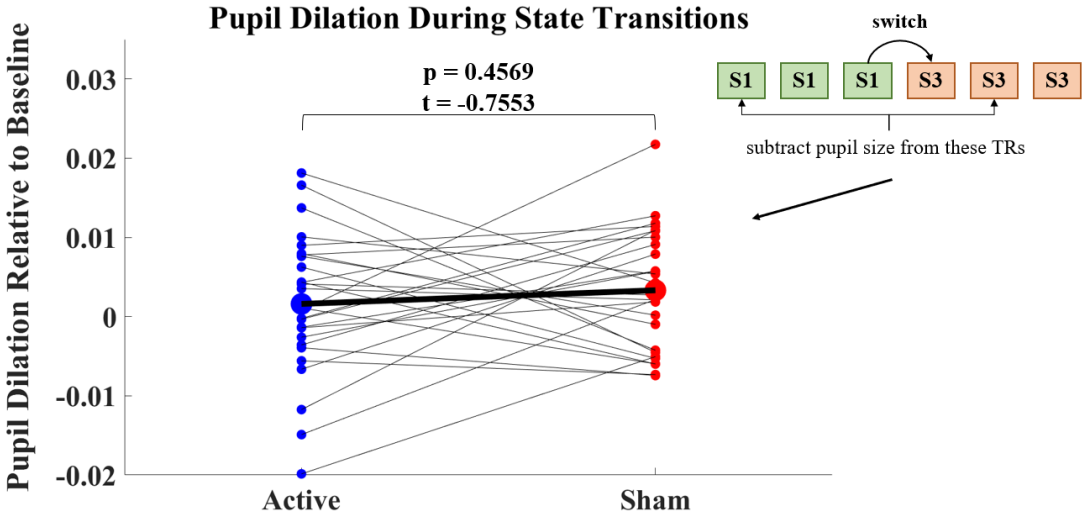

Figure 39: Pupil dilations during transitions between any states for active and sham sessions. The change in pupil size relative was found via the difference in pupil size two TRs before a switch and one TR after a switch when subjects remained in the same state for two TRs before and after the switch. The larger dots and heavier line indicate the mean of each measure. A paired t-test was conducted to determine whether the overall pupil dilation during any transition was significantly difference across conditions.

The pupil dilations were again calculated, but this time the state assignments during the transitions were recorded. The pupil dilations for state-specific transitions were computed for active (Fig. 40a) and sham (Fig. 40b) sessions, as well as the difference between them (Fig. 40c) to highlight distinct pupil size changes resulting from LC activity up-regulation. The difference in pupil dilation for state-specific transitions across conditions (prior to taking the global average across subjects) was Spearman correlated with LC MTC (Fig. 40d). The criterion that subjects must remain in the same state for two TRs before and after the switch was enforced resulting in only a handful of subjects meeting this standard for a transition. That is, the number of subjects used to calculate mean pupil dilation in Fig. 40abc and correlations in Fig. 40d were different for every square. Fig. 40e shows that although some differences in dilations between sessions (i.e., from S4 into S1 or from S3 into S5) or some correlations (i.e., when correlating LC MTC with the transition from S2 into S4 or from S5 into S4) appear to be distinct features, these

correlations were performed for 10-12 subjects causing their statistical power to diminish. Negative results should be interpreted with caution given the amount of data loss.

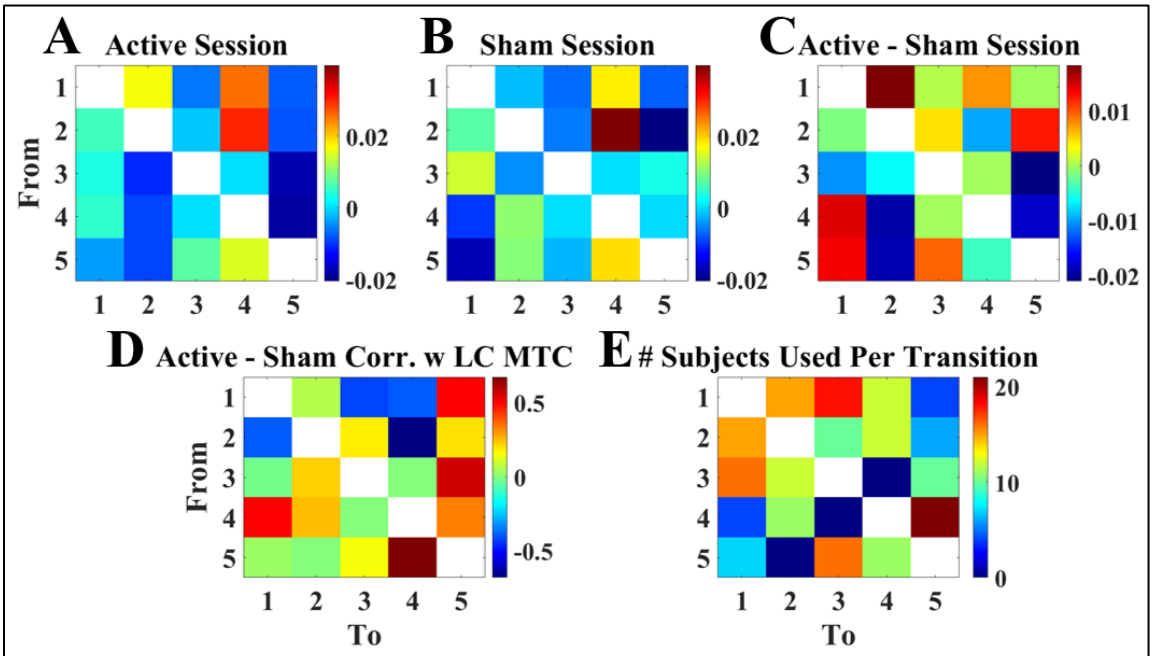

Figure 40: Pupil dilations for **(A)** active and **(B)** sham sessions as well as **(C)** the difference between them when subjects transitioned between specific states. **(D)** The difference in pupil dilation relative to baseline across conditions was Spearman correlated with LC MTC. We computed pupil dilation for a transition only when the subject remained in the same state for two TRs before and after an identified switch. Not all subjects underwent a switch that met this criterion, so **(E)** shows the number of subjects that did per transition.

