Computational Neuroscience Methods: Investigating the Relationship between Resting-State EEG Microstate Syntax and fMRI BOLD Signal

by

David Graham Haydock

Submitted to
the University of Hertfordshire
in Partial Fulfilment of the Requirement
of the Degree of Doctorate of Philosophy

Biocomputation Research Group
Centre for Computer Science and Informatics Research
January 2024

Declaration

I, David Graham Haydock, confirm that the work presented here is my own. All code used for this project can be found via the link to the repository here. Where information has been derived from other sources is indicated clearly in the thesis or in the repository. I declare also that no data or results were added to the repository after the time of submission of this thesis.

This study was based upon work supported by the United States Air Force Office of Scientific Research (AFOSR) under award number FA9550-19-1-7034 to Prof Chrystopher Nehaniv and Dr Elena Antonova. The funder had no role in study design; in data collection, analysis and interpretation; or in the writing of the thesis.

General Abstract

Simultaneous recording of Electroencephalography (EEG) and functional Magnetic Resonance Imaging (fMRI) has been used consistently in the past as a means of understanding EEG microstate function. EEG microstates are quasi-stable states of EEG activity. Here, I investigate existing methodologies that attempt to draw relationships between microstate classes and fMRI signal, shedding light on their limitations and proposing alternative methods which may better utilise the advantages of simultaneously recorded EEG-fMRI.

Three distinct studies are presented, each using a novel methodology which compares EEG microstates to the simultaneously recorded fMRI signal in resting state recordings. Each proposed method could be used and developed upon in the future to address gaps in the existing literature.

The first study shows how EEG microstate n-grams exhibit varied durations and frequencies in some participants during concurrent fMRI Co-Activation Patterns (CAPs). The second study employs a random forest regressor model, utilising microstate n-gram parameters as features per fMRI time point in a sliding window, attempting to predict patterns in fMRI activity in a low dimensional space. In the third study, the focus shifts to conceptualising the EEG signal as a continuous signal rather than sequence of microstates, with analysis of microstates occurring post-hoc; a novel means of investigating microstates which has not yet been attempted.

I also show how existing investigations of microstate syntax may benefit from adjustments to their processing pipelines in order to better retain the information apparent in EEG microstate sequences.

Keywords: Computational Neuroscience; EEG; fMRI; BOLD Signal; Simultaneous EEG-fMRI; EEG Microstates; EEG Microstate Syntax; fMRI Gradient Space; Methodological Developments.

Acknowledgements

I would like to begin by thanking my supervisory team. Elena, thank you for committing so much of your time to helping me grow as a researcher. I have learned a lot from you about EEG, mental states, academia, meditation and neurophenomenology. I didn't even know what neurophenomenology was until I met you! So thank you, and I look forward to working with you in the future. Rob, thank you for the countless calls where we went over code and I felt like I was taking up your time. Thank you for taking the time to explain complicated topics to me in simple terms, and giving me opportunities to meet other researchers. I know you probably will think that you didn't do much, but you did a lot. So thank you. I'm going to keep going to steal those hours away from you on calls on my post doc. Chrystopher, we were hoping to get into the epsilon machines and we didn't manage to quite get there. There ended up being so much to do on the path there that we never made it. But you still taught me a lot about them. Thank you for taking me through your teaching course, and for your thoughtful suggestions on methods development. I hope that we can keep in touch and dive into semi-group theory once again in the future. Shabnam, thank you for bringing your suggestions to the table each week and taking the time to think through and understand my code. Your insight was indispensable and I am grateful for your help. I would also like to thank Ting Xu and Jonny Smallwood for exposing me to useful methods.

Away from academia, I would like to thank my Mum and John, and my brother Ali. Mum, I know that when I first said I wanted to do a PhD, even when I said I wanted to do a Masters, you thought I might pulling your leg. But I am very grateful for you supporting me. It looks like it worked out after all. Thank you and John for helping with all the annoying moves during this journey, I will be forever in your debt. Ali, thanks for being around during my breaks to chat to and take my mind off work. I hope football went well this weekend. Luke, I wanted to put you in here so I might be the first neuroscientist to cite an astrophysicst, but I'm sure someone else has done that already. Thanks for being such a solid friend during this time in my life. I hope we can find something to work on together in the future. Some honourable mentions for thanking are Kat, Aaron and Antoine for your support during my work. Sorry I dropped off the face of the earth whilst I was writing this, I promise I'll get back in touch. Finally, I want to thank Harriet. You have been an inspiration for me to keep going throughout this process, and I don't think that anyone understands what I went through with this better than you do. You have been so supportive and helpful every day, especially since I started writing. Thank you so much. You mean the world to me.

Thank you to all of you. I really could not have done it alone. I am forever grateful.

Contents

List of Figures

List of Abbreviations

ACC Anterior Cingulate Cortex.

ADHD Attention Deficit Hyperactivity Disorder.

AG Angular Gyrus.

AI Anterior Insula.

AIF Auto-Information Function.

ASR Artefact Subspace Reconstruction.

BOLD Blood-Oxygenation Level Dependency.

CAP Co-Activation Pattern.

CRT Choice Reaction Time.

CV Cross-Validation.

DAN Dorsal Attention Network.

dFC dynamic Functional Connectivity.

dlPFC Dorsolateral Prefrontal Cortex.

DMN Default Mode Network.

DOF Degrees Of Freedom.

ECG Electrocardiography.

EEG Electroencephalography.

EOG Electrooculargraphy.

EPI Echo-Planar Imaging.

FC Functional Connectivity.

FDR False Discovery Rate.

fMRI Functional Magnetic Resonance Imaging.

FWHM Full-Width-at-Half-Maximum.

GEV Global Explained Variance.

GFP Global Field Power.

GLM General Linear Model.

HCP Human Connectome Project.

HRF Haemodynamic Response Function.

IC Independent Component.

ICA Independent Component Analysis.

ICN Intrinsic Connectivity Network.

IFG Inferior Frontal Gyri.

IPS Inter-Parietal Sulcus.

LGN Lateral Geniculate Nucleus.

LSTM Long Short-Term Memory.

mPFC medial Prefrontal Cortex.

MR Magnetic Resonance.

MRI Magnetic Resonance Imaging.

MSE Mean Squared Error.

MT Middle Temporal.

PCA Principle Component Analysis.

PCC Posterior Cingulate Cortex.

PFC Prefrontal Cortex.

REM Rapid Eye Movement.

RNN Recurrent Neural Network.

RSN Resting State Network.

SNR Signal-to-Noise Ratio.

TE Echo Time.

TMS Transcranial Magnetic Stimulation.

TR Repetition Time.

vmPFC ventromedial Prefrontal Cortex.

WM Working Memory.

1 General Introduction

1.1 Background and Context

Electroencephalography (EEG) is a means of measuring the electrical activity of the human brain non-invasively with high temporal resolution. Electrodes are evenly distributed across the scalp of participants to capture the topographical distribution of electrical activity. The signal captured at the scalp represents the local current flows of active neuronal assemblies. Only large populations of active neurons can generate electrical activity recordable on the head surface using EEG (Teplan, 2002).

Observed activity patterns across time commonly have sinusoidal wave shapes, which range in frequency. These ranges are categorised into bands: delta ($0.5-4$Hz), theta ($4-8$Hz), alpha ($8-13$Hz), beta ($13-30$Hz), and gamma (>30Hz). The alpha band frequency range is dominant in the human brain during quiet wakefulness and has been the most extensively investigated frequency range (Teplan, 2002). Whilst its precise neural origin is not known, it is widely believed that alpha activity is driven by layer V pyramidal cells within the cortex (Silva et al., 1991). The primary alpha pacemaker is believed to be the thalamus, with the posterior alpha rhythm being driven by the pulvinar nuclei and/or the Lateral Geniculate Nucleus (LGN) (Halgren et al., 2019).

1.1.1 Definition of EEG Microstates

EEG activity is somewhat discontinuous in that it is characterised by rapid changes in spatial configuration, followed by periods of quasi-stable topographic distribution or maps, referred to as EEG microstates (Lehmann et al., 1987). These quasi-stable topographies maintain a consistent location of the maximum positive and negative potential poles on the scalp for a short period (approx. $30-60$ms), with the poles swapping their locations intermittently. The formal definition of a microstate ignores the intermittent pole switching, only considering the pole locations, with a single microstate lasting approximately $80-120$ms in eyes-closed resting-state healthy human participants throughout their lifespan (Koenig et al., 2002). Microstates have historically been investigated in the alpha band due to past investigations of the resting state (Lehmann et al., 1987). A past study has suggested that each microstate may be primarily driven by a different alpha generator using source localisation methods (Milz et al., 2017) [1].

Four observed topographies are more common across time points than others and are highly replicable across studies in both healthy and clinical populations, accounting for about 80% of variance in eyes-closed resting-state EEG signal

[1] Although it may be the primary driver, all studies that have investigated microstates thus far apply bandpass filtering during preprocessing (usually between 2 and 20Hz Michel and Koenig (2018)), which may limit the impact of beta and gamma bands. Furthermore, the method used to derive microstates is a simple clustering algorithm, which could be applied to different frequency bands, yet I am unaware of any study that has attempted to identify similar microstates in other bands.

(Michel & Koenig, 2018). These four discrete topographic classes are referred to as "canonical" EEG microstates and are labelled as A, B, C, and D. Figure 1 shows four canonical microstate class topographies, as adopted from Milz et al. (2017). Microstate A is characterised by a right frontal-to-left posterior configuration, whilst B is a mirrored left frontal-to-right posterior configuration. Microstate C has a symmetric anterior-to-posterior configuration, and microstate D has a similar symmetric configuration, but has a more central pole than C's anterior pole.

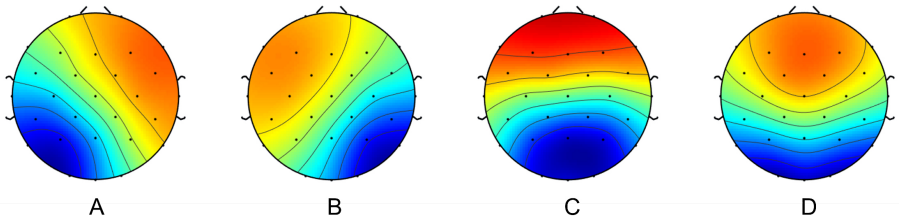

Figure 1: The canonical microstate classes A to D left to right adopted from Milz et al. (2017). EEG microstate classes are defined by the topographical distribution of electrical activity across the scalp. Note that the location of the poles defines a microstate class, with the polarity itself considered to be irrelevant.

Different methods have been used to derive EEG microstates (e.g., Yuan et al. (2012)), but the most common are forms of clustering analysis, which through a process of labelling only allow for one microstate class to be active at each time point (Pascual-Marqui et al., 1995). In the clustering analysis process, EEG time points where Global Field Power (GFP) peaks occur are used as input into the clustering algorithms. GFP is equal to the root mean square across the electrodes, i.e., the standard deviation across electrodes at each time point:

$$GFP(t) = \sqrt{\frac{\Sigma_{i=1}^{n}(v_i(t) - \bar{v}(t))^2}{n}} \tag{1}$$

where $v_i(t)$ is the measured voltage of electrode i at time point t, $\bar{v}(t)$ is the mean voltage across electrodes at time point t, and n is the number of electrodes (Murray et al., 2008).

Topographies which occur at local maxima of the GFP time series have a high Signal-to-Noise Ratio (SNR), and therefore the GFP peaks are used for the derivation of microstates, with the time points in between, which have low SNR, assigned a microstate observed at a neighbouring GFP peak (Lehmann et al., 1987). Clustering methods vary across the studies, with the most common being the modified k-means clustering algorithm (Pascual-Marqui et al., 1995). The algorithm modifies the standard k-means clustering by taking into account the polarity invariance of the EEG topographies used as input.

The canonical set of microstates have been demonstrated to be stable across different clustering algorithms, as well as different numbers of electrodes, with as few as 16 electrodes across the scalp being sufficient to derive four canonical classes (Khanna et al., 2014). Due to this demonstrated stability, some studies (e.g., Milz et al. (2017)) opt to label observed GFP peaks based on spatial similarity to microstates derived in other studies (Koenig et al., 2002; Milz et al., 2016), rather than using data-driven microstate classes, to increase generalisability across the studies when investigating the functional significance of EEG microstates.

Whilst the approach of assigning predefined microstate maps to a dataset ensures that the same microstate classes are investigated across studies, a data-driven approach to the derivation of microstates provides a more accurate representation of the data being investigated and is therefore more preferable for advancing the understanding of EEG microstates and their functional significance across cognitive states and populations (Michel & Koenig, 2018).

1.1.2 Choosing the Optimal Number of EEG Microstates

When using the data-driven approach of deriving microstates, the studies have applied a set of criteria to determine the best fitting number of microstates for their given dataset (e.g., Custo et al. (2017) and Michel and Koenig (2018)).

The Global Explained Variance (GEV) is commonly used to quantify the quality of class assignment of GFP peaks to each of the microstates (Murray et al., 2008). GEV is defined as:

$$GEV = \frac{\Sigma_{t=1}^{t_{max}} (GFP(t) \cdot r)^2}{\Sigma_{i=1}^{t_{max}} GFP^2(t)} \tag{2}$$

where r is the Pearson's cross-correlation between the given microstate topography and the current time point topography.

Another is the cross-validation criterion, which is a ratio between GEV and the degrees of freedom of the topographies in question (Pascual-Marqui et al., 1995). More recently, a meta-criterion has been developed which considers multiple existing clustering criterion and computes an overall fit for the given number of cluster centres (Michel & Koenig, 2018).

When implementing a data-driven approach to microstate derivation, different studies may arrive at different numbers of microstate classes, with potentially different cluster centres, making meta-analysis between studies difficult. To mitigate the issue, a database of data-driven microstates has been developed, where the input microstates are compared to those already found in previous studies through meta-analysis (Koenig et al., 2023). The microstate topographies found across studies were also clustered in a meta-analysis, referred to as "meta-microstates". When clustering microstates generated from over 50 studies using $k = 4$, the canonical set of A to D were indeed identified. However, when increasing to $k = 5, 6$ and 7, the meta-microstates E, F and G in addition to the canonical A-D set were defined. Microstate E is similar to microstate D, but has a more posterior location of its central pole. Microstates F and G show

similarities to B and A respectively with their left-right distributions, but with less frontal-posterior splits and more hemispheric pole locations.

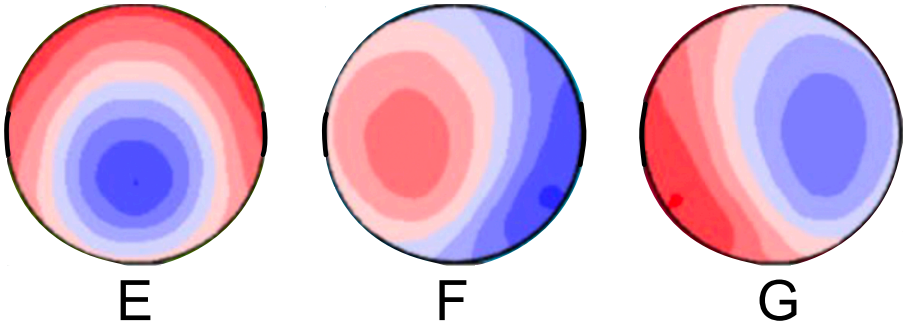

Figure 2: The meta-microstates additional to the canonical set of 4, generated by Koenig et al. (2023), with microstate classes E to G left to right. EEG microstate classes are defined by the topographical distribution of electrical activity across the scalp. Note that the location of the poles defines a microstate class, with the polarity itself considered to be irrelevant.

It should be noted, however,, that the microstate label assigned to a specific cluster centre in some of the studies was not necessarily the same as the meta-microstate label derived in the meta-analysis Koenig et al. (2023).

1.1.3 EEG Microstate Parameters

It is immediately apparent that due to their quasi-stable nature and relatively short duration, it is not sufficient to attempt to compare microstates between tasks, groups or participants simply by occurrence alone. It is common for all microstate classes derived in an analysis to occur multiple times within a 1-second window of EEG time series. For this reason, attempting to associate the individual occurrences of these states to cognitive processes or brain networks is untenable. Therefore, parameters have been derived to characterise how microstate behaviour varies between tasks and groups.

The standard parameters used are mean duration, occurrence, coverage and GEV. Mean duration is defined as the average amount of time that a instance of a microstate lasts for, formalised as:

$$\bar{d}_m = \frac{\Sigma_{i=1}^{s_m} d_{m_i}}{s_m} \tag{3}$$

where d_{m_i} is the duration of the given occurrence i of microstate m; s_m is the total count of occurrences of the given microstate m, $\bar{d}_m$ is the mean duration of the given microstate m. [2]

[2]How the duration of an individual microstate (d_{m_i}) is determined varies between methodological approaches, which can impact investigations of transitions between microstates. See Chapter 3 for a discussion of this area.

Occurrence is defined as the average number of times that a given microstate occurs within a 1-second window, formalised as:

$$o_m = \frac{s_m}{T} \tag{4}$$

where o_m is the occurrence rate of microstates m and T is the total amount of time in the time series in seconds.

Coverage is defined as the percentage of the whole time series that the given microstate is dominant, formalised as:

$$c_m = \frac{\Sigma_{i=1}^{s_m} d_{m_i}}{T_f} \tag{5}$$

where c_m is the coverage of microstate m and T_f is the amount of time in the whole time series in the appropriate sample frequency f.

Note the distinction made between T and T_f between Equations 4 and 5. This is because coverage is calculated using durations of microstates rather than a count of occurrences of microstates. This is done because the percentage of the time series that is dominated by a given microstate includes all time points, not just the GFP peaks.

GEV is also often used to characterise microstates individually (defined in Equation 2), where the overall GEV is the sum of each individual microstate's contribution to explained variance.