The exclusion criterion caused an insufficiently small subject pool to prevail from which to obtain information about pupil dilation during state-specific transitions. As a result, negative results should not necessarily be interpreted as a lack of effect, but as a potential lack of statistical power. This limitation can be rectified by increasing the overall scan duration, i.e., employing a one-hour long paradigm rather than one lasting less than 20 minutes, and by collecting data from more subjects. The chances of switching between states would increase and more subjects would survive the rejection standard consequently increasing the effect size. We cannot be less stringent with the criterion of remaining in the same state for two TRs before and after the switch because it ensures that changes in

pupil size were examined for specific transitions, and that dilations/contractions from previous or succeeding transitions do not bleed over into that calculation.

4.5 Conclusion

In this chapter we applied a probabilistic model to a dataset geared towards up-regulating LC activity to examine how latent brain states and related measured changed as a function of performing a handgrip task. The goal of these analyses was to obtain insight into the underlying neural mechanisms driving LC's modulation of arousal. Analyses of a series of different data types—functional, pupillometry, and structural—were related to the results from an HMM to quantity the impact that LC exerts upon attention. Fitting the HMM to the described LC dataset yielded five states consisting of various activation levels of five attention networks. S1 was a DMN-dominant state, S2 was an attention-dominant state, S3 showed all networks to be activated, S4 was an arousal state associated with the squeeze blocks of the experimental paradigm, and S5 showed all networks to be deactivated. Connectivity state patterns showed that in activation states showing certain networks to exhibit high levels of activity also exhibited below baseline correlations, i.e., functional disconnectedness. The outputted hidden state sequence clearly displayed the squeeze blocks because S4 was prevalent almost for the entire block and for almost all subjects.

Examining the noise-to-signal ratio (Fano factor) of all five networks within each state showed DMN and S1 to consist of the most noise, likely because the resting state dataset investigated was disrupted with either an active squeeze or with a prompt to squeeze. S4 also demonstrated a relatively higher Fano factor likely because it mostly

occurred during the brief squeeze blocks. The fractional occupancy of all states as a function of block and regardless of condition indicated that the scanner environment could have induced some arousal thereby causing subjects to reside more in S2 during RS0 than in PostAr. A possible indication of NE depletion resulting from the handgrip task may have been observed. The average duration of the attention-dominant state (S2) once subjects transitioned into it was shorter during PostAr compared to baseline when we expected it to increase due to LC's established modulation of attention. The whole-brain activation state (S3) exhibited higher fractional occupancy and average state duration than the whole-brain deactivation state (S5) likely because the handgrip task up-regulated LC activity which then impacted downstream attention networks. No conclusive insight into how LC drove arousal can be deduced because all three-way omnibus ANOVAs determined the difference in global measures between active and sham sessions to be insignificant. Step-down statistical tests showed that these measures were significantly different during PostAr compared to baseline. This allowed us to conclude that the handgrip task had some effect on Fano factor, fractional occupancy, and average state duration, but the degree to which the active squeeze had an effect cannot be determined from this investigation.

Although the global switching rate relative to baseline demonstrated no significant difference as a function of condition, we hypothesized that the switching rate between certain states would vary in response to active squeeze. We subsequently explored the transition probabilities for active and sham sessions separately, the transition probabilities for both sessions relative to baseline, and the difference in these relative probabilities

across condition. These measures allowed us to pinpoint whether apparent large likelihoods of transitions between sessions occurred as a result of an increase due to active squeezing or via a decrease during sham PostAr relative to baseline. From these results we were neither able to sufficiently establish an increase in arousal nor able to demonstrate an increase in the transition probability into an attention-dominant state (S2) as a result of active squeezing. Thus, using the tools and measures presented in this dissertation, we were not able to adequately show how a handgrip task impacted arousal or attention. Furthermore, analyses of the TPMs and RTPMs showed that subjects were more likely to reside in the whole-brain deactivation state (S5) prior to the active squeeze blocks.

The active squeeze increased the probability of transitioning out of S1 and into S4 relative to baseline, likely because it provoked subjects to transition out of the DMN network and into a state of arousal. The active squeeze also prevented a reduction in the probability of switching from S4 into S1 possibly indicating NE depletion and "resetting" attention (97–99). We also observed that the active squeeze prevented an increase in the probability of transitioning between S3 and S5. This indicates that subjects were more likely to transition between the whole-brain activation (S3) and whole-brain deactivation (S5) states in the absence of the arousing task, a phenomenon also observed by Chen et al. 2016 who utilized a pure resting state data uninterrupted by any task (11). Our results also suggested that a processing system unrelated to attention but related to resting state could have occurred since subjects transitioned from S1 into S5 in response to the squeeze.

Potential evidence of attentional allocation—an idea associated with the LC reset hypothesis—may have been observed as no subject transitioned between S3 and S4

(97,98). This might be why, in response to the squeeze, subjects were more likely to transition from the arousal state (S4) into S2 (where DAN and SN were activated), into S5 (where none of the networks investigated here were activated, but another attention-related network could be), or even into S1 (where a resting state network was activated) rather than into S3 where all attentional networks were concurrently activated and another LC-induced phenomenon not observed by the tools used in this investigation may have occurred.

Exploring the pupillometry data showed that accounting for subjects settling into a state to prevent contamination of pupil size changes between state transitions excluded many subjects. This greatly reduced the statistical power of our data and as a result we could not glean any useful information about the relationship between pupil dilation (an indirect measure of LC engagement) and transitions between HMM-derived brain states. Although examined in Chapters 2 and 3, fractional occupancy correlation was not pursued in this chapter because a 5-state model was chosen which contained too few states from which to identify potential metastates. Vidaurre et al. 2017 grouped six states per metastate (which are more than the total number of states identified from this dataset), Stevner et al. 2019 at least three states, and in Chapter 3, four states were grouped within each identified metastate (12,13). No major conclusions could have been made about state clustering, and thus, no additional information would have been gleaned from the fractional occupancy correlation that could not have been understood from the other measures described (i.e., transition probabilities or average duration). As a result, we decided not to investigate this measure, but could do so if a longer dataset is obtained.

All HMM-related measures explored in this investigation (Fano factor, fractional occupancy, average state duration, general switching rate, specific switching rate, and pupil dilations during state-specific transitions) were Spearman correlated with LC MTC to determine if they had a relationship with the amount of neuromelanin in LC. None of the correlation values found were likely to reach significance after multiple comparison corrections indicating that we potentially could not recognize a relationship between LC neuromelanin content and measures based on HMM outputs. This may have occurred because the subject pool in this dataset contained only healthy young adults where neuronal loss is expected to be minimal if occurring at all (72–74). Thus, we would not expect to see a relationship between LC structure and changes in HMM-based measures because our subject pool's neuromelanin content has not been compromised due to age or to disease which we screened for prior to data collection (72–74). It is possible that we would see significant results in older adults or in a diseased population (i.e., subjects with Parkinson's or Alzheimer's diseases) where LC neuromelanin content is more variable (10,100). This limitation also applied to all HMM-related measures explored in this investigation: little to no insight about LC's relationship with attention could be obtained likely because we examined a population where LC remained intact and where its function had not been compromised.

Another reason why no difference in measures were found between active and sham is because the handgrip task may not have elicited a strong enough LC response. Ideally, our measures would have related back to the Yerkes-Dodson curve to track subjects' progression along it as their LC was engaged via squeezing (29). However, this could not

be done using the measures presented in this investigation because three repeated measures ANOVAs and two paired t-tests showed that the measures analyzed could not be distinguished between active and sham sessions. To rectify this, a stronger external stimulus could be used that potentially evokes an unmistakable LC response. Alternative stimuli include subjects dipping their hand in cold water, administering electrical pulses, or presenting jarring sounds (101–105). On the other hand, if the squeeze did induce a recognizable LC response, it is possible that the measures used in this investigation were not sensitive enough to detect it. This may have occurred because the BOLD signal might not have been strong enough for the HMM to detect all possible changes in the signal. Future investigations could focus on improving the signal-to-noise ratio of the networks via denoising to obtain reliable BOLD signals and pupillometry data (106,107). LC BOLD is subject to physiological noise, specifically cardiac pulsation and respiration (108–111). These sources of noise were not regressed out in this investigation because physiological data were not collected in conjunction with functional data. In the future, cardiac pulsation and respiration data should be obtained in order to regress out physiological noise from LC BOLD prior to HMM fitting (109). This would consequently help LC to have more of a presence in the activation and connectivity state pattens which would help us better understand how it interacts with the other networks within each state. These corrective measures would also increase LC and network signal which would inadvertently allow the HMM to be more sensitive to BOLD fluctuations as active squeezing.

Another limitation of this study is the number of subjects, especially in the pupillometry analyses. A relatively small N was used because many subjects were

excluded as they did not meet the preset criterion of remaining in the same state two TRs before or after a switch. Increasing N would also increase the statistical power of our pupillometry analyses potentially allowing us to find a relationship between pupil dilation and attention-related latent brain states. Furthermore, although the dataset had been optimized to obtain subcortical signals, it had a relatively short scan duration and a relatively slow TR resulting in relatively poor temporal resolution. The HCP dataset explored in Chapters 2 and 3 had more than twice the total scan duration and less than half the TR. Because of this, we were even able to capture temporal dynamics when employing a probabilistic model based on a windowed analysis (which greatly reduced temporal resolution). Increasing the length of the scan and utilizing a faster TR would make the data more sensitive to fluctuations in BOLD. This potentially increases the effect size of HMM-related measures that rely on temporal resolution (i.e., average stat duration and switching rate) which were found not to be significant when using the data acquisition parameters presented in this chapter.

Overall, the analyses presented in this chapter could potentially provide useful quantitative insight into the dynamic relationship between LC and arousal if the many limitations discussed are elucidated.

Chapter 5: Conclusion

The work presented in this dissertation aimed to understand different subtypes of hidden Markov models (HMMs) and to apply this technique to a pseudo-resting state dataset in effort to analyze various characteristics of attention-related brain states. We introduced three types of HMMs: an activation-based HMM (AB HMM), a summed functional connectivity HMM (SFC HMM) taken from Ou et al. 2014, and a full functional connectivity HMM (FFC HMM) (11–13,28). All three model types, their direct outputs, and HMM-based measures subsequently calculated were theoretically compared and contrasted. Local and global analyses of AB, SFC, and FFC HMMs were analyzed empirically when applied to a resting state dataset. The results allowed us to infer which model type would be best to use for certain investigations based on temporal resolution, spatial patterns, and transitional dynamics. AB HMM was then applied to a dataset optimized to image the locus coeruleus (LC) in effort to obtain clearer insight into LC's dynamic relationship with attention and arousal.