A combination of these parameters allows characterising microstates' dynamics within- and between-datasets, experimental conditions, mental states, and/or populations.

1.1.4 Functional Significance of EEG Microstates: Findings from the Studies Of Clinical Populations

EEG microstates could potentially be developed into clinical bio-markers since differences in microstate parameters and their dynamics have been reported in the EEG microstate dynamics between clinical populations and healthy individuals. [3]

Firstly, microstate A has shown longer duration and higher coverage in groups with panic disorder versus controls (Kikuchi et al., 2011). Conversely, it has shown shorter duration and lower coverage in adult Attention Deficit Hyperactivity Disorder (ADHD) patients (Férat et al., 2021). A decrease in the duration of A has also been observed in Parkinson's patients from before to after drug administration which caused an increase in dopamine levels (Serrano et al., 2018).

[3] Since the inception of the meta-microstates software (see Section 1.1.2, or Koenig et al. (2023)), understanding which microstates are associated with which function or dysfunction has become difficult without the use of the software. A study published before the creation of this software may label a microstate with a C yet meta-analysis suggests it is labelled with meta-microstate D, for example. For this reason, the following section which covers the functional significance of microstates, will refer to the meta-microstate label of the microstate in question where possible.

Microstate B has shown a shorter duration, occurrence and coverage in bipolar patients versus healthy controls (Vellante et al., 2020), with all parameters showing a negative correlation with self-reported symptoms of dissociation and anxiety. Higher coverage has also been observed in Parkinson's disease patients versus healthy controls (Chu et al., 2020).

Microstate C has shown a lower occurrence in panic disorder patients versus healthy controls (Kikuchi et al., 2011), and shorter duration in fronto-temporal Dementia patients versus healthy controls (Nishida et al., 2013). A lower occurrence and coverage in Parkinson's disease patients versus healthy controls was also observed (Chu et al., 2020).

A longer duration of microstate D has been observed in ADHD patients (Férat et al., 2021), whilst shorter duration has been reported in schizophrenia patients across multiple studies (Kindler et al., 2011; Nishida et al., 2013; Soni et al., 2018; Strelets et al., 2003).

In general, a decrease in the duration across all canonical states in schizophrenia patients has been reported in multiple studies, (Nishida et al., 2013; Soni et al., 2018; Strelets et al., 2003). This may indicate a need for balance of the parameters of microstates C and D, with the dominance of microstate C resulting in detachment of mental states from environmental output.

Multiple studies have also shown differences in multiple microstate parameters between Alzheimer's disease patients and healthy controls (Schumacher et al., 2019; Strik et al., 1997; Tait et al., 2020), but the differences are inconsistent (Tait et al., 2020).

Canonical microstate parameters have also been shown to differ between age groups and sexes. Microstate D has been shown to have longer duration in males, and C longer duration in females. Both sexes show changes in microstate parameters between age groups, and some share changes with age, such as a general increase in microstate D occurrence from childhood to adulthood (Tomescu et al., 2018).

Overall, although EEG microstates hold promise as biomarkers of psychopathological and neuro-developmental/degenerative disorders, there is a lack of replicability across studies, with a notable exception for schizophrenia. Furthermore, earlier studies focused on the canonical set of 4, which might have contributed to inconsistency of findings across the studies. For example, when the data-driven approach was applied to the Parkinson's patients vs. healthy controls comparison, the observed microstate D in healthy controls did not occur as a cluster centre in Parkinson's patients, and instead was replaced by meta-microstate E when using four cluster centres(Chu et al., 2020; Pal et al., 2021).

Despite these preliminary findings differentiating clinical and non-clinical populations, it is difficult to ascertain the functional significance of EEG microstates from comparing individuals with complex disorders characterised by often co-morbid symptoms as well as overlapping and divergent cognitive dysfunction. Utilising cognitive manipulation processes in healthy controls has yielded further insight into the functional significance of EEG microstates.

1.1.5 Functional Significance of EEG Microstates: Findings from Cognitive Manipulation, Source Localisation and fMRI BOLD Signal Association Studies

Several studies have investigated the functional significance of canonical EEG microstates by comparing EEG microstate parameters during different experimentally manipulated cognitive processes/mental states in healthy participants. Although all microstates occur during all cognitive processes/experimental conditions, there are a few consistent findings across the studies regarding the relative dominance of a particular microstate during a particular cognitive process. Additionally, given that more is known about the relationship between mental states/cognitive processes and brain regions/large-scale networks (Damoiseaux et al., 2006; Karapanagiotidis et al., 2020; Yeo et al., 2011), studies have attempted to better understand the functional significance of EEG microstates by using source localisation techniques (Custo et al., 2017; Grech et al., 2008), or by associating EEG microstates with Blood-Oxygenation Level Dependency (BOLD) signal using simultaneously recorded EEG/Functional Magnetic Resonance Imaging (fMRI)data (See Chapter 2 for an in-depth review of the association methods used in previous research, including the overview of their shortcomings).

Thus, microstate *A* parameters have been found to increase during tasks requiring visualisation. A higher duration and occurrence during both object and spatial visualisation tasks as compared with a verbalisation task and no-task resting state (Milz et al., 2016), as well as during visualisation tasks when compared to a spatial reasoning task (Zanesco et al., 2021) have been reported. Seitzman et al. (2017) reported decreased duration of microstate *A* in an eyes-open vs. eyes-closed condition during a serial subtraction task, possibly reflecting decreased visualisation with eyes-open due to the increased visual sensory input. An indirect evidence for the association of microstate *A* with visualisation comes from the reported increased duration, occurrence, and coverage when moving from rest to light to deep hypnosis (Katayama et al., 2007), with the latter known to be associated with increased visual imagery (Lanfranco et al., 2021).

However, other findings appear to be inconsistent with the visualisation association. Lower duration and higher occurrence of microstate *A* (i.e., shorter but more frequent) have been reported during self-reported verbal thoughts such as 'I thought in words' or 'I imagined talking to myself' (Tomescu et al., 2022). Other studies reported no significant differences in microstate *A* parameters between experimental conditions requiring visual vs. verbal processing (Antonova et al., 2022; D'Croz-Baron et al., 2021). Microstate *A* duration was also found to positively correlate with alertness ratings during rest and a verbalisation task but not a visualisation task (Antonova et al., 2022), contradicting association with visualisation.

Some have suggested that associations made between microstates and fMRI networks may further support microstate *A*'s association with visualisation. Britz et al. (2010) associated microstate *A* with negative BOLD signal in the

left-lateralised phonological network, with source localisation studies showing similar networks (Custo et al., 2017). It has been suggested that these two associations could potentially be reconciled by Antonova et al. (2022). The reasoning is the following: microstates may be driven primarily by sources in the alpha band (Milz et al., 2017), and the alpha band has shown inhibition of neural activity in the past. This combined with the suggestion that a negative BOLD signal may be associated with an inhibition of neural activity (Sten et al., 2017). Hence, the association reported by Britz et al. (2010) could be interpreted as an inhibition of the language processing network during visualisation (Antonova et al., 2022).

It should be pointed out however that association of alpha band activity with neural inhibition is perhaps an oversimplification. O'Gorman et al. (2013) highlighted alpha activity may be associated with negative BOLD signal. Whilst some associate negative BOLD signal with inhibition of neural activity (Sten et al., 2017), it may also be due to regions that are functionally connected to areas exhibiting positive BOLD responses, reflecting complex neural interactions within the brain (Braga & Leech, 2015; Leech et al., 2014), or indeed, neural desynchronisation (Mantini et al., 2007).

Microstate B parameters (duration, occurrence, and coverage) have been reported to increase during tasks requiring verbal processing as compared to visuospatial processing tasks and no-task rest (Milz et al., 2016). However, other cognitive manipulation studies have reported a higher coverage in a visualisation task versus a verbalisation task and rest (Antonova et al., 2022) and an increased coverage and occurrence during the eyes-open vs. eyes-closed resting state (Seitzman et al., 2017).

Microstate B was associated with negative BOLD signal in the visual network (Britz et al., 2010). Source localisation studies have also placed microstate B generators in the visual cortex (Bréchet et al., 2019; Custo et al., 2017; Diezig et al., 2022). The same argument regarding reconciliation of these two findings as with microstate A have also been made (Antonova et al., 2022), and the same criticisms as above are also apparent here.

Increases in microstate B parameters have been reported during other cognitive states, including higher class B duration being associated with less effort during a visualisation task (Antonova et al., 2022), stage 2 non-Rapid Eye Movement (REM) sleep (Brodbeck et al., 2012), and meditation versus rest (Faber et al., 2005). A recent study has also found a positive correlation between microstate B and simultaneously recorded fMRI BOLD signal in the auditory network during sleep (Xu et al., 2020). These findings suggest that there is more to be understood about the functional significance of microstate B.

Microstate C occurrence was reported to decrease during visualisation tasks relative to resting state (Milz et al., 2016; Seitzman et al., 2017). Class C was also less frequent during a mathematics task and more frequent in a memory task (Bréchet et al., 2019), as well as decreased in duration, occurrence, and coverage during a serial subtraction task (Seitzman et al., 2017). Based on these findings and microstate C topography, Seitzman et al. (2017) proposed microstate C reflects the anterior part of the

The association of microstate C with positive BOLD in the salience network, the Anterior Cingulate Cortex (ACC), bilateral Inferior Frontal Gyri (IFG), and right Anterior Insula (AI) (Britz et al., 2010), implicated in various aspects of executive function, including cognitive control, error monitoring, attention and working memory (Tops & Boksem, 2011), also appears to contradict the proposed relationship between microstate C and mind-wandering. It has been suggested that microstate C may be associated with the interplay between resting state and attention/executive control networks, or attention reorientation (Michel & Koenig, 2018).

Supporting the proposal that microstate C may be associated with attention orientation, a Transcranial Magnetic Stimulation (TMS) study (Croce et al., 2020) found after the TMS inhibition of the IPS, a key region in the Dorsal Attention Network (DAN), microstate C of the four canonical microstates derived using a data-driven approach during rest was replaced with microstate G. Similarly, when the Angular Gyrus (AG), a key region of the Default Mode Network (DMN), was interfered with, microstate C was replaced by microstate F.

Furthermore, it might be necessary to further split microstate C into two sub-clusters as was done by Custo et al. (2017). Although topographically similar to class C, the reported C' appears to be associated with the activity of the anterior nodes of the DMN. With the canonical set of four microstate classes, the anterior and posterior activation gradient subsumed under the class topography might in fact reflect the activity of different networks: an anterior saliency network, associated with error detection/attention reorienting back to task, vs. a posterior network, associated with mind-wandering/self-referencing. Overall, further studies are required to determine the association between microstate C and underlying functional networks to understand its role in attention orientation and mind-wandering.

Microstate D has also been associated with attention orientation and mind-wandering. Cognitive manipulation studies suggest microstate D's association with the resting state, reporting an increase in duration (Antonova et al., 2022) and occurrence (Milz et al., 2016) during the resting state versus visualisation and verbalisation tasks. Lower class D occurrence was associated with higher alertness ratings during mind-wandering were, whilst higher occurrence with a higher level of spontaneous mind-wandering during verbalisation and visualisation tasks (Antonova et al., 2022). However, Seitzman et al. (2017) reported class D parameters to decrease during the resting state versus task conditions. As noted by others (Antonova et al., 2022; Michel & Koenig, 2018), class C and D topographies in the Seitzman et al. (2017) study resemble class D and C topographies reported by other researchers (see Michel and Koenig (2018) and Milz et al. (2016), with the two classes being highly spatially correlated more generally (Antonova et al., 2022), potentially explaining conflicting findings from behavioural studies in relation to both classes C and D.

The suggested generator of microstate D being the DAN has been consistent between studies (Britz et al., 2010; Custo et al., 2017), with Britz et al. (2010) showing the association with negative BOLD activity in the DAN, indicating an

inhibition of this network, supporting class D association with mind-wandering and attention reorientation during resting state.

The meta-microstates outside the canonical set are understudied, with no confirmed associations of classes E, F and G with cognitive processes. A recent literature review (Tarailis et al., 2023) associated microstate E with interoception/sensorimotor processes based on the findings of twelve studies. Classes F and G were only reported in two studies, with inconsistent findings., Therefore, further investigation of these classes is required to understand their functional significance.

Whilst it might be possible and fruitful to determine the functional significance of each microstate class of the canonical set of four and beyond, these functional associations can only be relative to other cognitive processes or mental states given the very short microstate duration, on average. All microstates are observed during all mental states/cognitive processes in the behavioural studies (Antonova et al., 2022; Milz et al., 2017). Since brain network activity is constantly changing, it is expected that EEG microstate activity would reflect the constant switching between brain networks (Abreu et al., 2021).

Furthermore, it has been shown that observed microstate sequences are non-Markovian at short-term time lags up to 1 second (von Wegner et al., 2017), suggesting the importance (non-randomness) of microstate sequences. However, the behavioural studies to date only looked at the pair-wise transitions between the microstates (Antonova et al., 2022; Milz et al., 2016). Exploring the associations between a more complex microstate syntax with brain network activity might aid a better understanding of the functional significance of EEG microstates.

1.1.6 EEG Microstate Syntax

There is a growing literature which investigates the syntax of microstate sequences. That is, the dynamics of transition between microstate classes, both at the state-to-state level, and across short sub-sequences of states. Multiple studies have demonstrated that these dynamics can be associated to different groups and cognitive states. Transition probabilities are often used as an additional parameter of microstate investigation alongside duration, occurrence, coverage and GEV (see Section 1.1.1 on parameters).

Differences in transition dynamics between groups may elucidate some of the shortcomings of studies that aim to understand the function of microstates. Multiple studies have shown that the frequency of specific transitions changes between groups. Vellante et al. (2020) demonstrated for example that transitions from microstate B to C were more prominent in patients with bipolar disorder versus controls. Other studies have demonstrated differences in transition probabilities between various task conditions in healthy participants (Antonova et al., 2022).

Whilst it is clear that differences in transition probabilities point to microstate syntax being associated to function in some way, investigating transitions between individual states may be too brief of a period of investigation.

Important steps in this direction came from von Wegner et al. (2017), where properties of the transition matrix between microstates were analysed using information theory. The analysis showed a non-Markovian element to resting state microstate sequences at the zeroth, first and second order. They also demonstrated that the information retained in the microstate sequence at different time intervals showed peaks in contribution every 50ms periodically. The Auto-Information Function (AIF) (a method developed by the authors) showed a devolving into a first order Markov model after 1000ms into the future, with a decrease in peak size from 500ms to 1000ms. von Wegner et al. (2021) then demonstrated that the periodic peaks in retained information could be explained by phase patterns of alpha oscillations, suggesting that the microstate syntax retained the periodicities of the underlying alpha band in its sequence.

Developing upon this Sikka et al. (2020) used a Recurrent Neural Network (RNN) which modelled the microstate sequences at multiple time scales from $200 - 2000$ms, and captured stably recurring microstate patterns. Importantly, they highlighted that the traditional uni-variate measures of microstates could not delineate statistically between a group at rest, and the same group that were in a stressed state. Transition probabilities were also not significantly different between rest and stress conditions. When training an RNN on the microstate sequences of both rest and stress conditions however, they observed that the model could correctly classify the sequence as rest or stress $63 - 73\%$ of the time, demonstrating that microstate syntax changes with cognitive state, and hence has a role to play in understanding microstate function. The approach used by Sikka et al. (2020) also highlights the need for an interpretable method. As it was pointed out by the authors, the hidden layers within the RNN model were too complex for simple visualisation, and hence interpreting the classification process of the model was difficult.

Past studies suggest that microstate transition ratios are not sufficient to understand microstate syntax. Both Tait et al. (2020) and Musaeus et al. (2019) showed that transitions between microstates were not significantly different in Alzheimer's or mild cognitive impairment patients respectively. Tait et al. (2020) highlighted that a non-Markovian difference in syntax may have been observed in the data, but was not identifiable by investigating transition probabilities alone, perhaps suggesting that investigation of longer sequences of microstates is required to better understand microstate syntax.

1.1.7 EEG Microstate n-Grams

Past studies have also investigated shorter length sequences of microstate classes beyond the use of transition probabilities. It was shown by Lehmann et al. (2005) that the frequency of specific microstate sequences of length 4 differed between a control group and schizophrenic patients. The sequence $ACDA$ was more frequent in controls, whereas the opposite order of $ADCA$ was more frequent in patients. Schlegel et al. (2012) has since demonstrated a similar pattern between two groups that showed personality differences. It is the case that these shorter length sequences fall within the period of time where memory is retained

in the sequence. With the average duration of a microstate being $80-120$ms, a length 4 sub-sequence would be expected to on average of $320-480$ms. Since it is clear that longer sequences of microstates need to be characterised in order to understand microstate function and microstate syntax function, a clear definition of sub-sequences is needed. Here we define a sub-sequence of microstates as a microstate "n-gram", where n is the number of consecutive microstate labels in the sub-sequence. An illustration of n-grams of different lengths is shown in Figure 3.

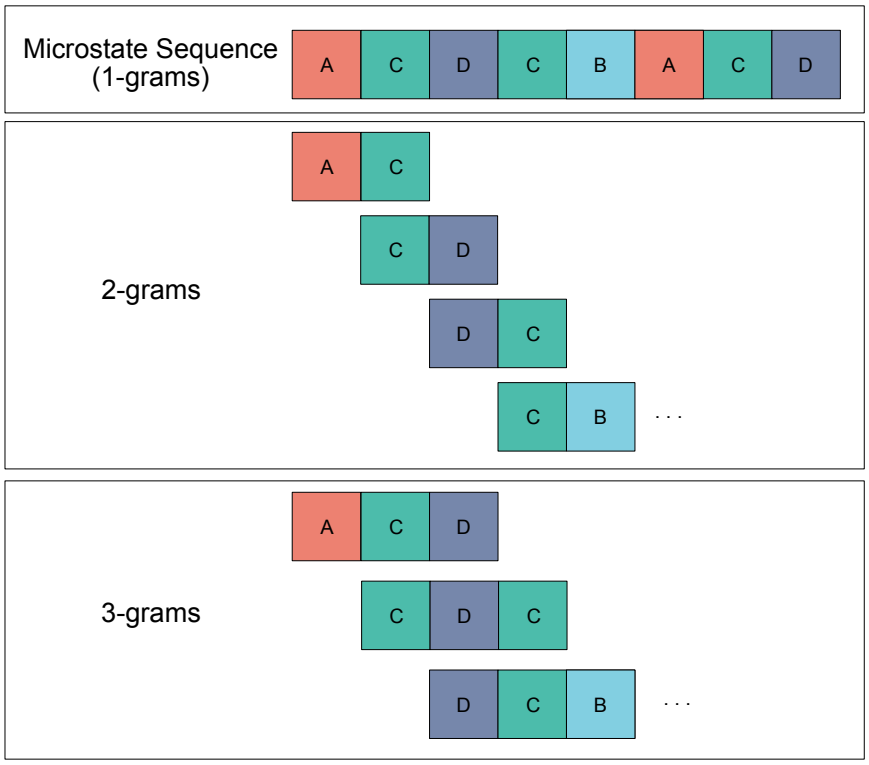

Figure 3: Illustration of event-mode microstate n-grams. Top shows sequence of microstates from observed EEG time series, equivalent to 1-grams. Second row shows 2-grams, 3rd shows 3-grams. n-grams are effectively derived using a sliding window of size n with step size 1.

The n-gram approach has the problem of "dictionary size". If considering a sequence where five microstate classes are the possible states at each time point, the possible number of 4-grams that could be generated (without repetition, $ABAB$, $ABAC$, $ABAD$, $ABAE$, $ABCA$,.. etc) would be 320; its dictionary size. If considering every possible n-gram as n increases, the number of n-grams to consider explodes combinatorially, limiting the length of n-grams that can be investigated.

Artoni et al. (2022) demonstrated a method which circumvented this problem. In the study, 6-grams were investigated in resting state participants, and participants under general anaesthesia. All length 6 sequences of microstates that occurred in the participants were recorded, and were categorised into what were referred to as "entropy classes". Entropy classes were defined by how much new information was contained within an n-gram. For example, the 6-gram $ABDCBA$ shows that within the n-gram, microstate-to-microstate transitions are unpredictable. This would be a high entropy class 6-gram. In contrast, the 6-gram $ABABAB$ shows a predictable transition back and forth between two microstates, and would hence be in a low entropy class.

It was demonstrated by Artoni et al. (2022) that high entropy class 6-grams were more common than expected in the awake, at rest participant, and the same was the case for low entropy class 6-grams in those under general anaesthesia. Such methodologies that reduce the number of n-grams by assigning them to categories are useful and require further investigation. A shortcoming of this approach however is that such approaches assume that the transition between states holds the same weight. It may be the case that the transition from A to B is much more common in the observed data than from B to A, but classification in complexity classes does not take this into consideration.

Whilst there has been significant development regarding the investigation of microstate syntax, I am not aware of any investigation of microstate syntax which associates syntax to brain regions or networks through simultaneous fMRI recording. There are multiple reasons why this may be, but the most prominent are regarding the disparity between the temporal resolutions of EEG and fMRI, as well as a lack of existing methodologies that would allow for a direct comparison between microstate sequences and fMRI signal. It is also apparent that there is yet to be a robust EEG processing and analysis pipeline that considers the syntax of the microstate sequence at every stage. The following two chapters will hence be literature reviews of; methodological approaches for the association of EEG and fMRI, and; methodological considerations regarding the investigation of microstate syntax.

1.2 Aims and Objectives

Microstates have been associated with multiple different processes in different studies through comparison of microstate parameters, and through association with brain regions and networks. It may be the case that a microstate is more active during a certain state, or may exhibit higher coverage when a specific brain region is more active. But all microstates are active during all tasks, and whilst all brain regions are active. For this reason, attempts to associate microstate parameters with individual processes will always leave unexplained associations. Investigations into microstate syntax have begun to show these shortcomings, as the interaction between microstates may be a correlate with functions or brain region activities, and circumvent the need for a one-to-one association between microstate and functional process. However, current methodologies that do investigate microstate syntax use preprocessing and analysis methods which do

not consider the microstate syntax in their pipeline, and therefore may provide misleading results. Therefore, the first aim of this thesis is to demonstrate the shortcomings of current approaches that are used to investigate microstate syntax, and to propose a suitable microstate pipeline that updates the traditional approach so that it can investigation of microstate syntax can be prioritised.

Past studies that have attempted to associate EEG microstates with fMRI have generally done so through reduction of the temporal resolution of the EEG signal to that of the fMRI signal, therefore completely destroying the relationship between microstates and time. More recently, studies that have investigated microstates have clearly shown that microstate syntax retains information with a short term memory, meaning that retention of the relationship between microstates and time is extremely important, and therefore could be used to better understand their function. For this reason, the second aim of this project is to first demonstrate these shortcomings of existing methodologies, and suggest more suitable methods that associate microstates to simultaneously recorded fMRI, whilst retaining the high temporal resolution of EEG, and therefore retaining the grammar inherent in microstate syntax.

Whilst multiple studies have shown that microstates and their parameters may be used as valuable bio-markers for neurological disorders, attempting to understand the function of microstates by only using their duration, occurrence and coverage, may be too simple of an approach. This is exemplified by the multiple contradicting studies that were highlighted in Section 1.1.1. It has been suggested in the past that microstates are better conceptualised as attractor points in a multidimensional space. The parameters are perhaps a surface level investigation of this claim. The "duration of a microstate" in this context can be thought of as how long the EEG time series is under the influence of the given microstate attractor point. For this reason, it may be beneficial to develop methodological approaches which use this concept to investigate microstates. Hence, the third aim of this thesis is to develop a method that uses microstates as a dimensionality reduction of the EEG time series, using them as labels that associate EEG activity to the influence of an attractor point, rather than take the microstates as "winner-takes-all" states.

The final aim of this thesis is to provide a proof of concept to each of these methodological developments by applying them to our own dataset, and providing positive results that associate fMRI activity to microstate syntax in a meaningful way.

1.3 Thesis Overview

The thesis will begin with a literature review which is split into two chapters. The first investigates the current methodological approaches to associating EEG (and more specifically EEG microstates) and fMRI. The second reviews current methodological approaches to the analysis of microstate syntax in the field. Both chapters end with suggested methodological practices based on their findings. The suggested practices are then applied in three separate studies, which all investigate the association between EEG microstate syntax and fMRI ac-

tivity in different ways. Each of the three studies is presented as a standalone investigation, with its own abstract, introduction, methodology, results, and discussion. A general discussion of overall findings between the three studies is then provided in the next chapter, followed by some preliminary investigations into future directions that have already been conducted. The thesis is then ended with a general conclusion chapter.

2 Literature Review: Methodological Approaches to the Association of EEG Microstates and fMRI BOLD Signal Time-Series

Past studies have attempted to characterise simultaneously recorded Electroencephalography (EEG)-Functional Magnetic Resonance Imaging (fMRI) Blood-Oxygenation Level Dependency (BOLD) through various means. Attempts to adapt these existing methodologies for application to the investigation of EEG microstates have been made, with some novel developments also proposed. Here, I review existing methodologies which attempt to associate fMRI BOLD signal to EEG microstates and suggest potential developments.

2.1 fMRI Networks

In order to compare microstates to fMRI, most approaches attempt to simplify the BOLD signal through dimensionality reduction techniques. I separate these techniques into two categories: spatial and temporal. A spatial network is derived through an analysis method which assumes that spatial regions are independent. In these analyses, a single voxel cannot belong to two different networks, making the networks functional parcellations of the brain. The temporal approach instead assumes that time points are independent, and so the whole of the brain is included in a global state. The regions of each global state show different levels of activation. Hence the temporal approach uses dimensionality reduction that summarises differences in activations between each Repetition Time (TR), rather than between voxels. In general, these networks are referred to under the umbrella term of Functional Connectivity (FC). The term refers to spontaneous fluctuations from functionally related regions that show correlation with one another (Biswal et al., 1995).

2.1.1 Spatial fMRI Networks

Spatial fMRI networks have been derived using various analysis methods. Early investigation of BOLD signal during rest suggested a baseline activity of the brain that was represented across studies by functionally relevant states, each referred to as a Resting State Network (RSN) (Damoiseaux et al., 2006). Another name is the Intrinsic Connectivity Network (ICN) (Seeley et al., 2007), used more generally when the generated networks are not necessarily from resting state data. [4]

An example of such spatial networks are those generated by Yeo et al. (2011), where clustering was used to derive parcellations of the cortex into functionally distinct networks based on their connectivity profile. The authors found 7 and 17 to be the most stable numbers of networks in their analyses. Figure 4 shows a visualisation of the 7 networks. Note that these networks are indeed a parcellation, and hence no one region can belong to more than one network.

The spatially distinct networks derived by Yeo et al. (2011) and others (Gordon et al., 2017; Smith et al., 2009; Yeo et al., 2011) have since been used in meta-analyses as the starting point of a proposed "taxonomy" of functional states of the brain that reliably occur across multiple studies (Uddin et al., 2019). Such a technique has been proposed due to difficulty cross-referencing identified spatial networks between studies. Whilst utilisation of this method streamlines meta-analyses, if not used cautiously it may lead to an enforcement of the taxonomy onto spatial networks identified in the future, causing smaller differences between networks to be overlooked.

Although the networks identified by Yeo et al. (2011) used clustering methods in their generation, earlier investigation of ICNs generally used methods such

[4] They were first referred to as ICNs by Seeley et al. (2007) to point out that the brain itself was not at rest when the participant was in the "resting state", since the brain was shown to always be active.

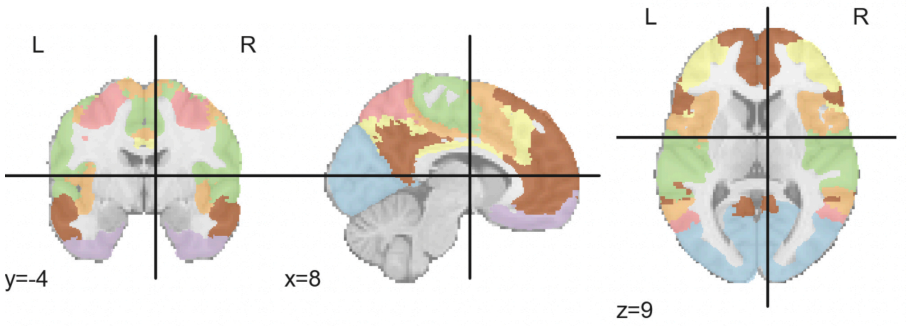

Figure 4: Visualisation of the 7 Yeo networks in volumetric MNI152 standard space provided by Yeo et al. (2011) using the *nilearn* database package (Abraham et al., 2014). Each colour indicates a different network. Note that networks are parcellations, with no voxel belonging to more than one network.

as Independent Component Analysis (ICA), to identify functional networks in the fMRI data (Damoiseaux et al., 2006; Seeley et al., 2007). The benefits of ICA are in its model-free approach. The decomposition of a dataset into a set of independent one-dimensional time series and associated three-dimensional spatial maps allows for a description of the temporal characteristics of a specific spatial patterns signal across the time series (Beckmann et al., 2005). Attempts to associate EEG microstates with spatial fMRI networks have generally used the ICA approach when deriving fMRI networks (Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012).

Seed-based approaches are another spatial network method. This approach is employed when a specific region of interest is expected to be active (Fox et al., 2006). Whilst this approach is more rare in microstate analysis, recent studies have begun investigation in this direction. Case et al. (2017) used ICA outputs of fMRI recordings in sickle cell disease patients using a seed in the cerebellum, since it was expected that connectivity would differ in that region in patients. Bréchet et al. (2019) derived fMRI sub-networks by investigating the activation of different regions during specific tasks, and focused specifically on sub-networks within regions that were expected to be more active during each task. Conceptually there are some differences between local seed-based maps and ICA derived networks, but the difference has been modelled simply and can be accounted for when comparing between methods (Joel et al., 2011).

Reducing the dimensionality of the BOLD time series from thousands of voxels to a few functional parcellations provides an overview of global activity across TRs. However, the major shortcoming of this approach is in the parcellation itself. If two distinct brain regions belong to two different parent networks, their activity will be summarised by the time series of their parent network. If those two regions were active simultaneously, yet overall their parent networks were not active simultaneously, information regarding the functional connec-

tivity of the child regions would be lost. Since the relationship between EEG microstates and fMRI states is not well known, oversimplification of the fMRI into a rigid taxonomy may limit the associations that can be made between microstates and fMRI networks, which is not advisable. These parcellated networks should instead perhaps be utilised as descriptors of common patterns of activity in global activity networks.

2.1.2 Global Temporal fMRI Patterns

The temporal ICA was a method suggested by Smith et al. (2012), after it was highlighted that two ICNs could not have the same brain region in their respective networks, even if they had patterns that proved to be similar. The approach generates global maps where output Independent Component (IC)s are global temporally independent states, rather than spatially independent parcellations. All brain regions are included in all networks in this approach, with differences between states being differences in activation patterns across the whole brain, rather than between different parcellations.

Investigation of connectivity in this area is referred to as dynamic Functional Connectivity (dFC) (Hutchison et al., 2013; Preti et al., 2017). Connectivity matrices are derived either in a windowed or frame-wise manner across the BOLD signal time series, providing a time series of connectivity between regions of interest (Preti et al., 2017). This has been employed recently in the investigation of microstates (Abreu et al., 2021), and is the first to utilise fMRI temporal networks in the investigation of EEG microstates to the authors knowledge.

There are various other temporal network derivation techniques that have not been utilised by the EEG microstate literature however. A relatively simple approach is the Co-Activation Pattern (CAP). CAPs are derived using k-means clustering across time points (Gutierrez-Barragan et al., 2019), with the set of CAPs being the cluster centres. This is an equivalent derivation to the most common methodological derivation of EEG microstates (Michel & Koenig, 2018; Pascual-Marqui et al., 1995). Unlike temporal ICA, the number of CAPs derived in an analysis is participant to the users decision on number of clusters to use. Additionally, generation of CAPs across the time series leads to each TR being labelled with the cluster it belongs to. Hence a discrete sequence of CAPs can be generated, similar to the EEG microstate sequence outlined in previous sections (see Section 1.1.6).

A new functional descriptor of fMRI that is yet to be utilised in the microstate literature is the gradient analysis approach. The term "gradient" is used to describe the internal structure of brain areas. Unlike traditional parcel-based methods, gradients characterise the brain area's organisation in terms of spatial variations between regions. Margulies et al. (2016) derived a set of functional gradient maps by applying diffusion embedding (Coifman et al., 2005) to the overall connectivity matrix of fMRI from the Human Connectome Project (HCP) dataset (n=32,492) (Van Essen et al., 2013). The first three connectivity gradients derived by Margulies et al. (2016) are shown in figure 5.

After meta-analysis of the patterns identified in the first gradient, the au-

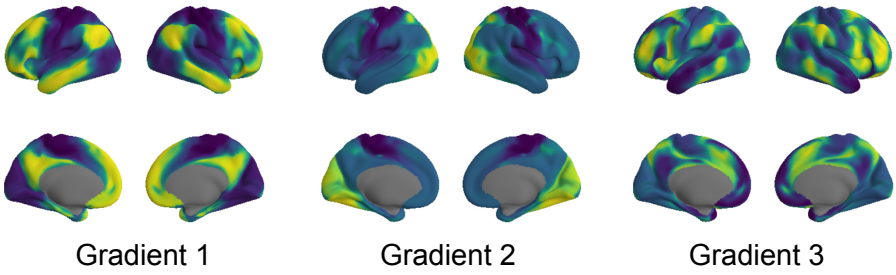

Figure 5: Visualisation of the first three gradient axes, on cortex surface derived by Margulies et al. (2016), 1 to 3 left to right. Note the distribution of activity across regional boundaries.

thors suggest that each of these functional gradients identified is a spectrum of some functional activity. It was demonstrated that regions along the "principle" gradient axis (the axis which explains the most variance in the connectivity matrix, gradient 1 in Figure 5) associated with the Default Mode Network (DMN) were active at one end of the spectrum, whereas primary sensory/motor regions were active at the other end. This principle gradient has been proposed as a dimension of functional organisation in the cortex, as it supports existing theories that activity in the cortex is organised along a "unimodal-transmodal" axis (Mesulam, 1998). The second axis is also consistent with this past theory, which differentiates regions solely within the unimodal end of the principal gradient on a "visual-somoatosensory/motor" axis (Margulies et al., 2016; Mesulam, 1998). The principle gradient has also been shown to co-correlate with gradients that were derived in other observations, such as in intra-cortical myelin (Huntenburg et al., 2018), microstructure (Paquola et al., 2019), structural connectivity (Vos de Wael et al., 2021), cortical thickness (Wagstyl et al., 2015) and gene transcription (Burt et al., 2018).

Haak and Beckmann (2020) discuss how since the inception of systems neuroscience, the brain has been subdivided into regions in different ways, and that each region has been used as reference for functional localisation. Yet brain regions have shown clear "functional heterogeneity" (Jbabdi et al., 2013), that is, even though areas may exhibit a homogeneous cytoarchitecture or connectivity, a functional diversity can be apparent within that region (Purves et al., 2019). For this reason, the use of parcellation can lead to signal mixing. If it is assumed that the signal output of a brain region is due to a single function, when a functional heterogeneity is present may lead to mistakes in inference. In the context of understanding the relationship between EEG microstate syntax, and fMRI activity, use of a global networks which considers gradients rather than parcellations when attempting to understand how microstate syntax relates to spatial patterns may be beneficial.