In Chapter 2, we noted SFC and FFC HMM to have a poorer temporal resolution than AB HMM because they utilized a windowed analysis which severely smoothed data over time, restricted the information fitted by the HMM to strictly range between -1 and 1, and included the effects of autocorrelation. SFC HMM's temporal resolution was reduced even further because summed correlations within each TR were inputted into the model rather than the raw R^2 values. Although all three models could produce activation and connectivity state spatial patterns, not all were feasible for interpretations, nor did they exhibit similarities. The connectivity-based states generated from AB HMM were

expected not to display the same patterns as those derived from SFC and FFC HMMs because the HMM was fitted to varying patterns of the blood oxygen level dependent (BOLD) signal. The activation spatial patterns acquired from averaging the BOLD signal where SFC and FFC HMM's Viterbi path labeled a state to be active was not analogous to AB HMM's activation state patterns. The Viterbi path from the connectivity-based model types were based on changes in connectivity patterns (dynamic functional connectivity), and therefore poorly replicated activation patterns. Published literature discussed that an ideal window size for a sliding window analysis should include less than 60 seconds of data per window, so our analyses accordingly used ~40 seconds of data (14). The two methods of model order determination described in this chapter (and implemented in Chapter 3) could be applied to all three HMM subtypes.

Chapter 3's focus was on applying all three HMM subtypes to a resting state dataset from the Human Connectome Project (HCP). Both stability analysis methods of model order determination showed that eight states were best for this investigation. Although we determined an 8-state model should be used for AB, SFC, and FFC HMMs, we crucially noted that identifying the same model order for all three HMM subtypes might not occur for all investigations. Local analyses (characterizing state patterns from different methods of acquisition) and global analyses (measures that can be extracted from all HMM outputs) allowed us to make empirically based informed decisions regarding which model type is best pursued under certain circumstances. AB HMM satisfactorily captured temporal dynamics, likely because it had the leading temporal resolution. This model type also allows users to explore activation and connectivity state patterns in conjunction. However,

it was able to identify only weak connectivity patterns, so if the focus of an investigation is to examine connectivity spatial patterns, then SFC HMM or FFC HMM would be preferred. SFC HMM and FFC HMM produced connectivity states based on the same dynamic functional connectivity sliding window analysis with slight variations. We determined that these discrepancies did not drastically change the spatial patterns but did generate enough differences that users must take it into consideration when choosing between the two. SFC HMM is preferred if an investigation concentrates on examining general nodal connections since its inputs were vectors representing the overall connectivity strength between each node and all others. Conversely, FFC HMM is advantageous if a study were to focus on specific nodal connections since this model type was fitted to every correlation value between a node and all other nodes stipulated. We found that FFC HMM severely smoothed the data over time, so it was not as satisfactory as AB HMM, or even SFC HMM, in capturing temporal dynamics. It is possible that eight may have been an unfavorable model order for this particular subtype, which must also be taken into consideration when comparing each HMM's temporal resolution. Furthermore, each subtype was distinct in that they recognized changes only in the data type inputted. We observed that switching rate could not be meaningfully interpreted across subtypes because of the acute smoothing factor but could be evaluated within a model type—particularly in AB HMM since switching rate relies on temporal dynamics.

AB HMM was fit to a pseudo-resting state dataset where LC activity was up-regulated via a handgrip task in Chapter 4 (36,37). The goal of these analyses was to understand how HMM-related measures changed as a function of performing a squeeze

task thereby providing us with insight into LC's dynamic relationship with attention and arousal. A 5-state model was preferred for this study. State 1 (S1) was a default-mode network (DMN) dominant state, state 2 (S2) was an attention-dominant state because dorsal attention network (DAN) and salience network (SN) were activated, state 3 (S3) was a whole-brain activation state because all networks (DMN, fronto-parietal control network (FPCN), DAN, SN, and LC) were activated, state 4 (S4) was the arousal state prevalent during the squeeze blocks of the paradigm, and state 5 (S5) was a whole-brain deactivation state. We performed the local and global analyses defined in previous chapters and examined pupil data's relationship with the HMM outputs since changes in pupil size have correlated with LC activity (76–78). Only a few statistically significant results were procured. We found that subjects generally reside more in S2 prior to the introduction of squeeze blocks, which could indicate norepinephrine (NE) depletion as a result of the handgrip task (37,94–96). We also observed that subjects never transitioned between S3 and S4. Because all attention-related networks were activated in S3, this could allude to the hypothesis that LC allocates attention which may be modulated by NE (97,98). In response to the handgrip task subjects transitioned between S4 and S1 or S2 where some attention networks were activated, but never to S3 where they were concurrently activated. Thus, it is possible that the LC allocated attention as subjects transitioned into S1 or S2 where DMN and DAN/SN were respectively active, but induced a different phenomenon not captured in this investigated when subjects transitioned into S3. Additional tests focusing on neuroendocrinological factors including NE are needed to verify these ideas.

Importantly, we found that none of the measures analyzed were statistically significantly different between an active squeeze session and a control session. Thus, no information regarding how HMM-related measures changed as a function of active squeeze could be gleaned, only how they changed relative to baseline. Although there is literature to support the idea of active squeezing engaging the LC, it may not have evoked a strong enough LC response for the design used in this study (36,37). In the future, a more stressful stimulus could be utilized such as cold pressors, jarring sounds, or electrical pulses (101–105). It is possible that informative insight can be made about LC's dynamic modulation of attention and arousal when using HMM-related measures if the limitations discussed are amended.

All measure analyzed were correlated with LC magnetization transfer contrast (MTC) to determine whether a relationship between LC structure and changes in HMM-related measures as a function of active squeeze could be established. However, none of the correlations were definitively, statistically significant. It is possible that no significant relationship between LC MTC and Fano factor, fractional occupancy, or average state duration occurred because the subject pool of this dataset consisted of only young healthy adults where LC neuronal loss is minimal if occurring at all (72–74). Similarly, we could not make informative observations about active squeezing's impact on HMM-related parameters when analyzing the pupillometry data. A stringent criterion was imposed where subjects must have remained in the same state two TRs before and after a switch between states. This excluded many subjects which greatly reduced the effect size and statistical

power of computations. As a result, no informative conclusions could be drawn about LC's impact on arousal when inspecting LC MTC or pupil dilations.