2.2 Association of fMRI Networks with EEG Signal

Both EEG and fMRI are non-invasive neuroimaging methods that show different aspects of human brain activity. EEG records electrical activity at the scalp, whereas fMRI BOLD signal is generated by changes in the ratio of oxyhemoglobin to deoxyhemoglobin in the brain (Heeger & Ress, 2002).

The main advantage of EEG over other neuroimaging techniques is its comparatively high temporal resolution, generally recorded at the millisecond time scale. Due to electrical activity being captured at the scalp however, intermediate tissues and the skull separate the recording electrodes from the neuronal tissue, making it difficult to identify the exact locations of neuronal generators (Jorge et al., 2014) and decrease the Signal-to-Noise Ratio (SNR) of the neuronal signal (Mulert & Lemieux, 2023).

While fMRI offers high spatial resolution (millimetre-scale voxels), its temporal resolution is relatively low, commonly with a TR of 2 to 3 seconds (0.33-0.5Hz sample frequency). Unlike EEG, which provides direct insights into neuronal activities, fMRI indirectly records these activities through changes in oxygenation of blood within the brain. When a brain area is active it requires more oxygen, which leads to an increase in the ratio of oxygenated to deoxygenated blood flow to that area (Heeger & Ress, 2002). fMRI detects these changes by measuring the magnetic properties of oxygen-rich blood as it flows through the brain. The BOLD signal is a complex combination of hemodynamic and metabolic effects, making it more difficult to interpret than EEG (Logothetis, 2008). The Haemodynamic Response Function (HRF) is a mathematical model that describes how the BOLD signal changes in the brain over time after a neural event such as a change in neural activity or a stimulus (Buckner, 1998). The HRF is typically characterised by a delay between the neural event and the change in BOLD signal, followed by a gradual increase in blood flow over several seconds to the active region.

In the context of associating EEG microstates with simultaneously recorded BOLD signal, the HRF is convolved with the EEG signal at each time point to estimate the change in blood flow over time that would be due to the neural activity observed at the scalp at each instance (Artoni et al., 2022; Britz et al., 2010; Faber et al., 2005; Musso et al., 2010; Yuan et al., 2012).

It is crucial to highlight that both EEG and fMRI need to be aligned in a consistent mathematical space before analysis can proceed. Here, I use terminology coined by Manganas and Bourbakis (2017) to categorise these alignment methods into two groups: asymmetrical and symmetrical. Asymmetrical approaches adjust the resolution of one recording to fit the other, whereas symmetrical approaches do not change the resolution of either recording. [5]

[5] Where Manganas and Bourbakis (2017) used the terms asymmetrical and symmetrical to describe the *overall analysis approach* of associating EEG and fMRI, here I refer to these categories when referring to how the resolution of each recording is modified *before analysis proceeds*.

2.2.1 Asymmetrical Methods of Alignment

Alignments which use asymmetrical methods would either be characterised by a resolution change of EEG to an fMRI space, or vice versa. The former has been attempted in various capacities in the context of EEG microstates.

The most common means by which the asymmetrical approach has been incorporated is the voxel-wise General Linear Model (GLM) (Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012). After convolution with the HRF, the EEG time series is correlated with each microstate topography at each time point, generating a correlation time series for each microstate class. The time series of each microstate is then down-sampled to the same resolution as the recorded fMRI, and each down-sampled microstate time series is used as a regressor in a GLM. A model is constructed for each voxel in the fMRI, with the set of microstates as regressors for each voxel. The map of voxels that correlate with a microstate is then investigated as a generator of the microstate.

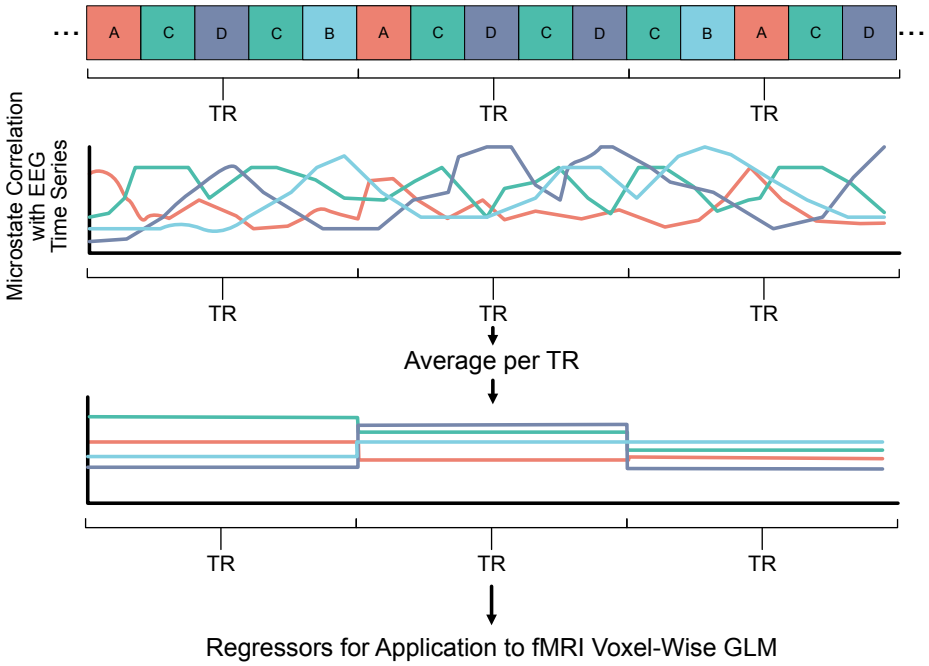

Figure 6: Illustration of the standard process of generating regressors from microstate sequences. First the microstate sequence is derived on the EEG time series (top). The microstate topographies are correlated with each EEG time point (middle), which are then down-sampled to the same sample frequency as the simultaneously recorded TR (bottom). This down-sampled signal is used in voxel-wise GLM analysis. TR denotes the amount of time for a single fMRI time point. Note that duration of microstates is not accurate to observation.

Investigations using this approach use multiple comparison corrections across

voxels to ensure statistical viability (Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012). Additionally, the conversion of the discrete sequence of microstates into a continuous correlation time series for each microstate allows for an investigation of the dynamics of microstate dominance.

The shortcomings of this approach are primarily in its asymmetry. Downsampling the EEG time series to match the temporal resolution of the concurrently recorded fMRI yields a single microstate correlation value for the entire TR. For instance, with a 250Hz EEG sample frequency and TR of 2s, 500 correlation values within a TR are condensed into one. This simplification erases the microstate sequence, which not only undermines any relationship that the microstates may have to time, but also impedes the analysis of their syntax, which may be crucial in understanding their function (see Section 1.1.6).

It was previously outlined that source localisation is used in attempts to estimate sources of EEG microstates (see Section 1.1.5). Source localisation attempts to infer the position of the current sources in the brain from the electrode potentials at the scalp using algorithms like "LORETA" (Pascual-Marqui, 1999), which use "inverse modelling" to estimate the location and strength of current sources of the scalp activity. The term inverse modelling refers to the "inverse problem" (Grech et al., 2008), which states that many different current density distributions in a 3D volume can produce the same potential distribution on the surface of the volume. Source localisation models are built on prior assumptions such as the number of underlying generators in the brain that lead to the spontaneous occurrence of a microstate at the scalp, resulting in an uncertainty regarding the true nature of said generators.

Whilst multiple studies have used source localisation in place of the simultaneous recording of some higher spatial resolution imaging technique to estimate the sources of EEG topographies, the method has also been utilised in simultaneous EEG-fMRI, comparing the estimated sources of EEG microstates with the observed functional states in the fMRI (Bréchet et al., 2019).

Whilst asymmetrical methods of alignment have provided insight into the function of EEG microstates, projecting microstate activity into an fMRI space removes the syntactic structure of the EEG microstate sequence. If a method aimed to investigate microstate sequences and their syntax, it should perhaps instead utilise methods which employ symmetrical alignment methods.

2.2.2 Symmetrical Methods of Alignment

Alignments which ensure that both EEG and fMRI are aligned temporally but do not alter the resolution of either recording are referred to here as symmetrical. In the context of EEG microstate analysis this is a lesser used approach.

Even in the general context of EEG-fMRI studies which do not employ microstate analysis, methods which investigate EEG spectrum analysis for example, tend to use the asymmetrical approach of GLMs (Jorge et al., 2014), or at the very least will down sample the EEG data to one data point per fMRI scan. Methods which simply compare the activity of EEG to fMRI without forcing one resolution to the other are few. One example is joint ICA (Huster et al.,

2012), where the simultaneously recorded EEG-fMRI are subject to the same ICA after concatenation. Some studies have also applied separate ICAs to EEG and fMRI and subject the components to Bayesian-formulated criteria to match the components derived in each recording modality (Lei et al., 2011).

The only paper I am aware of which uses symmetrical methods of alignment for the investigation of EEG microstates is Abreu et al. (2021). The authors refrained from altering the resolution of their simultaneously recorded EEG and fMRI by deriving microstate within each TR window, and comparing microstates between simultaneously occurring fMRI dFC states, and showed that microstates could reliably predict said states. The derivation of microstates within sliding windows raise questions regarding EEG sample sizes. Regardless, this approach clearly demonstrates the strength of using symmetrical methods of alignment.

Despite these existing methodologies exhibiting symmetrical alignment, these approaches do not retain the syntax of the microstate sequence. A potential reason for a lack of such methodology may be down to the difficulty of association post-alignment. If for example fMRI was recorded at TR= 2s, and EEG at 250Hz, the 500 EEG time points that must be compared to the single fMRI the question how the comparison should be conducted arises. Regardless, if the aim of investigation is how EEG microstate syntax relates to fMRI signal, then it is absolutely necessary to retain the EEG time series and not reduce it to a more coarse representation. Each of the three studies presented in this thesis (an overview of which can be found in Chapter 4) suggest a novel methodology for this investigation.

It is the case that existing methodologies which process and analyse EEG signal for the generation of microstates also do not consider the necessity of retaining microstate syntax prior to its investigation. In the next chapter we will review existing methods, of processing EEG and of microstate analysis, and will identify points where syntax is not prioritised, suggesting alternatives that do.

3 Literature Review: Methodological Considerations regarding Microstate Syntax

The primary objective of microstate analysis has been to categorise microstates as bio-markers or cognitive markers (see Sections 1.1.4 and 1.1.5). Approaches used to derive microstate topographies have, therefore, been employed to those ends. Over time, studies have also begun to include the investigation of microstate syntax by investigating one-to-one transitions between microstates (e.g., Vellante et al. (2020)). Despite this, as well as preliminary investigations into longer sequences and more complex characteristics of microstate syntax (see Section 1.1.6 and 1.1.7), methods of microstate derivation still need to be updated to account for possible pitfalls in the observation of syntax.

Here, I investigate existing microstate derivation methods, point out their shortcomings in application to syntax analysis, and suggest updates to existing preprocessing and analysis pipelines that will ensure that microstate syntax is retained by the pipeline up to the analysis phase.

3.1 EEG Pre-processing for Microstate Syntax Investigation

3.1.1 Standard EEG Pre-processing Pipeline

The steps for standard pre-processing prior to Electroencephalography (EEG) microstate analysis include; re-referencing to a reference electrode or the average potential across electrodes, bandpass filtering, downsampling from the recorded frequency, artefact removal if applicable, removal of epochs of noise, removal of noisy channels, and Independent Component Analysis (ICA) followed by rejection of noise Independent Component (IC)s (Kleinert et al., 2023; Michel & Koenig, 2018; Poulsen et al., 2018; Tait & Zhang, 2022). These steps are intended to remove any signal in the EEG that may be from sources other than the brain, such as eye movement, heartbeat, head movements, or muscles in the head, and to generally increase Signal-to-Noise Ratio (SNR). Bandpass filtering is used to isolate the alpha band. In the case of simultaneous EEG-Functional Magnetic Resonance Imaging (fMRI), algorithms for Magnetic Resonance (MR) noise removal are also employed (Allen et al., 2000).

While these methods adequately clean the EEG for microstate derivation, their suitability for investigating microstate syntax requires careful examination. It is immediately apparent that most of these steps do not affect the temporal resolution of the EEG and, therefore, are not cause for concern. However, The two steps that do affect it are down-sampling and epoch removal.

3.1.2 Downsampling

Downsampling from the recorded sample frequency may initially appear to affect microstate syntax. However, microstate studies tend to down-sample to approximately 250Hz from recordings commonly recorded in kHz (Michel & Koenig, 2018). With the average length of a microstate consistently between $60 - 120$ms, a sampling frequency of 250Hz is unlikely to affect transitions between microstates.

3.1.3 Epoch Removal

Moving to epoch removal however, such a practice destroys the overall syntax of the microstate sequence. Microstate sequences that would have otherwise contributed to syntax are removed, meaning either that the periods before and after a "cut" section of the time series are concatenated together, or that the time series before and after the cut have to be considered as individual and separate sequences during analysis. This damages the syntactic structure, and may make identification of the underlying grammar of the syntax more difficult.

Additionally, in the context of simultaneous EEG-fMRI, the alignment of two recording modalities becomes difficult when epoch removal occurs. Periods of noise removed from the EEG time series will be removed in sections that are recorded at a different resolution to fMRI. This difference adds further complexity to the comparison if using symmetrical alignment methods (see Section

2.2.2).

If epoch removal is ill-advised for investigating microstate syntax, the question of how to deal with epochs of noise arises. Including periods of noise in the time series for later microstate analysis would add noise to microstate cluster centres, affecting the microstate syntax differently. Some potential solutions to this problem begin first with cleaning algorithms.

Sinusoidal noise removal techniques are commonly applied using Fourier transforms on a sliding window (Mitra & Bokil, 2007). Bad channel identification algorithms can remove or clean noise that is isolated to a single channel by different forms of interpolation, using neighbouring channels, or using more complex approaches such as random sample consensus (Bigdely-Shamlo et al., 2015). Windows where noise is inherent for a period of time can be cleaned using approaches such as Artefact Subspace Reconstruction (ASR) (Mullen et al., 2013), which interpolates noisy windows based on the rest of the EEG signal.

3.1.4 Improved Pre-processing Pipeline for Microstate Syntax investigations

On top of cleaning algorithms is the change in the role of ICA in this analysis. The standard EEG pre-processing pipeline first uses epoch removal and follows that with noise IC identification and removal. When epochs of noise are not removed but ICA is still applied for data cleaning, noise IC's may include noise inherent in the epochs that would have been removed previously. Algorithms such as *ICLabel* (Pion-Tonachini et al., 2019) or *MARA* (Winkler et al., 2011) automatically categorise IC topographies into their expected source, be that brain or different sources of noise, based on comparison with a dataset of topographies which have had their sources previously identified.

A combination of ICA and cleaning algorithms may allow for a sufficiently clean EEG time series for microstate analysis without needing epoch removal, hence retaining the microstate syntax.

3.2 Microstate Definition and Analysis Methods for Application to Syntax Investigation

Following pre-processing, the process of microstate derivation contains potential pitfalls regarding syntax investigation. Here, the microstate analysis pipeline is reviewed, and potential remedies to these pitfalls are suggested.

3.2.1 Canonical vs Data-Driven Microstates

Section 1.1.1 discussed the difference between canonical and data-driven microstates, focusing on the debate regarding this point in the field (Michel & Koenig, 2018). There is a need for more consideration as to how this difference would affect syntax, however.

Using the canonical set of microstates ensures that a consistent set of states is investigated between studies (Milz et al., 2017). Parameters are derived in

each study on the same classes, meaning that they can be compared across task conditions and between populations. This extends to the investigation of syntax. If multiple studies fit their observed data to a pre-derived set of microstates, then transitions, n-grams, and syntactic rules identified in each of those studies will also be more comparable. However, when fitting the observed data to microstates that were not derived from that data, the possibility of a sub-optimal fit for that specific data means that the syntax observed may not reflect the syntax that would be observed had the microstates been derived in a data-driven manner.

However, a potential issue to consider when investigating data-driven microstates is the number of selected microstates. If using a data-driven method, the optimal number of microstates is high; this increases the number of possible transitions that can occur from one state to another. Increasing the number of microstates causes a combinatorial explosion in the number of n-grams to consider.

3.2.2 Smoothing Destroys Syntactic Structure

There are different approaches by which the duration of a microstate is defined. Worded another way, there are different approaches by which the boundaries between consecutive microstates are defined. The most popular means of deriving microstate boundaries is referred to as *back-fitting* (Poulsen et al., 2018). After microstate (cluster centres) derivation from Global Field Power (GFP) peaks, the topography of each microstate is correlated with each time point in the participant time series, resulting in a correlation time series for each microstate. Each time point is then labelled with the microstate that shows the highest correlation with it in a winner-takes-all fashion, providing a discrete sequence of microstates where the labelled microstate is given for each time point.

This process shows predominant long sequences where a microstate label repeats. The number of consecutive labels of a single microstate is called the microstate occurrence's duration. Due to the conception that a microstate generally lasts approximately $60-120$ms, however, a smoothing step is usually carried out to remove microstate occurrences that are considered too short (Kleinert et al., 2023). This smoothing step removes microstates that last less than a user-defined duration (which varies between studies) through a re-labelling process. The microstates that occur before and after the "short-microstate" are identified, and the short occurrence is assimilated into whichever of the two neighbours is most similar.

While this does homogenise the duration of microstates and controls for "noisy" short microstates, it must be pointed out that the smoothing step introduces unnecessary noise into the microstate sequence. If a section of the time series shows a short period that is most similar to microstate A, for example, but is then re-labelled with C, the transitional structure of the microstates is damaged. This has been highlighted previously (von Wegner et al., 2017), but multiple studies which have investigated microstate sequences since have yet to adopt the method.

Additionally, whilst unlikely, this approach contains a theoretical possibility that would damage the syntactic structure of the sequence further. The microstates themselves are generated using the topographies at GFP peaks with clustering methods (see Section 1.1.1). The back-fitting process then uses these cluster centres (which are only generated using the GFP peaks) to label the whole time series. This approach intends to get a more precise transition point between GFP peaks where the microstate label switches from the former to the latter. It is possible, however, that the smoothing process could occur across a GFP peak if the back-fitting process found the duration of the microstate around the GFP to be under the user-defined threshold. This would result in a GFP peak that was clustered with one centre being labelled with another.

Therefore, if the back-fitting procedure was adopted for syntax analysis, smoothing should be avoided (von Wegner et al., 2017).

An alternative which some have adopted is the interpolation approach (Poulsen et al., 2018; Tait & Zhang, 2022). It is referred to as the interpolation approach because rather than correlating each time point with each microstate, the centre points between GFP peaks are assumed as the transition point. Whilst this approach does lose the more precise boundaries between microstates that appear with the back-fitting approach, the microstate sequence here becomes an analysis of the interactions between consecutive GFP peaks, the points used in the microstate derivation in the first place, avoiding potential noise in the low SNR regions between peaks.

3.2.3 A Formal Definition of Sequence Types - Event, Clock and Peak Mode

While multiple recent studies have investigated microstate syntax in different contexts, the definition of how the sequence should be analysed varies. There are three clear means by which a microstate sequence is defined. An example sequence of microstates will be used to illustrate the difference between these "modes" by which the sequence can be conceptualised, A, shown in Figure 8.

A generated sequence of microstate labels is derived (either by back-fitting or interpolation), with each label being one of the four microstates corresponding to the microstate. The n-gram is $A- > C- > D$, with A lasting 100ms, C 80ms, and D 120ms, with the EEG in question observed at a sampling frequency of 250Hz (each time point lasting 4ms). The first means by which this sequence can be described (and how it is generally described in microstate literature, which does not investigate syntax, see Michel and Koenig (2018) for a review) is the "event-mode". Event-mode refers to the occurrence of a transition from one microstate to another different microstate. The example sequence would be represented simply as ACD, and could hence be described as a 3-gram (see Section 1.1.7 for a definition of n-grams).

The second means by which the sequence can be represented is the "clock-mode". Clock-mode refers to a sequence where each label represents a clock tick on the time series. In this example, there would be 250 microstate labels per second, one label for each time point. Hence, ACD would be represented

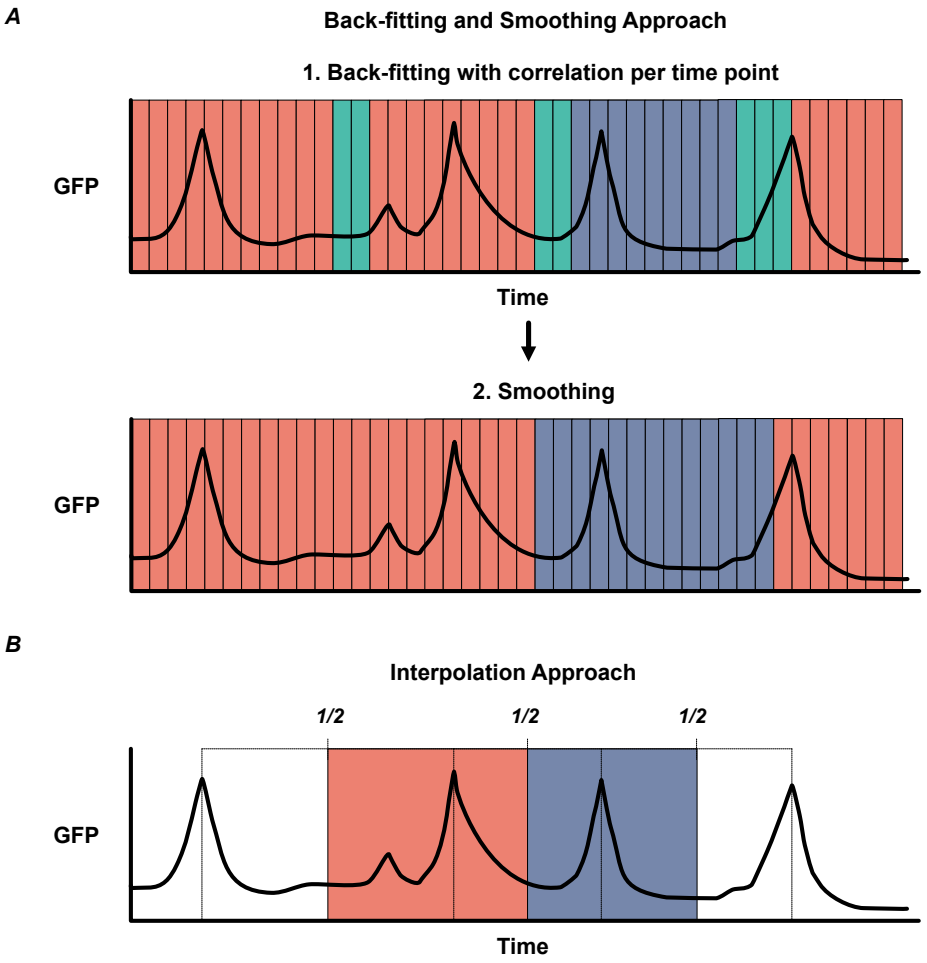

Figure 7: Comparison of the back-fitting/smoothing and interpolation approaches. Colours are representative of different microstates. (A) shows the back-fitting and smoothing method. First on top, each time point is correlated with each microstate, and is then labelled with the microstate it is most similar to (note labels are not to scale temporally on the illustration). Bottom then shows the smoothing process. Where the number of consecutive time points for a microstate is under the user defined number of time points, the short microstate is assimilated into neighbouring microstates. (B) shows the interpolation method. The GFP peaks are subject to clustering and hence are already labelled with microstates. The durations are defined by the mid-points between microstates.

as a sequence of 25 As ($100ms/4ms = 25$ time points), 20 Cs ($80ms/4ms = 20$ time points) and 30 Ds ($120ms/4ms = 30$ time points). In this analysis, the

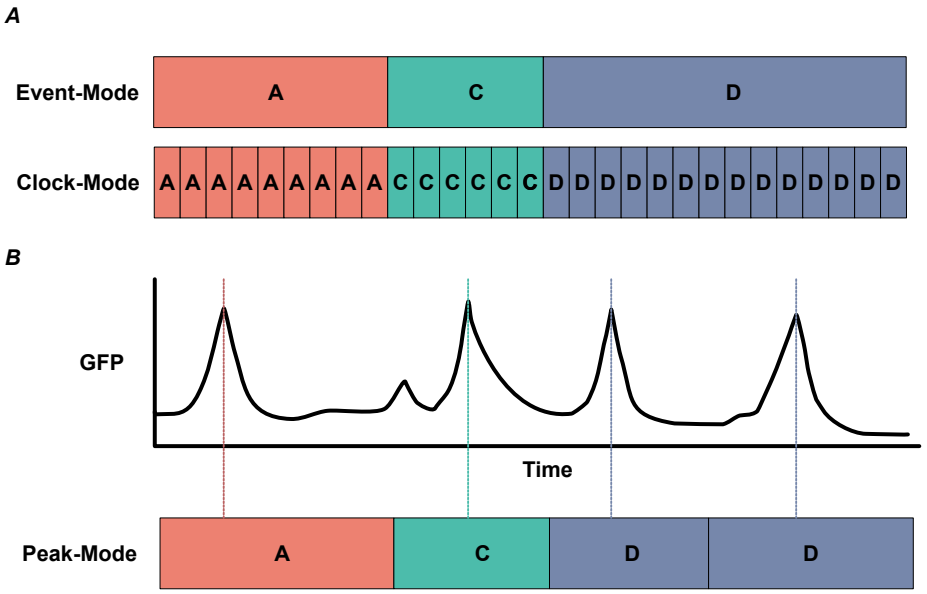

Figure 8: Visualisation of sequence derivation modes, event, clock and peak. (A) Event-mode microstate sequences are shown on the top. Transitions between microstates are where boundaries are defined. Clock-mode microstate sequences are on the bottom, where each clock tick is regarded as a separate state in the sequence (note that illustration is not to scale temporally). (B) Top, illustration of the GFP time series that generated the example microstate sequences. Two GFP peaks in a row labelled with microstate D would constitute a single occurrence of D in the event-mode. Bottom, the peak-mode sequence derived from the GFP time series. Note how there is a single microstate label for each GFP peak.

sequence is tied to the sampling frequency. The only paper that directly defines clock- and event-mode is that of von Wegner et al. (2017), where the authors refer to them as permanence and non-permanence.

The third means by which the sequence can be investigated is referred to as "peak-mode" and, to my knowledge, has not been previously defined. Peak-mode only considers GFP peak labels in the sequence of microstates investigated. Continuing with the example, we see in Figure 8 that the microstate D is comprised of two GFP peaks. The event-mode sequence of ACD would be represented in the peak-mode as $ACDD$, with each D lasting approximately 60ms. Note the possibility for repeating microstate labels, which is impossible in event-mode. Note that peak-mode cannot be derived using the back-fitting and smoothing approach since the point that shows the transition from one D to the next cannot be identified. Peak-mode therefore, is only derivable using the interpolation approach.

The definition of n-grams varies between these modes. The event-mode

3-gram of ACD would be a 75-gram in clock-mode ($AAA..CCC...DDD...D$), and a 4-gram in peak-mode ($ACDD$). Event-mode reduces the dimensionality of the sequence drastically but changes the sequence's relationship to time, as each label in the event-mode sequence lasts a different period. Additionally, to investigate shorter sequences like 3-grams in clock-mode by using a sliding window of size 3, the window would only ever observe either 3 of the same letter or a transition from one microstate to another (i.e., the movement from A to C in the above example using a sliding window which moves one label at a time, would observe AAA, AAC, ACC, CCC).

Further investigation into each of these modes is required for a better understanding of microstate syntax. Here, I propose that investigation using peak-mode is preferred. This is not least because it has a natural development from the preferred interpolation approach but also opens up a preliminary investigation into "within-cluster" dynamics with regards to syntax (microstate D to microstate D) whilst still reducing the dimensionality of the sequence from that of clock-mode.

3.3 Within-Cluster Differences - Microstates as Attractor Points

It has been suggested that microstates could be conceptualised as attractor states in a multidimensional space (Milz et al., 2017). Existing investigations primarily use parameters to characterise microstates in various contexts, the measures of duration, occurrence, coverage and Global Explained Variance (GEV) being the most prominent (see Section 1.1.3 for an overview of microstate parameters). In the context of microstates as attractor points, whenever a microstate occurs on the time series, that is, whenever the EEG time series is most similar to that microstate for a given amount of time, the duration of the occurring microstate could be conceptualised as how long the EEG time series is "under the influence of" the given attractor point. The EEG time series can be conceptualised as a trajectory in an attractor space. This concept allows us to step back from existing approaches and consider how syntax analysis may be improved.

The first change using this new perspective is what microstate parameters mean in this context. The duration of a microstate is how long the EEG time series is under the influence of the microstate attractor point. Coverage is how much of the overall trajectory is spent under the influence of each attractor state. A one-to-one transition from one microstate to another can be conceptualised as the EEG time series moving from the influence of one attractor point to another. Whilst these points may appear to be obvious, relatively few existing methodologies have considered the relationship between microstate classes and time beyond their average duration (Lehmann et al., 2005; Schlegel et al., 2012; von Wegner et al., 2017).

Additionally, this perspective points out the shortcomings of the "event-mode" (see Section 3.2.3 for the definition of modes) approach employed in studies which do consider microstate syntax. Whilst microstate duration does

show in this context how long the EEG signal is under the influence of a microstate, it does not consider how many GFP peaks occur within that duration. In two separate occurrences of microstate A, which last 100ms, one may contain one GFP peak and the other 3, yet existing methods do not detect this. Using the peak-mode in conjunction with a model of an attractor space may allow for an understanding of "within-microstate" dynamics. That is, if there are two consecutive GFP peaks within the influence of microstate B, for example, it may be the case that *where* within the cluster of microstate B each of those peaks could affect where the trajectory goes next (i.e., it may affect the syntax of the microstates).

Finally, although it is undoubtedly the case that microstates explain a high percentage of variance in GFP peaks (Michel & Koenig, 2018), the use of an attractor space shows the difference between the observed EEG and the microstate attractor points. In the existing winner-takes-all method of microstate sequence labelling, each time point is assigned a single label. In the context of an attractor space, the microstate that each EEG time point is closest to would be the microstate that it is traditionally labelled with. However, there is a possibility for the investigation of each time point's distance from each microstate, meaning that the influence of each microstate can be modelled at each time point rather than using the winner-takes-all approach.

If a methodology were derived which generated a space that the EEG time series occupied and could be analysed as a trajectory, this would allow for a more flexible investigation of the EEG time series. Upon understanding the dynamics of the EEG as a multidimensional trajectory in that space, a post-analysis step of re-labelling the EEG time series with microstates could be used to understand microstate syntax. This would be a practical approach as it allows for investigating the dynamics of microstate transitions without using discrete sequences of states. Other analysis methods are possible with the EEG time series in a space where it is considered a continuous signal. This is somewhat akin to the step used in General Linear Model (GLM) analyses, where microstates are made into continuous regressors using a correlation value (Britz et al., 2010; Musso et al., 2010; Yuan et al., 2012).

3.4 Proposed Methodology for Microstate Syntax Analysis

It is suggested that for the investigation of microstate syntax, first, in the preprocessing stage, epoch removal should be avoided and replaced with noise artefact reconstruction methods, or at the very least, epoch removal is minimised wherever possible.

Then, in the microstate analysis phase, interpolation can be utilised to investigate peak-mode sequences so that within-microstate syntax can be better understood. Additionally, if the back-fitting technique is employed, this is not followed by smoothing, which destroys the syntactic structure of the sequence.

It is suggested that these considerations be utilised in conjunction with symmetrical EEG-fMRI alignment techniques when comparing microstate syntax to

fMRI, as a failure to do so adds complexities to the comparison, and damages the microstate syntactic structure in and of itself.

Finally, it is suggested that a methodology which conceptualises the EEG time series as a trajectory in an attractor space be developed and used in conjunction with syntax analysis methods as a complementary means of understanding EEG signal dynamics.

4 Overview of Studies

Following the identification of a proposed methodology and pipeline for the investigation of Electroencephalography (EEG) microstate syntax with simultaneously recorded Functional Magnetic Resonance Imaging (fMRI) Blood-Oxygenation Level Dependency (BOLD) signal, three studies were conducted which utilise these methodologies, proposing new directions for these methods, whilst also identifying preliminary results with regards to comparisons between EEG microstate syntax and BOLD signal in resting state, eyes-open recordings.

4.1 Data and Recording Paradigm

All studies used the same dataset. Information regarding the capture of this data is as follows.

4.1.1 Participants

Simultaneously acquired EEG and fMRI data was collected from 18 participants (11 males) in a previous study with mean age SD 28.02 +/- 5.44 (age range 20 – 41). Two of the 18 participants were excluded due to incomplete data acquisition. Participants had no history of neurological/psychiatric disorders. All participants were right-handed with either normal or corrected-to-normal vision. All participants gave written informed consent and the study was approved by the Hammersmith Hospital (London, UK) research ethics committee (Fagerholm et al., 2015). All participants underwent an at-rest recording in an awake, eyes-open state lasting 480s. An additional one participant was removed from analysis during processing, showing characteristics of an outlier during microstate analysis. This resulted in fifteen of the total eighteen participants being used for the analysis.

4.1.2 EEG Recording

EEG data were collected simultaneously with fMRI scanning with an electrode cap consisting of 30 electrodes placed according to the extended international 10-20 system (Klem et al., 1999). In addition to EEG, one channel of Electrooculargraphy (EOG) and Electrocardiography (ECG) were recorded respectively. Clock synchronisation hardware between EEG and fMRI ensured simultaneous recording (Fagerholm et al., 2015). The EEG was sampled at a rate of 5 kHz.

4.1.3 fMRI Recording

Images were acquired using a 3.0 Tesla Magnetic Resonance Imaging (MRI) scanner. A whole-brain Echo-Planar Imaging (EPI) sequence was used (T2*-weighted gradient-echo). Voxel size was $3.00 \times 3.00 \times 3.00mm$, with a field of view of $192 \times 192 \times 105mm$ and flip angle of 80°. The repetition time/echo time (Repetition Time (TR)/Echo Time (TE)) ratio was 2000/30 ms, using 35

ascending slices with 3.00 mm thickness. Whole-brain structural images (T1-weighted) were also acquired for all participants (Fagerholm et al., 2015).

4.2 Code Availability

Code used for all studies in the project is available via the link here. Additional figures for results can be found there also, in the "Results" section of each study.

4.3 Plan of Investigation

4.3.1 Study 1 (Chapter 5)

I first derive data-driven microstates in resting state EEG data, reporting their observed parameters. I then derived microstate n-grams at various lengths in the event-mode and parameters for each length. I then derive fMRI Co-Activation Pattern (CAP)s, and after alignment of the recording modalities, conduct statistical analyses which show differences in microstate n-gram parameters between fMRI CAPs, which are simultaneously active. The objective is to show preliminary results that demonstrate the microstate syntax is related to fMRI signal and to demonstrate that the methodological developments proposed in Chapters 2 and 3 alleviate some shortcomings of previous studies.

4.3.2 Study 2 (Chapter 6)

I then go on to derive an fMRI gradient space to place the observed resting state fMRI BOLD signal in order to reduce the dimensionality of the signal without reducing TRs into categories as in Study 1. This results in a distribution of microstate parameters across the fMRI space for each participant individually. The objective is to show more precisely how the microstate syntax is associated with the fMRI signal, with the potential to associate n-gram parameter changes with more precise cognitive function using the fMRI gradient space.