Altogether this dissertation explored the theoretical foundation for different HMM subtypes, empirically assessed which model type is best used under certain circumstances, and utilized HMM-related measures in an attempt to characterize LC's dynamic relationship with arousal.

References

1. Raichle ME. The Restless Brain. Brain Connectivity. 2011 Jan;1(1):3–12.

2. Berger H. Über das Elektrenkephalogramm des Menschen: XIV. Mitteilung. Archiv f Psychiatrie. 1938 Jun;108(3):407–31.

3. Zylberberg A, Roelfsema PR, Sigman M. Variance misperception explains illusions of confidence in simple perceptual decisions. Consciousness and Cognition. 2014 Jul;27:246–53.

4. McCurdy LY, Maniscalco B, Metcalfe J, Liu KY, de Lange FP, Lau H. Anatomical Coupling between Distinct Metacognitive Systems for Memory and Visual Perception. Journal of Neuroscience. 2013 Jan 30;33(5):1897–906.

5. Yu AJ, Dayan P. Uncertainty, neuromodulation, and attention. Neuron. 2005 May 19;46(4):681–92.

6. Posner MI, Petersen SE. The Attention System of the Human Brain. Annu Rev Neurosci. 1990 Mar;13(1):25–42.

7. Langley J, Huddleston DE, Liu CJ, Hu X. Reproducibility of locus coeruleus and substantia nigra imaging with neuromelanin sensitive MRI. Magn Reson Mater Phy. 2017 Apr;30(2):121–5.

8. Chen X, Huddleston DE, Langley J, Ahn S, Barnum CJ, Factor SA, et al. Simultaneous imaging of locus coeruleus and substantia nigra with a quantitative neuromelanin MRI approach. Magnetic Resonance Imaging. 2014 Dec;32(10):1301–6.

9. Langley J, Hussain S, Flores JJ, Bennett IJ, Hu X. Characterization of age-related microstructural changes in locus coeruleus and substantia nigra pars compacta. Neurobiol Aging. 2020 Mar;87:89–97.

10. Langley J, Hussain S, Huddleston DE, Bennett I, Hu XP. Impact of locus coeruleus and its projections on memory and aging. Brain Connect. 2021 Jun 17;

11. Chen S, Langley J, Chen X, Hu X. Spatiotemporal Modeling of Brain Dynamics Using Resting-State Functional Magnetic Resonance Imaging with Gaussian Hidden Markov Model. Brain Connect. 2016 May;6(4):326–34.

12. Vidaurre D, Smith SM, Woolrich MW. Brain network dynamics are hierarchically organized in time. Proc Natl Acad Sci USA. 2017 Nov 28;114(48):12827–32.

13. Stevner ABA, Vidaurre D, Cabral J, Rapuano K, Nielsen SFV, Tagliazucchi E, et al. Discovery of key whole-brain transitions and dynamics during human wakefulness and non-REM sleep. Nat Commun. 2019 Mar 4;10(1):1035.

14. Lurie DJ, Kessler D, Bassett DS, Betzel RF, Breakspear M, Kheilholz S, et al. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. Network Neuroscience. 2020 Jan;4(1):30–69.

15. Liu X, Chang C, Duyn JH. Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. Front Syst Neurosci [Internet]. 2013 [cited 2021 Jun 26];7. Available from: http://journal.frontiersin.org/article/10.3389/fnsys.2013.00101/abstract

16. Liu X, Duyn JH. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proc Natl Acad Sci USA. 2013 Mar 12;110(11):4392–7.

17. Beckmann CF, DeLuca M, Devlin JT, Smith SM. Investigations into resting-state connectivity using independent component analysis. Phil Trans R Soc B. 2005 May 29;360(1457):1001–13.

18. Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, et al. Temporally-independent functional modes of spontaneous brain activity. Proceedings of the National Academy of Sciences. 2012 Feb 21;109(8):3131–6.

19. Schlösser R, Gesierich T, Kaufmann B, Vucurevic G, Hunsche S, Gawehn J, et al. Altered effective connectivity during working memory performance in schizophrenia: a study with fMRI and structural equation modeling. NeuroImage. 2003 Jul;19(3):751–63.

20. Petridou N, Gaudes CC, Dryden IL, Francis ST, Gowland PA. Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. Hum Brain Mapp. 2013 Jun;34(6):1319–29.

21. Sakoğlu Ü, Pearlson GD, Kiehl KA, Wang YM, Michael AM, Calhoun VD. A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. Magn Reson Mater Phy. 2010 Dec;23(5–6):351–66.

22. Chang C, Glover GH. Time–frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage. 2010 Mar;50(1):81–98.

23. Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. Proc IEEE. 1989 Feb;77(2):257–86.

24. Rabiner L, Juang B. An introduction to hidden Markov models. IEEE ASSP Mag. 1986;3(1):4–16.

25. Eddy SR. Hidden Markov models. Curr Opin Struct Biol. 1996 Jun;6(3):361–5.

26. Eddy SR. What is a hidden Markov model? Nat Biotechnol. 2004 Oct;22(10):1315–6.

27. Vidaurre D, Abeysuriya R, Becker R, Quinn AJ, Alfaro-Almagro F, Smith SM, et al. Discovering dynamic brain networks from big data in rest and task. NeuroImage. 2018 Oct;180:646–56.

28. Ou J, Xie L, Jin C, Li X, Zhu D, Jiang R, et al. Characterizing and Differentiating Brain State Dynamics via Hidden Markov Models. Brain Topogr. 2015 Sep;28(5):666–79.

29. Aston-Jones G, Cohen JD. AN INTEGRATIVE THEORY OF LOCUS COERULEUS-NOREPINEPHRINE FUNCTION: Adaptive Gain and Optimal Performance. Annu Rev Neurosci. 2005 Jul 21;28(1):403–50.