4.3.3 Study 3 (Chapter 7)

Finally, I derive a preliminary data-driven EEG gradient space, and following alignment with the fMRI signal, attempt to predict co-occurring fMRI CAPs based on the EEG gradient trajectory, using deep-learning techniques. Microstates and microstate n-grams could be used as categories for post-analysis description rather than as the object of analysis. This approach may better associate general patterns of microstate syntax with function, developing upon previous methodologies which have associated individual microstates with cognitive functions.

5 Study 1 - Microstate n-Gram Parameters during fMRI "Macro-states"

Past investigations into associations between Electroencephalography (EEG) microstates and cognitive states have provided a general overview of the relationship between microstates and mental activity. By some, this has been done by associating microstates with Functional Magnetic Resonance Imaging (fMRI) networks and leveraging those networks' knowledge to understand microstate function better. Some studies have investigated how microstate sequences may be related to different cognitive states through cognitive manipulation approaches, but there has yet to be a methodology that attempts to associate microstate sequences with fMRI states. Here, I derive EEG microstate sequences and align them to simultaneously occurring fMRI Co-Activation Pattern (CAP)s, using CAPs as labels to categorise sub-sequences of microstates, referred to as n-grams. I demonstrate apparent differences in microstate n-gram parameters between fMRI CAP labels that would not be expected in random labelling.

5.1 Introduction

EEG activity can be described by quasi-stable epochs of scalp potential topography lasting approximately $60 - 120$ms before rapidly changing to another topographical pattern of quasi-stability. Each epoch is called an *EEG microstate* (Lehmann et al., 1987). Microstates may reflect states of global neuronal activity, and hence transitions between these topographies may reflect dynamic, synchronous patterns of global neuronal activity (Antonova et al., 2022; Nehaniv & Antonova, 2017).

There are several scalp topographies which have been identified as consistent across individuals. The "canonical" topographies of A, B, C and D have been identified in multiple studies, explaining approximately $70 - 80\%$ of variance in the data (Michel & Koenig, 2018). These states are highly replicable across multiple studies (Khanna et al., 2014). Other studies which have used data-driven approaches to derive microstates that better fit specific datasets have been clustered across studies and identified as "meta-microstates" outside the canonical set (Koenig et al., 2023). They are referred to as meta-microstates due to their derivation from a meta-analysis that investigated microstates across multiple studies.

The standard parameters used in many studies are *duration*, the average amount of time that the given microstate remains stable; *occurrence*, the number of times that a microstate appears within 1s, regardless of the duration of each of those microstates, and *coverage*, the fraction of the time series that the given microstate is dominant.

Additionally, there have been preliminary investigations into the *syntax* of microstates, including transition probabilities (Bréchet et al., 2019) and the frequency of microstate sequences in the time series (Lehmann et al., 2005; Schlegel et al., 2012).

Despite the canonical microstates being replicable across multiple studies, the functional significance of microstates remains to be determined. Cognitive manipulation studies have attempted to elucidate function, generally associating microstate A with inhibition of language processing, microstate B with inhibition of visual processing, microstate C with self-relevant thought, and microstate D with attention orientation (Antonova et al., 2022; Sten et al., 2017).

However, many of these associations are heavily debated, with different studies reporting conflicting findings (Antonova et al., 2022; Bréchet et al., 2019; Milz et al., 2017; Seitzman et al., 2017). Furthermore, all EEG microstates are active during all cognitive states, so the association of each microstate to an individual process may be an oversimplified explanation of their function (see Section 1.1.5 for a more in-depth description of speculated microstate function).

Aspects of microstate syntax, such as transitions between microstates, differ between populations (Lehmann et al., 2005; Schlegel et al., 2012). It has been shown that specific microstate transitions are significantly more likely to occur in different task conditions (Bréchet et al., 2019). Whether there is a difference in microstate syntax between simultaneously occurring intrinsic brain networks has not been investigated. Therefore, investigating how microstate sequences

differ between intrinsic brain networks can clarify microstate function.

fMRI has been used in past studies to relate EEG microstates to intrinsic brain networks using voxel-wise regression models in an attempt to derive microstate function (Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012). However, this approach has two key issues: [6]:

1. EEG and fMRI are recorded at different sample frequencies, with EEG microstates generally down-sampled to 250Hz (Khanna et al., 2014; Michel & Koenig, 2018) and taking an example of fMRI being recorded with a standard Repetition Time (TR) of 2 seconds means a sampling frequency 0.5Hz. The EEG signal, which must be convolved with the Haemodynamic Response Function (HRF), is often down-sampled to the same resolution as the fMRI (Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012). This dismantles any relation to time that the microstates may have, including the complexity of their syntax, and does not utilise the high *temporal* resolution of EEG effectively.

2. The fMRI states that have been compared to the voxel-wise EEG maps used in regression studies have historically been networks derived from spatial Independent Component (IC)s (Britz et al., 2010; Xu et al., 2020); static networks which are not temporally independent from one another, but spatially independent from one another. When investigating differences in EEG microstates, which are derived as temporally *independent*, comparing them to fMRI states, which are temporally *dependent* does not utilise the high *spatial* resolution of the fMRI effectively. Associating each microstate with spatially independent fMRI states also assumes that one fMRI state cannot be associated with multiple EEG microstates or, indeed, that a single brain region can contribute to more than one microstate.

Therefore, it may be more beneficial to use a methodology which considers the dynamics of EEG microstates *and* fMRI states simultaneously to understand better the function of EEG microstates.

A recent study used dynamic Functional Connectivity (dFC) fMRI states in relation to EEG microstates using simultaneously recorded EEG-fMRI (Abreu et al., 2021). The study aimed to predict the current dFC state given the microstates that co-occurred alongside it. This highlighted that the occurrence of all canonical microstates correlated with all dFC states, but each microstate contributed a different proportion of explained variance to the prediction of each dFC state.

It is apparent that the relationship between intrinsic brain networks and EEG microstates is more complex than previously indicated one-to-one relationships (Britz et al., 2010; Custo et al., 2017; Musso et al., 2010). However, the sliding window approach used in this study limits the possibility for robust investigation of microstate syntax, since the sequences of microstates that could

[6] See Chapter 2 for a discussion of this area in more detail.

be investigated are bound by the arbitrary size of the sliding window, rather than a functionally relevant boundary.

Here, I aim to investigate whether there are distinct differences between microstate parameters in the resting state when microstates are active during different simultaneously occurring dynamic fMRI states. To this end, I employ CAPs as the dynamic fMRI states, which are derived by clustering in a frame-wise fashion to retain global activity patterns that occur across the time series; a process equivalent to the process of derivation of EEG microstates (Gutierrez-Barragan et al., 2019; Pascual-Marqui et al., 1995).

Past studies have yet to understand the neural origins of CAPs. Dynamic fMRI states, in general, likely reflect the continuous reorganisation of brain networks in response to internal and external stimuli (Calhoun & Adali, 2016; Hutchison et al., 2013). Whilst EEG microstates have not been directly linked to neuronal origins, associating CAPs and EEG microstate syntax to one another may assist in understanding the neural origins of CAPs and microstates simultaneously.

Here, I collected parameters of short EEG microstate sequences (referred to as microstate n-grams) that co-occurred with each of the derived CAPs to investigate whether microstate syntax differed between fMRI states. I demonstrate that the parameters of microstate n-grams are significantly different between some pairs of CAPs.

These findings suggest that if applying fMRI to the investigation of EEG microstates, future work should use dynamic rather than static fMRI networks and voxel-wise regression models should be used cautiously. It is suggested that the methodology proposed be developed by first analysing the continuous signals of EEG and fMRI before reducing them to discrete sequences of states. This could allow for more flexible analysis that may be investigated with state sequences post hoc.

5.2 Methodology

5.2.1 EEG Preprocessing

All EEG data were processed using the *EEGLAB* toolbox (Delorme & Makeig, 2004). The average potential across all electrodes was calculated and used as a reference for all other electrodes for each participant individually (Croft & Barry, 2000). The data were down-sampled from their recorded sample frequency of 5kHz to 250Hz. The *FMRIB* suite was then used to remove noise in the EEG signal caused by the scanner (Iannetti et al., 2005; Niazy et al., 2005). Magnetic Resonance (MR) noise was first removed from the EEG signal using the *FASTR* artefact slice removal template. Using the simultaneously recorded Electrocardiography (ECG) channel as a reference, the heartbeat detection algorithm was used to detect and remove irregularities in the EEG signal that correlated with observed ECG peaks.

The dataset contained an unusual noise component caused by vibrations from the scanner. The vibration caused a consistent frequency artefact through-

out the EEG time series at approximately 17Hz. Using *EEGLAB CleanLine* (Delorme & Makeig, 2004), which estimates and removes sinusoidal artefacts for each channel using frequency domain regression techniques (Mitra & Bokil, 2007), sinusoidal signal of excessive amplitude was regressed from the time series with a sliding window of 2 seconds across individual time series.

In order to minimise the number of epochs removed due to bursts of noise, Artefact Subspace Reconstruction (ASR) was used to clean epochs of the data that the algorithm labelled as noisy (Miyakoshi et al., 2020).

The signal was bandpass filtered between 2 and 20Hz to isolate alpha band activity before being subject to Independent Component Analysis (ICA). The analysis was applied at the subject level, and ICs were removed from participants using the *ICLabel* algorithm, which cross-references IC topographies with a database of existing IC topographies which have been labelled as genuine brain signal or noise. The algorithm measures the likelihood that a given EEG topography has its source from brain signal, eye movement, heartbeat, muscles, or noise. All ICs with a brain correlation less than 0.7 were assumed to be noise and removed from the signal (Khanna et al., 2014; Milz et al., 2017).

5.2.2 EEG Microstate Analysis

EEG microstates were generated using a data-driven technique with application of the *EEGLAB* microstate plug-in *MST* (Poulsen et al., 2018). For each participant, the EEG topographies at time points of maximum Global Field Power (GFP) (Skrandies, 1990) were collected as input into a modified k-means clustering algorithm, applied here to classify EEG time points based on spatial similarity. The modified version of k-means differs from the traditional algorithm in that the polarity of channels is ignored to derive the microstates (Pascual-Marqui et al., 1995).

All GFP peaks over the defined threshold of one standard deviation above the mean across the fifteen participants were applied to k-means up to a maximum of 500 iterations, with 100 repetitions of the given number of clusters. The run with the highest explained variance was used as the result at each k. The process was repeated with 3, 4, 5, 6, 7 and 8 clusters. This range was used to keep the number of microstates down for n-gram analysis without using so few microstates that variation across classes was an issue.

The number of clusters to use was chosen using Global Explained Variance (GEV) (defined in Equation 2) and the Cross-Validation (CV) criterion (Pascual-Marqui et al., 1995). The CV criterion calculates an estimator of the variance of residual noise in the fit (Pascual-Marqui et al., 1995). The resulting cluster centroids were then assessed on their physiological feasibility.

After choosing the number of clusters, the GFP peaks were labelled with their associated microstate, and these labels were placed at the corresponding GFP index on the time series. The interpolation approach was used to determine the transition points between microstates (see Section 3.2.2). The centre index between each peak index was taken. A microstate's duration was then defined as the number of time points between the before and after the centre point of

the given GFP peak. The first and last GFP peak in each participant time series were excluded since the centre point of their before and after peaks could not be accounted for. The standard parameters were calculated for each microstate class (see formal definition of duration in Equation 3, coverage in Equation 5 and GEV in Equation 2).

5.2.3 EEG Microstate n-gram Derivation and Parameter Calculation

An n-gram is defined as a sequence of states of a given length n. Here, we derived microstates using event-mode analysis (see Section 3.2.3 for an explanation of temporal modes of microstate derivation. Comparison between event- and peak-mode can be found in Chapter 6). These n-grams were generated for $n = 1$, 2, 3, 4 and 5.

The mean duration of each possible n-gram was calculated. The duration of a single n-gram occurrence is the sum of the durations of the constituent microstates of the n-gram. After the calculation of each n-grams duration, the mean duration of each n-gram is calculated (i.e., mean duration of ABC, of BCD, etc).

Coverage is the proportion of the time series that each microstate is active. Since microstate n-grams overlap with one another in the time series when $n > 1$ (Figure 3), the proportion of the time series that each sequence takes up is dependent upon the n-grams before and after it. This makes the calculation of an equivalent percentage for n-grams of length 2 or more infeasible.

Instead, I define a novel parameter - *frequency*. Microstate n-gram frequency is defined as:

$$f_i^n = \frac{o_i^n}{\sum_{i=1}^{I} o_i^n} \tag{6}$$

and:

$$\sum_{i=1}^{I} f_i^n = 1 \tag{7}$$

where f_i^n is the frequency of the given n-gram i of length n, o_i^n is the number of occurrences of the given n-gram i of length n. I is the total number of n-grams at the given length n.

The sum of frequencies at a given n-gram length is equal to 1. For example, if the entire sequence of microstate observations is $ABCDE$, the possible number of occurrences of an event-mode 3-gram would be 3; ABC, BCD, and CDE. ABC, BCD and CDE would have one occurrence each over the total of three possible occurrences and hence would have a frequency of 33.3% each.

This definition of n-gram frequency accommodates the length of the n-gram considered but should not be considered equivalent to the traditional microstate coverage, as its derivation makes it independent from microstate duration, which microstate coverage is not. Instead, frequency should be considered a normalised

n-gram occurrence, returning a percentage for the whole time series sample whilst also being independent of duration.

5.2.4 FMRI Preprocessing

Processing of fMRI was carried out using *FEAT* version 6.00, part of the *FSL* software library (Jenkinson et al., 2012) and visually checked using *FSLeyes*.

First, brain and non-brain tissues were identified in the images using the brain extraction tool (*BET*) (Smith, 2002). This algorithm identifies voxels associated with the brain by generating a surface mesh. The mesh is updated through an iterative process, which begins as a sphere in the centre of the brain volume and is expanded and refined to fit the shape of the brain with each iteration.

MCFLIRT is next applied (Jenkinson et al., 2002), a registration and motion correction tool which aligns the 3D images of each time point across a participant's time series to ensure that a single voxel location implies the same anatomical location across participants. To do this, between Echo-Planar Imaging (EPI) scans across TRs, *MCFLIRT* minimises dissimilarity by a set of transformations. The set of transformations used here can be rigid-body transformations (3 rotations and 3 translations) or affine transformations (the 6 rigid-body transformations plus 3 or 6 additional transformations that allow for linear non-rigid-body manipulations), which are referred to as Degrees Of Freedom (DOF) also defined by the user. Here 12 DOF were used.

Applying this alignment between EPI images corrects for motion across the time series of the given participant. The EPI scans are then aligned with their structural data (T1 image) for registration. The structural image was registered to the MNI152 2mm standard space so participants could be compared.

Spatial smoothing was then applied. This process assigns each voxel a weighted average of its neighbouring voxels, which is weighted based on its closeness to the voxel of interest. This was applied using a Gaussian kernel of Full-Width-at-Half-Maximum (FWHM) 5mm. The process reduces the spatial resolution of the data but increases Signal-to-Noise Ratio (SNR), and a minimum smoothness is required for Gaussian random field theory to apply to analysis, a prior that must be adhered to for processing to be valid.

A high-pass temporal filter was applied at 0.01Hz with a weighted least squares regression line and sigma equal to 50s to remove low-frequency components commonly associated with signal drift. The *ICA-AROMA* package was applied to reduce noise further (Pruim et al., 2015). The algorithm applies the it FSL package *MELODIC*; an ICA, and detects components which would be considered noise based on comparison to known motion artefacts. This was applied at the participant level to avoid generalising noise across the group.

A group mask was quantified to ensure homogeneity across participants. With all participants in the MNI152 space, voxels that were not active in every participant and were not in the standard brain were removed across the group. The ventricles of each participant's signal were also removed using the MNI152 template. The global signal of each participant's time series was calculated as

the mean Blood-Oxygenation Level Dependency (BOLD) signal across voxels at each TR, for each participant individually. The global signal was removed from the fMRI time series for the following CAP analysis. The global signal is the mean BOLD signal across all voxels at each time point. [7]

Finally, a parcellation was applied to the data for particular analyses. Parcellation takes the average activity across a set of voxels labelled with the same region and assigns the mean value across voxels to the region rather than each voxel retaining its observed value. The Schaeffer-1000 cortex parcellation was used here (Schaefer et al., 2018), meaning that the CAPs derived here only include the cortex.

5.2.5 EEG Microstate fMRI General Linear Model

To compare previously applied methods to those developed here, the General Linear Model (GLM) approach described in Section 2.2.1 was applied to the EEG microstate time series and fMRI signal.

Each microstate map was correlated with the EEG topography at each time point to establish a microstate correlation time series for each microstate. Using *FSL's FEAT* analysis (Jenkinson et al., 2012), each correlation time series was then down-sampled to 0.5Hz to match the fMRI recording of TR of 2 seconds. The design matrix of the GLM was set up with each down-sampled microstate correlation time series as regressors, aiming to model the BOLD responses associated with each microstate map. Each regressor was convolved with the canonical double gamma HRF in *FEAT* with a peak at 6 seconds.

Statistical maps were obtained using t-tests and corrected for multiple comparisons using False Discovery Rate (FDR) at $p = 0.05$.

5.2.6 FMRI Co-Activation Pattern Derivation

Co-activation patterns were generated by applying a standard k-means clustering algorithm to the time series of fMRI using the python packages *scikit-learn* (Pedregosa et al., 2011) and *nilearn* (Abraham et al., 2014).

The individual participants were concatenated into a single time series. Each participant's 4D time series was then reduced into a 2D matrix, the size of which was the number of parcels by the number of time points across participants. Clustering was applied in the temporal dimension, using a maximum of 1000 iterations for 100 repetitions. The best fit was identified for 6, 8 and 10 clusters. Only even numbers are considered due to the typical activation/attenuation CAP pairs (Liu et al., 2018).

Upon generation of cluster centres (CAPs), a matrix of spatial similarity between CAP volumes was then applied to the *Munkres* algorithm (Munkres, 1957). This combinatorial optimisation algorithm was used to assign pairs of

[7] For ease of understanding in any descriptions of analysis throughout the thesis, the global signal should be assumed as included in fMRI time series unless specifically stated as it is here.

CAPs based on the highest level of dissimilarity. Each TR was then labelled with the cluster it belonged to to generate a sequence of CAPs for each participant.

5.2.7 Microstate and Co-Activation Pattern Sequence Alignment

The EEG and fMRI time series are aligned by their recording timings in the data acquisition phase (see Section 4.1). Following sequence generation of both EEG microstates fMRI CAPs, the microstate sequence was offset against the CAP sequences by 6 seconds; the time to the peak of the canonical double-gamma HRF (Buckner, 1998). This offset is used to align the haemodynamic response expected from the EEG microstate activity with the observed BOLD response of the fMRI.

Any EEG that was recorded before and after the aligned fMRI was removed from the EEG microstate sequence since it could not be compared to a corresponding TR. This process is repeated for each participant individually.

5.2.8 Obtaining Microstate Sub-Sequences per Co-occurring fMRI TR

In order to compare the EEG microstate time series to each CAP, the microstate sequence which occurred during each CAP was isolated through a "cutting" process. Figure 9 shows the process. Figure 9A shows the alignment of each modalities sequence, with 9B showing the cutting process described here.

Each CAP which occurs in the time series could potentially be multiple TRs long, so if the study aims to compare how microstate sequences differ between CAPs, it is not suitable to investigate microstates and microstate n-grams per TR, as this can create arbitrary separation of microstate sequence within a single occurrence of a CAP. Therefore, the microstate sequence is cut only at the transitions between CAPs.

At transition boundaries from one CAP to the next, it is rare for a microstate to end precisely at the boundary. When an EEG microstate occurs across the boundary of transition between CAPs, rather than treat that microstate as two separate shorter microstates, the boundaries between CAPs were altered slightly. Whichever CAP a boundary microstate occurred the longest in, its whole duration was taken as part of that CAP. If the microstate was in both CAPs for an equal amount of time, the microstate was assigned to the pre-transition CAP.

Finally, once microstate sequences were aligned and cut for each CAP class, the microstate and n-gram parameters were calculated using the n-gram parameters within the CAP (Figure 9C). This was done for n-grams length 1 to 5 in event-mode.

5.2.9 Comparison of Microstate n-gram Parameters between Co-occurring fMRI CAPs

Differences in n-gram parameters between CAPs were calculated. Using frequency as an example parameter, differences were measured by calculating the

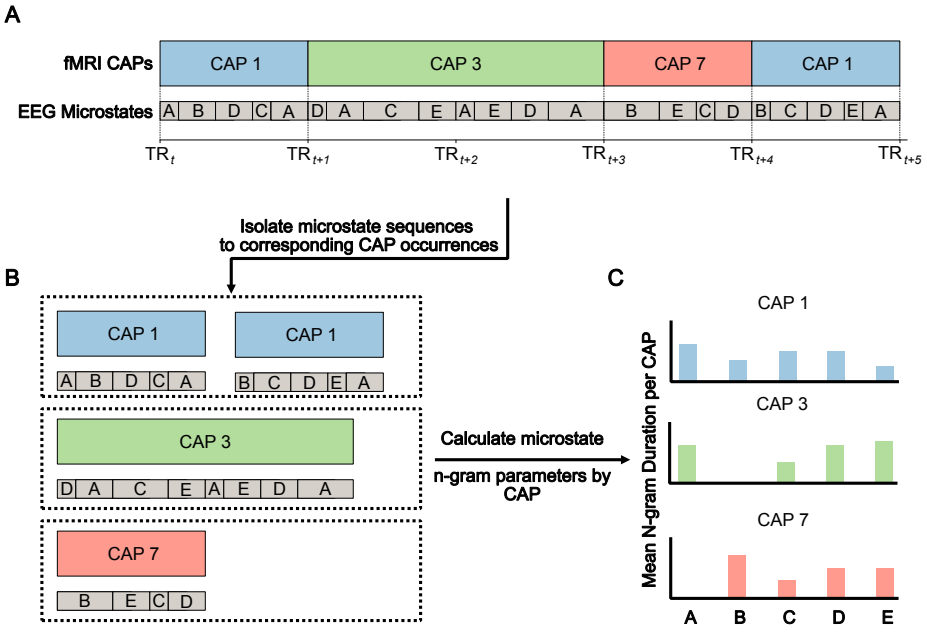

Figure 9: Pipeline Schematic of alignment and cutting process between EEG Microstate sequences and fMRI CAPs. The time series of EEG and fMRI are aligned using the recording offset and the time to the peak of the HRF (A). The microstates which occur during a given CAP are cut from the time series and isolated (B). The microstates and n-grams which occur simultaneously with each CAP state type then have their parameters calculated across all the state type occurrences to obtain a summary of microstate and n-gram parameters for the given CAP type (C).

L_1 distance between the distribution of n-gram frequencies during each CAP:

$$L_1(x-y) = ||\mathbf{f}_x^n - \mathbf{f}_y^n||_T \tag{8}$$

where $\mathbf{f}_x^n$ and $\mathbf{f}_y^n$ are vectors of frequencies of length I for the given n-gram length during CAPs x and y respectively. The n-grams compared between CAPs must be of the same length and mode. For example, if $n = 1$ and the canonical microstates are used, the vector $\mathbf{f}_x^1 = (f_A^1, f_B^1, f_C^1, f_D^1)_x$ holds the frequency of each microstate during the CAP x. $L_1(x-y)$ is the distance between the microstate n-gram frequencies that occur during CAP x and y. The distance was calculated for all pairs of CAPs, and was applied at all n-gram lengths individually, 1 to 5 for all parameters.

When considering the mean duration of n-grams, if an n-gram did not occur in one CAP but did occur in another, taking the non-occurrences mean duration as zero would have resulted in a huge difference between n-grams, when in fact, should that n-gram have occurred in both CAPs (which it may have done had

more data been recorded), the difference would have been much smaller. As such, when the n-gram did not occur in a given CAP, its mean duration was replaced by the mean duration of that n-gram across the whole time series of the given participant. Hence, when the L_1 distance between the CAPs was calculated, the L_1 distance would include how far away the n-grams in the occurring CAP were from the global mean. If the n-gram did not occur anywhere in the time series, replacement was unnecessary since the difference between CAPs would be zero. This process was only necessary for mean duration, as a frequency of zero would not cause such an offset, and coverage can only be calculated in $n = 1$ where non-occurrences would not happen.

5.2.10 Microstate n-Gram Parameter Differences between CAP Labels as an Undirected Graph

Following the generation of a matrix of pairwise L_1 distances between CAPs, the CAP labels were conceptualised as nodes in a network, with L_1 distances being weights on edges that connect each label. A sparse matrix was generated in each case by setting an edge weight threshold and removing any edges where the weight was less than the threshold to establish meaningful connections.

The threshold was defined as a matrix's mean L_1 distance plus one standard deviation. The corresponding edge was removed if an observed weight was less than this value. The original all-to-all connected graph was hence reduced to a sparse representation of the largest L_1 distances between CAPs. The degree centrality of each node in the sparse graph was recorded (total number of remaining connections), and the sum of degree centrality across nodes for the graph was also recorded. This process was repeated for each participant individually, for each n-gram length, for each parameter.

5.2.11 Comparison of Distances between CAPs against a Data-Driven Null Distribution

In order to establish whether the differences in parameters between CAPs were meaningful, a null distribution of comparison L_1 distances was generated to be compared against the observed L_1 distance between CAPs, as shown in Figure 10.

The order of CAP labels was first shuffled on the time series, with the EEG microstate sequence kept the same, hence changing the CAP labels relationship to microstate labels. The parameters mean duration, frequency (and coverage for $n = 1$) were calculated within each of the shuffled CAP labels as described in Section 5.2.9, from $n = 1$ to 5 (Figure 10B). The L_1 distances between pairwise CAPs of a single shuffle were then calculated. The process outlined in the previous Section regarding the derivation of sparse graphs was also carried out on each shuffle. The threshold used for the shuffle was the same as the observed threshold for the given parameter at the given n-gram length for the given participant, i.e., the threshold in the observed data for frequency in participant 1 for $n = 2$, was used in the shuffled CAPs for the same parameter, participant

and n-gram length. Again, once the sparse graph was created for a shuffle, the degree centrality of each node was recorded, and the total sum of degree centrality for the sparse graph was also recorded.

This process was repeated in each case P times (in this case $P = 1000$). The set of P sum degree centrality of a sparse graph was used as a null distribution and compared against the total degree centrality in the same graph in the observed data to evaluate the significance of labelling microstate sequences with the observed CAP sequence (Figure 10C). More connections in the observed sparse graph than in the graphs generated from shuffled CAP data would suggest that differences between n-gram parameters in each CAP are not random.

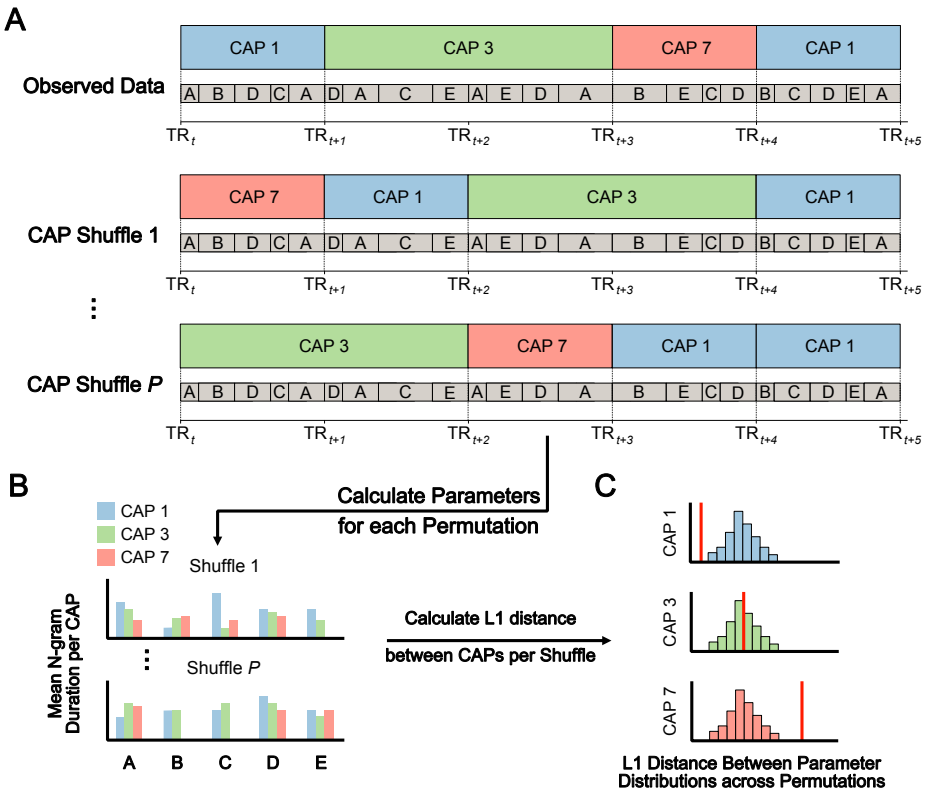

Figure 10: Pipeline Schematic of alignment and cutting process between EEG Microstate sequences and fMRI CAPs. The time series of EEG and fMRI are aligned using the recording offset and the time to the peak of the HRF (A). The microstates which occur during a given CAP are cut from the time series and isolated (B). The microstates and n-grams which occur simultaneously with each CAP state type then have their parameters calculated across all the state type occurrences to obtain a summary of microstate and n-gram parameters for the given CAP type (C).

A p-value was computed to represent the fraction of shuffled CAP labels

where a total distance between CAP n-gram parameter distributions is equivalent to the original data. This was calculated as:

$$p = \frac{X + 1}{P + 1} \tag{9}$$

where X is the number of permutations which had a total number of connections in their sparse graph that was greater than the number in the observed data (one-tailed test), this process was repeated for each parameter at each n-gram length.

Benjamini-Hochberg FDR (Benjamini & Yekutieli, 2005) was applied for each participant individually across n-gram lengths to correct for multiple comparisons. This process was repeated for all calculated parameters.

5.3 Results

5.3.1 EEG Microstate Derivation

The EEG microstates were derived from $k = 3$ to $k = 8$. The topographies for each k are shown in Figure 11. The left side of panel A shows the change in GEV and CV as k increases. From $k = 3$ to $k = 5$ GEV steadily increased to 69.5% with five microstates. Beyond five showed a plateau around approximately 70%. CV showed a plateau from three to five microstates, approximately 50% at three, up to 51% at five microstates (Figure 11A). Onward from five, a steady increase is observed, with 8 microstates measuring 61% CV. Since the aim is to maximise GEV and minimise CV, five microstates were chosen as the optimal number of microstates for this dataset.

The chosen microstates are shown in Figure 11B, where the microstates have been labelled with their corresponding meta-microstate labels (Koenig et al., 2023). The meta-microstate labels are derived from a meta-analysis which clustered the microstates generated across multiple studies (see Section 1.1.2 for information).

5.3.2 EEG Microstate Parameters

The mean duration, frequency, coverage and GEV of each microstate across the 15 participants is shown in Figure 12. The mean duration of each microstate of each participant is given in Figure 12A. All microstate group mean durations were in the range of $70 - 80$ms. The highest was microstate D at 79.32ms, and the lowest was C at 70.64ms.

Frequency is proposed here as a novel parameter independent of duration and can be considered a normalised occurrence (defined in Equation 6). Group frequency was the highest for microstate D at 23.3%, and lowest at C at 15.4%. Figure 12B shows the distribution of each microstate frequency across the group.

Group coverage was highest for microstate D, at 24.4%, and C again showed the lowest value at 14.8% (Figure 12C). Group GEV was highest for microstate F at 16.6%, followed by B at 15.02%. Then microstates A and D, which showed group GEV of 12.4% and 12.2% respectively. Microstate C was the lowest at

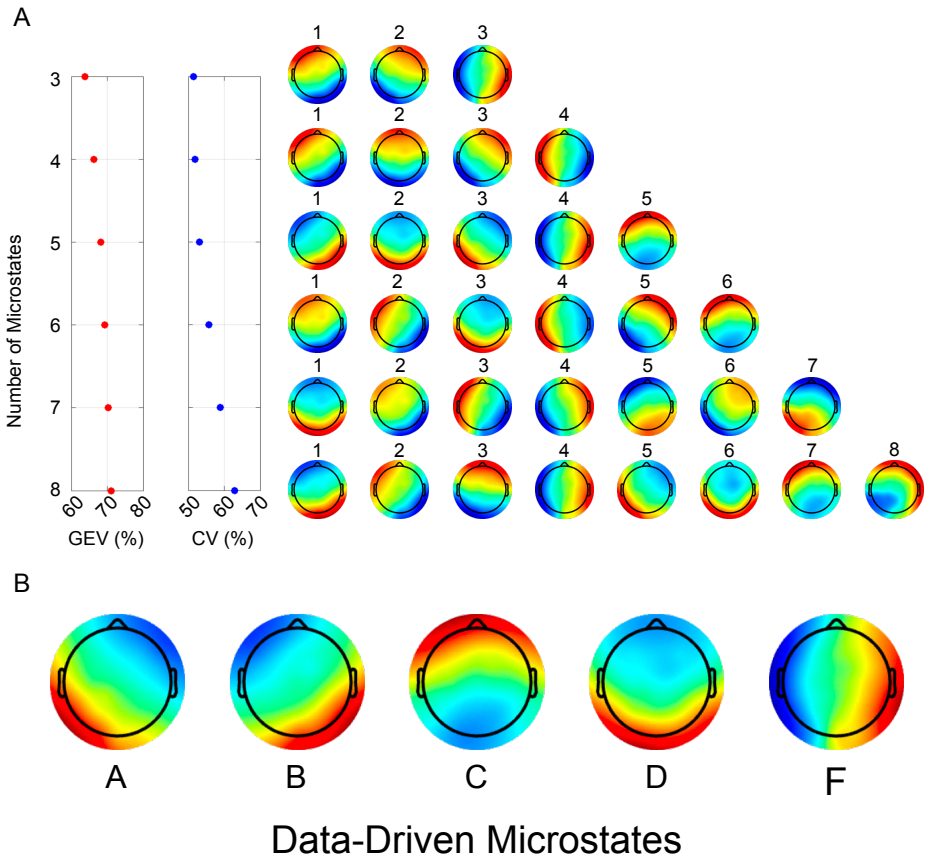

Figure 11: Selection of number of microstates. (A) Left side sub-figures give the GEV and CV criterion. Output of multiple runs of microstates are given on right side of figure from 3 to 8 clusters. Microstates are numbered and ordered for each number of microstates by the amount of explained variance that each contributes. (B) The five data-driven microstates chosen after k-means clustering from $k = 3$ to $k = 8$. The states are re-ordered and labelled with meta-microstate labels (Koenig et al., 2023).

6.9% (Figure 12D). Consistently microstate D showed the largest parameter values, and microstate C showed the smallest.

5.3.3 EEG Microstate fMRI General Linear Model

The statistical maps for each of the microstates across the group are found in Figure 13. Each microstate shows an observed z-map which summarises the voxels that correlated with the given microstate sufficiently across all participants.