30. Song AH, Kucyi A, Napadow V, Brown EN, Loggia ML, Akeju O. Pharmacological Modulation of Noradrenergic Arousal Circuitry Disrupts Functional Connectivity of the Locus Ceruleus in Humans. J Neurosci. 2017 Jul 19;37(29):6938–45.

31. Sara SJ. The locus coeruleus and noradrenergic modulation of cognition. Nat Rev Neurosci. 2009 Mar;10(3):211–23.

32. Guedj C, Meunier D, Meunier M, Hadj-Bouziane F. Could LC-NE-Dependent Adjustment of Neural Gain Drive Functional Brain Network Reorganization? Neural Plasticity. 2017;2017:1–12.

33. Braak H, Thal DR, Ghebremedhin E, Del Tredici K. Stages of the Pathologic Process in Alzheimer Disease: Age Categories From 1 to 100 Years. J Neuropathol Exp Neurol. 2011 Nov;70(11):960–9.

34. Braak H, Tredici KD, Rüb U, de Vos RAI, Jansen Steur ENH, Braak E. Staging of brain pathology related to sporadic Parkinson's disease. Neurobiology of Aging. 2003 Mar;24(2):197–211.

35. Yang Z, LaConte S, Weng X, Hu X. Ranking and averaging independent component analysis by reproducibility (RAICAR). Hum Brain Mapp. 2008 Jun;29(6):711–25.

36. Nielsen SE, Mather M. Comparison of two isometric handgrip protocols on sympathetic arousal in women. Physiology & Behavior. 2015 Apr;142:5–13.

37. Mather M, Huang R, Clewett D, Nielsen SE, Velasco R, Tu K, et al. Isometric exercise facilitates attention to salient events in women via the noradrenergic system. NeuroImage. 2020 Apr;210:116560.

38. Muehlhan M, Lueken U, Siegert J, Wittchen H-U, Smolka MN, Kirschbaum C. Enhanced Sympathetic Arousal in Response to fMRI Scanning Correlates with Task Induced Activations and Deactivations. Chialvo DR, editor. PLoS ONE. 2013 Aug 15;8(8):e72576.

39. Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD. Tracking Whole-Brain Connectivity Dynamics in the Resting State. Cerebral Cortex. 2014 Mar;24(3):663–76.

40. Leonardi N, Van De Ville D. On spurious and real fluctuations of dynamic functional connectivity during rest. NeuroImage. 2015 Jan;104:430–6.

41. Vergara VM, Mayer AR, Damaraju E, Calhoun VD. The effect of preprocessing in dynamic functional network connectivity used to classify mild traumatic brain injury. Brain Behav [Internet]. 2017 Oct [cited 2021 Jun 28];7(10). Available from: https://onlinelibrary.wiley.com/doi/10.1002/brb3.809

42. Jurafsky D, Martin JH. Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition. 2nd ed. Upper Saddle River, N.J: Pearson Prentice Hall; 2009. 988 p. (Prentice Hall series in artificial intelligence).

43. Khiatani D, Ghose U. Weather forecasting using Hidden Markov Model. In: 2017 International Conference on Computing and Communication Technologies for Smart Nation (IC3TSN) [Internet]. Gurgaon: IEEE; 2017 [cited 2021 Jun 28]. p. 220–5. Available from: http://ieeexplore.ieee.org/document/8284480/

44. Zhang M, Jiang X, Fang Z, Zeng Y, Xu K. High-order Hidden Markov Model for trend prediction in financial time series. Physica A: Statistical Mechanics and its Applications. 2019 Mar;517:1–12.

45. Vidaurre D, Hunt LT, Quinn AJ, Hunt BAE, Brookes MJ, Nobre AC, et al. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. Nat Commun. 2018 Dec;9(1):2987.

46. Vidaurre D, Quinn AJ, Baker AP, Dupret D, Tejero-Cantero A, Woolrich MW. Spectrally resolved fast transient brain states in electrophysiological data. NeuroImage. 2016 Feb;126:81–95.

47. Vidaurre D, Hunt LT, Quinn AJ, Hunt BAE, Brookes MJ, Nobre AC, et al. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. Nat Commun. 2018 Dec;9(1):2987.

48. Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Probert Smith PJ, et al. Fast transient networks in spontaneous human brain activity. eLife. 2014 Mar 25;3:e01867.

49. Eavani H, Satterthwaite TD, Gur RE, Gur RC, Davatzikos C. Unsupervised Learning of Functional Network Dynamics in Resting State fMRI. In: Gee JC, Joshi S, Pohl KM, Wells WM, Zöllei L, editors. Information Processing in Medical Imaging [Internet]. Berlin, Heidelberg: Springer Berlin Heidelberg; 2013 [cited 2021 Jun 28]. p. 426–37. (Hutchison D, Kanade T, Kittler J, Kleinberg JM, Mattern F, Mitchell JC, et al., editors. Lecture Notes in Computer Science; vol. 7917). Available from: http://link.springer.com/10.1007/978-3-642-38868-2_36

50. Shappell H, Caffo BS, Pekar JJ, Lindquist MA. Improved state change estimation in dynamic functional connectivity using hidden semi-Markov models. NeuroImage. 2019 May;191:243–57.

51. Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Probert Smith PJ, et al. Fast transient networks in spontaneous human brain activity. eLife. 2014 Mar 25;3:e01867.

52. Lindquist MA, Waugh C, Wager TD. Modeling state-related fMRI activity using change-point theory. Neuroimage. 2007 Apr 15;35(3):1125–41.

53. Liu W, Awate SP, Anderson JS, Fletcher PT. A functional network estimation method of resting-state fMRI using a hierarchical Markov random field. Neuroimage. 2014 Oct 15;100:520–34.

54. Robinson LF, Wager TD, Lindquist MA. Change point estimation in multi-subject fMRI studies. Neuroimage. 2010 Jan 15;49(2):1581–92.

55. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research. 2011;12:2825–30.

56. Guggenmos M, Sterzer P, Cichy RM. Multivariate pattern analysis for MEG: A comparison of dissimilarity measures. NeuroImage. 2018 Jun;173:434–47.

57. Goucher-Lambert K, McComb C. Using Hidden Markov Models to Uncover Underlying States in Neuroimaging Data for a Design Ideation Task. Proc Int Conf Eng Des. 2019 Jul;1(1):1873–82.

58. Van Essen DC, Smith SM, Barch DM, Behrens TEJ, Yacoub E, Ugurbil K. The WU-Minn Human Connectome Project: An overview. NeuroImage. 2013 Oct;80:62–79.

59. Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, et al. The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage. 2013 Oct;80:105–24.

60. Deshpande G, Santhanam P, Hu X. Instantaneous and causal connectivity in resting state brain networks derived from functional MRI data. NeuroImage. 2011 Jan;54(2):1043–52.

61. Lancaster JL, Tordesillas-Gutiérrez D, Martinez M, Salinas F, Evans A, Zilles K, et al. Bias between MNI and Talairach coordinates analyzed using the ICBM-152 brain template. Hum Brain Mapp. 2007 Nov;28(11):1194–205.

62. Laird AR, Lancaster JL, Fox PT. BrainMap: The Social Evolution of a Human Brain Mapping Database. NI. 2005;3(1):065–78.

63. Brett M, Johnsrude IS, Owen AM. The problem of functional localization in the human brain. Nat Rev Neurosci. 2002 Mar;3(3):243–9.

64. Deshpande G, Hu X, Stilla R, Sathian K. Effective connectivity during haptic perception: A study using Granger causality analysis of functional magnetic resonance imaging data. NeuroImage. 2008 May;40(4):1807–14.

65. Deshpande G, LaConte S, James GA, Peltier S, Hu X. Multivariate Granger causality analysis of fMRI data. Hum Brain Mapp. 2009 Apr;30(4):1361–73.

66. Stilla R, Deshpande G, LaConte S, Hu X, Sathian K. Posteromedial Parietal Cortical Activity and Inputs Predict Tactile Spatial Acuity. Journal of Neuroscience. 2007 Oct 10;27(41):11091–102.

67. Nielsen SE, Barber SJ, Chai A, Clewett DV, Mather M. Sympathetic arousal increases a negative memory bias in young women with low sex hormone levels. Psychoneuroendocrinology. 2015 Dec;62:96–106.

68. Lake CR, Ziegler MG, Kopin IJ. Use of plasma norepinephrine for evaluation of sympathetic neuronal function in man. Life Sciences. 1976 Jun;18(11):1315–25.

69. Wallin BG, Mörlin C, Hjemdahl P. Muscle sympathetic activity and venous plasma noradrenaline concentrations during static exercise in normotensive and hypertensive subjects. Acta Physiologica Scandinavica. 1987 Dec;129(4):489–97.

70. Wallin BG, Esler M, Dorward P, Eisenhofer G, Ferrier C, Westerman R, et al. Simultaneous measurements of cardiac noradrenaline spillover and sympathetic outflow to skeletal muscle in humans. The Journal of Physiology. 1992 Jul 1;453(1):45–58.

71. S Kozlowski, Brezezinska Z, Nazar K, Kowalski W, Franczyk M. Plasma Catecholamines During Sustained Isometric Exercise. Clinical Science and Molecular MEdicine. 1973 Jun 15;(45):723–31.

72. Manaye KF, McIntire DD, Mann DMA, German DC. Locus coeruleus cell loss in the aging human brain: A non-random process. J Comp Neurol. 1995 Jul 17;358(1):79–87.

73. Zucca FA, Bellei C, Giannelli S, Terreni MR, Gallorini M, Rizzio E, et al. Neuromelanin and iron in human locus coeruleus and substantia nigra during aging: consequences for neuronal vulnerability. J Neural Transm. 2006 Jun;113(6):757–67.

74. Shibata E, Sasaki M, Tohyama K, Kanbara Y, Otsuka K, Ehara S, et al. Age-related Changes in Locus Ceruleus on Neuromelanin Magnetic Resonance Imaging at 3 Tesla. MRMS. 2006;5(4):197–200.

75. Mather M, Harley CW. The Locus Coeruleus: Essential for Maintaining Cognitive Function and the Aging Brain. Trends in Cognitive Sciences. 2016 Mar;20(3):214–26.

76. Murphy PR, O'Connell RG, O'Sullivan M, Robertson IH, Balsters JH. Pupil diameter covaries with BOLD activity in human locus coeruleus. Hum Brain Mapp. 2014 Aug;35(8):4140–54.

77. Joshi S, Li Y, Kalwani RM, Gold JI. Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. Neuron. 2016 Jan;89(1):221–34.

78. Keren NI, Lozar CT, Harris KC, Morgan PS, Eckert MA. In vivo mapping of the human locus coeruleus. NeuroImage. 2009 Oct;47(4):1261–7.

79. Vecht RJ, Graham GW, Sever PS. Plasma noradrenaline concentrations during isometric exercise. Heart. 1978 Nov 1;40(11):1216–20.

80. Hussain S, Alizadeh Shalchy M, Yaghoubi KC, Langley J, Chen X, Bennett IJ, et al. Locus Coeruleus Engagment Drives Network Connectivity Dynamics in Humans and Rats. In Berlin, Germany; 2019.

81. Hussain S, Menchaca I, Alizadeh Shalchy M, Yaghoubi KC, Langley J, Seitz AR, et al. Locus coeruleus is associated with brain state switching. In 2021.

82. Yaghoubi KC, Alizadeh Shalchy M, Hussain S, Chen X, Bennett IJ, Mather M, et al. Computational fMRI Reveals Separable Representations of Simulus and Choice in Auditory Cortex: A Tool for Studying the Locus Coeruleus Circuit. In Berlin, Germany; 2019.

83. Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, et al. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage. 2004;23 Suppl 1:S208-219.

84. Woolrich MW, Jbabdi S, Patenaude B, Chappell M, Makni S, Behrens T, et al. Bayesian analysis of neuroimaging data in FSL. Neuroimage. 2009 Mar;45(1 Suppl):S173-186.

85. Langley J, Huddleston DE, Chen X, Sedlacik J, Zachariah N, Hu X. A multicontrast approach for comprehensive imaging of substantia nigra. Neuroimage. 2015 May 15;112:7–13.

86. Greicius MD, Menon V. Default-Mode Activity during a Passive Sensory Task: Uncoupled from Deactivation but Impacting Activation. Journal of Cognitive Neuroscience. 2004 Nov;16(9):1484–92.

87. Dixon ML, De La Vega A, Mills C, Andrews-Hanna J, Spreng RN, Cole MW, et al. Heterogeneity within the frontoparietal control network and its relationship to the default and dorsal attention networks. Proc Natl Acad Sci USA. 2018 Feb 13;115(7):E1598–607.

88. Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. Brain Struct Funct. 2010 Jun;214(5–6):655–67.

89. Fano U. Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions. Phys Rev. 1947 Jul 1;72(1):26–9.

90. Deco G, Hugues E. Neural Network Mechanisms Underlying Stimulus Driven Variability Reduction. Behrens T, editor. PLoS Comput Biol. 2012 Mar 29;8(3):e1002395.

91. Hahn G, Zamora-López G, Uhrig L, Tagliazucchi E, Laufs H, Mantini D, et al. Signature of consciousness in brain-wide synchronization patterns of monkey and human fMRI signals. NeuroImage. 2021 Feb;226:117470.

92. Transformations [Internet]. Available from: https://www.ndsu.edu/faculty/horsley/Transfrm.pdf

93. Barkhof F, Haller S, Rombouts SARB. Resting-State Functional MR Imaging: A New Window to the Brain. Radiology. 2014 Jul;272(1):29–49.

94. Nakagawa R, Tanaka M, Kohno Y, Noda Y, Nagasaki N. Regional responses of rat brain noradrenergic neurones to acute intense stress. Pharmacology Biochemistry and Behavior. 1981 May;14(5):729–32.

95. Shinba T, Ozawa N, Yoshii M, Yamamoto K. Delayed Increase of Brain Noradrenaline After Acute Footshock Stress in Rats. Neurochem Res. 2010 Mar;35(3):412–7.

96. Britton KT, S. Segal D, Kuczenski R, Hauger R. Dissociation between in vivo hippocampal norepinephrine response and behavioral/neuroendocrine responses to noise stress in rats. Brain Research. 1992 Mar;574(1–2):125–30.

97. Sara SJ. Locus coeruleus reports changes in enviornmental contingencies.

98. Sara SJ. Locus Coeruleus in time with the making of memories. Current Opinion in Neurobiology. 2015 Dec;35:87–94.

99. Mather M, Clewett D, Sakaki M, Harley CW. GANEing traction: The broad applicability of NE hotspots to diverse cognitive and arousal phenomena. Behav Brain Sci. 2016;39:e228.

100. Sulzer D, Cassidy C, Horga G, Kang UJ, Fahn S, Casella L, et al. Neuromelanin detection by magnetic resonance imaging (MRI) and its promise as a biomarker for Parkinson's disease. npj Parkinson's Disease. 2018 Dec;4(1):11.

101. Marmon AR, Enoka RM. Comparison of the influence of two stressors on steadiness during index finger abduction. Physiology & Behavior. 2010 Mar;99(4):515–20.

102. Schwabe L, Schächinger H. Ten years of research with the Socially Evaluated Cold Pressor Test: Data from the past and guidelines for the future. Psychoneuroendocrinology. 2018 Jun;92:155–61.

103. Stark R, Wolf OT, Tabbert K, Kagerer S, Zimmermann M, Kirsch P, et al. Influence of the stress hormone cortisol on fear conditioning in humans: Evidence for sex differences in the response of the prefrontal cortex. NeuroImage. 2006 Sep;32(3):1290–8.

104. Oyarzún JP, Lopez-Barroso D, Fuentemilla L, Cucurell D, Pedraza C, Rodriguez-Fornells A, et al. Updating Fearful Memories with Extinction Training during Reconsolidation: A Human Study Using Auditory Aversive Stimuli. El-Deredy W, editor. PLoS ONE. 2012 Jun 29;7(6):e38849.

105. Redondo J, Fraga I, Padrón I, Piñeiro A. Affective ratings of sound stimuli. Behavior Research Methods. 2008 Aug;40(3):784–90.

106. Quiroga RQ, Garcia H. Single-trial event-related potentials with wavelet denoising. Clinical Neurophysiology. 2003 Feb;114(2):376–90.

107. Kundu P, Inati SJ, Evans JW, Luh W-M, Bandettini PA. Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI. NeuroImage. 2012 Apr;60(3):1759–70.

108. Liu KY, Marijatta F, Hämmerer D, Acosta-Cabronero J, Düzel E, Howard RJ. Magnetic resonance imaging of the human locus coeruleus: A systematic review. Neuroscience & Biobehavioral Reviews. 2017 Dec;83:325–55.

109. Glover GH, Li TQ, Ress D. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magn Reson Med. 2000 Jul;44(1):162–7.

110. Mather M, Joo Yoo H, Clewett DV, Lee T-H, Greening SG, Ponzio A, et al. Higher locus coeruleus MRI contrast is associated with lower parasympathetic influence over heart rate variability. Neuroimage. 2017 Apr 15;150:329–35.

111. Clewett DV, Lee T-H, Greening S, Ponzio A, Margalit E, Mather M. Neuromelanin marks the spot: identifying a locus coeruleus biomarker of cognitive reserve in healthy aging. Neurobiology of Aging. 2016 Jan;37:117–26.