Microstate A shows activation in the right cingulate gyrus and the right

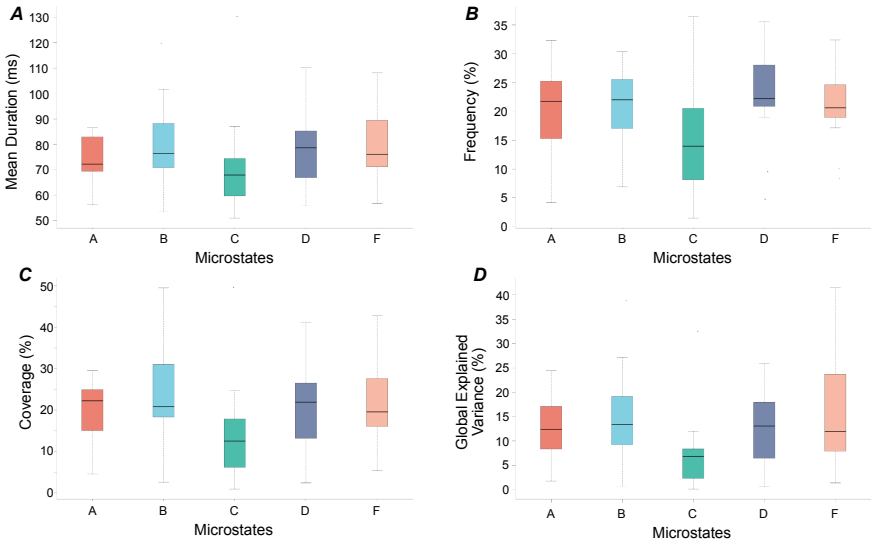

Figure 12: Box plots of microstate parameters across participants. Box indicates interquartile range across participants, whiskers show maximum and minimum participant values. Mid-line indicates median of participants. Parameters calculated are mean duration (A), frequency (B), coverage (C) and GEV (D).

supplementary motor area. Activation is also found on the left and right superior temporal gyrus and the left and right frontal pole. Microstate B shows a very similar activation pattern, which may indicate that these activity patterns are due to noise. Microstate C showed strong correlations across the group with a small region on the border of the right inferior and middle frontal gyrus and the frontal pole. Microstate D shows a group correlation with the right amygdala and the left occipital pole. Finally, microstate F showed group correlations with the right temporal occipital fusiform cortex and the left occipital pole.

The output of each participant's z-map can be found here. In general microstates A and B showed the most voxel correlations, but those correlations were not consistent between participants. Regions of correlation varied between participants. In some participants, microstate A showed correlations with temporal cortices, whereas in others, it correlated with the visual cortex, the Posterior Cingulate Cortex (PCC) and frontal regions. Similarly, microstate B showed correlations in some participants with the PCC and frontal regions, sub-cortical structures in one participant and visual regions in others.

Microstate C showed visual cortex correlations in one participant, the supramarginal gyrus in three participants, and PCC correlations in one participant. Microstate D showed two participants with correlations with the Default Mode Network (DMN), one participant correlating with visual regions, and two with sub-cortical structures. Microstate F showed correlations in five participants

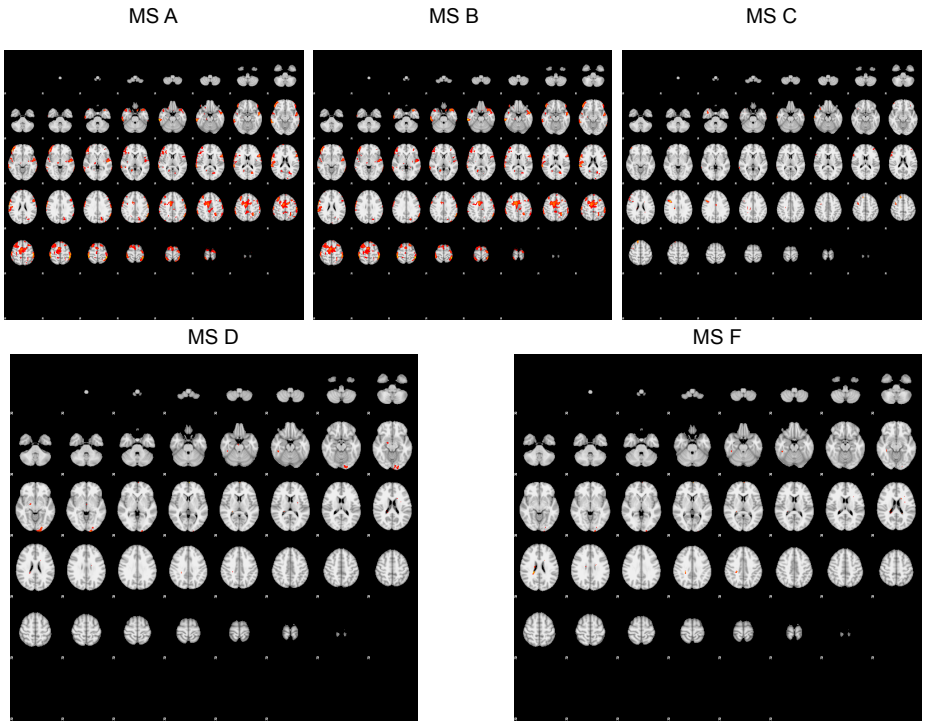

Figure 13: Output of group GLM z-maps for each microstate on the MNI152 standard brain, generated using the down-sampled time series of the EEG microstates as regressors. Red voxels highlight those under p-value threshold of 0.05.

with sub-cortical regions and the cerebellum in two participants.

5.3.4 fMRI CAPs and Parameters

Six CAP pairs are shown in Figure 14. Due to the nature of the approach, the number of CAPs was to be kept to a minimum (see Section 5.2.6). Clustering was run for 6,8 and 10 CAPs. Six demonstrated the highest level of functional feasibility and so were selected for further analysis. Each row of Figure 14 shows the pairings of anti-correlated CAPs (visualisations of 8 and 10 CAPs can be found here).

The parameters of each CAPs were also calculated for each participant, shown in Figure 15. The average dwell time of each CAP (average number of consecutive labels, equivalent to microstate duration) across participants is given in Figure 15A. The pair CAPs 6 and 5 had the first and second highest mean dwell time across the group at 2.50 and 2.14 TRs, respectively. CAPs 4 and 2 had the lowest average dwell times, at 1.66 and 1.64 TRs, respectively.

The coverage of each CAP across each participant is also given in Figure 15B.

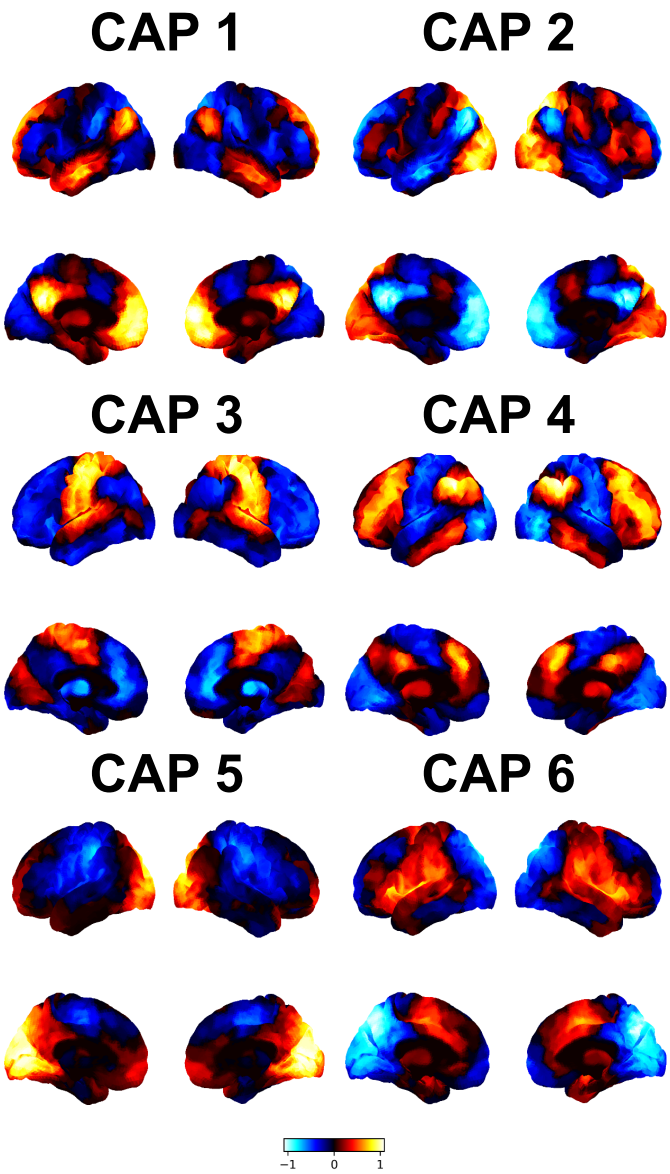

Figure 14: Six Co-Activation Patterns from clustering of the resting-state eyes-open data across participants. CAPs are numbered in their pairs, with each row showing a pair. Brighter red-yellow values indicate higher activation, brighter blue values indicate higher deactivation. CAPs are represented on surfaces with a colour range of -1 to 1 for visualisation purposes. Ll and Lm are left hemisphere lateral and medial views respectively. Rl and Rm are right hemisphere lateral and medial views respectively.

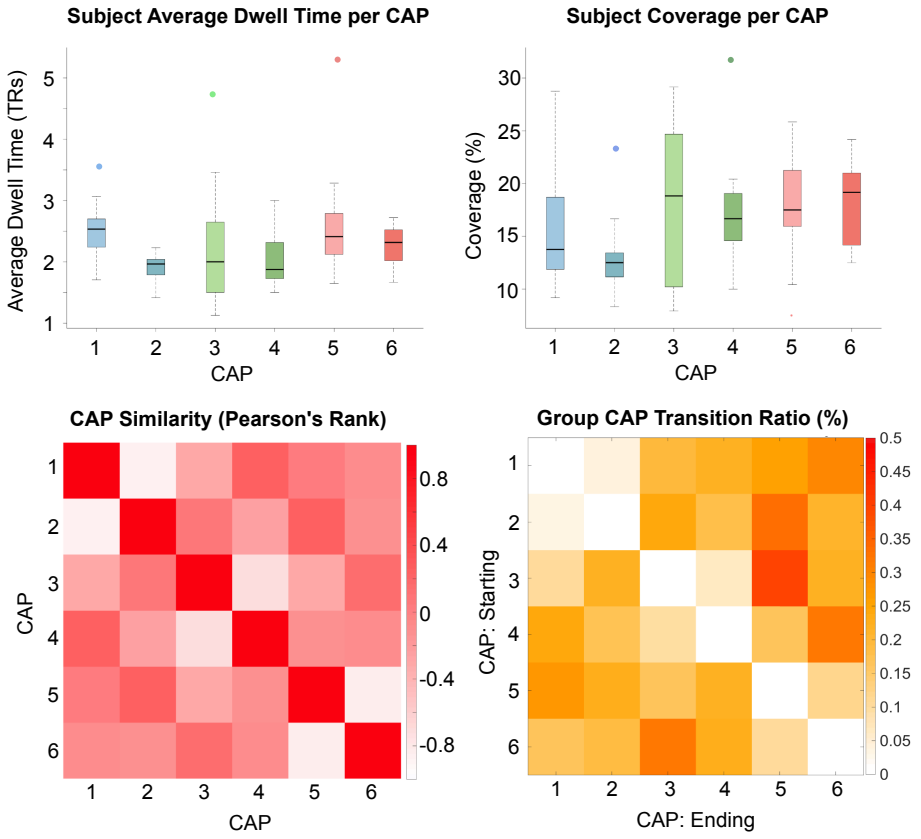

Figure 15: Parameters of each of the six derived fMRI CAPs. (A) shows the average dwell time of each CAP by participant in box plots. Dots which are not included in the plots are participants which were greater or less than the mean of the group, plus or minus 1.5 times the interquartile range. Note the averages given in the text are of the whole group and include these values. (B) shows the coverage of each CAP across the time series by participant. Box plots are represented in a equivalent way to in (A). Each CAP pair uses a similar colour in both (A) and (B) to indicate their pairing. (C) shows the Pearson's Rank Correlation of each CAP, white showing most dissimilar and dark red showing most similar. Note the matrix is symmetrical. (D) shows the transition ratio between each of the CAPs, the y-axis giving the starting CAP and the x-axis giving the ending CAP. Note the diagonal would be transition of a CAP to itself, and so in this case is zero in each instance.

All CAPs showed relatively stable coverage across the group, with medians of each ranging from $12 - 20\%$. The lowest coverage across the group was CAP 2, which showed relatively consistent coverage across the group, other than a single outlier at 25%. CAP 3 showed the highest variability across participants,

with the minimum participant showing coverage of 9% and the maximum at 30%.

The similarity matrix in Figure 15C highlights the dissimilarity between CAP pairs and was used to pair the CAPs using the *Munkres* algorithm (see Section 5.2.6). The pair with the most minor dissimilarity were CAPs 3 and 4. The least common transitions between CAP also happened to be between the pairs (Figure 15D). The most common transition across the group was from 3 to 5, which interestingly occurred 45% of the time that CAP 3 was transitioned out of. Other common transitions were from CAP 2 to 5, 6 to 3 and 4 to 6.

5.3.5 Functional Significance of CAPs

Observed CAPs were applied to *Neurosynth* in order to associate their patterns with the broader field (Yarkoni et al., 2011), shown in Figure 16.

CAP 1 was associated with the DMN and various prefrontal regions as well as more abstract cognitive labels such as theory of mind, social, and autobiographical thoughts. Its pair CAP 2 was associated with visual processing, as well as visual regions such as V1 and the occipital cortex and spatial processing, and other regions such as those in the parietal cortex and fusiform gyrus.

CAP 3 showed labelling with somatosensory and motor tasks, the movement, motor and premotor cortices, whereas CAP 4 showed a clear correlation with more cognitive states, such as memory and working memory, and tasks. The associated regions were the angular gyrus, frontal, and prefrontal regions.

CAPs 5 showed an association with visual regions and some association with temporal regions. CAP 6 labelled sensory regions primarily associated with somatosensory and pain.

For further clarification on function, the cortex of each CAP was correlated with the first three functional connectivity gradients derived by (Margulies et al., 2016), which are visualised in Figure 5.

Figure 17 shows the correlation between each CAP and each gradient, with each gradient correlation represented as an axis, creating a 3D correlation space, or "gradient space". Each pair is coloured similarly and is connected with a line.

Each of the CAP pairs was strongly correlated with the positive or negative of one of the first three gradients derived in Margulies et al. (2016).

5.3.6 EEG Microstate n-gram Parameters

The transition probabilities between each of the data driven microstates across the group are shown in Figure 18.

Abnormally low transitions were observed into microstate D across the group. The most common transitions were generally from B to A and A to B, as well as C to B, B to C and F to C.

Microstate n-gram parameters mean duration and frequency were derived from lengths 2 to 5 in event-mode. A visualisation of duration distribution per n-gram and frequency is shown in Figure 19 for 2- and 3-grams. Visualisation

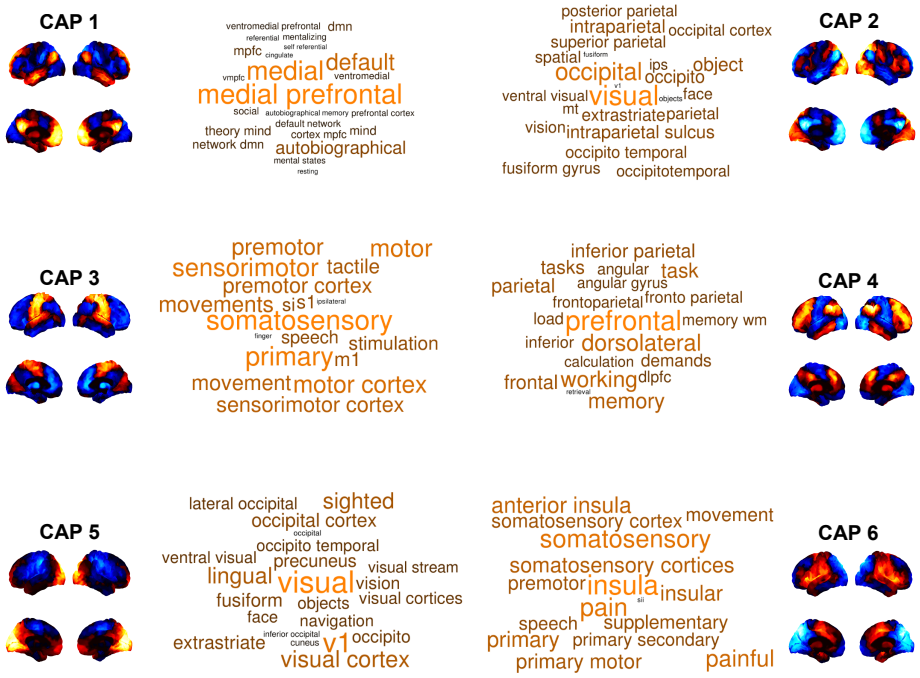

Figure 16: Word clouds indicating the most correlated topics to each CAP generated using Neurosynth image decoder. Each CAP has their visualisation shown next to its word cloud. Larger words and more orange words indicate most strongly correlated topics for each CAP, where black text indicates less correlated topics. Abbreviations are as follows: medial Prefrontal Cortex (mPFC), ventromedial Prefrontal Cortex (vmPFC), Inter-Parietal Sulcus (IPS), Middle Temporal (MT), Dorsolateral Prefrontal Cortex (dlPFC), Working Memory (WM).

of 4- and 5-grams is not given here due to the number of n-grams to view (320 for 4-grams and 1280 for 5-grams).

On each subplot, the left axis shows mean duration, and the right shows frequency. The violin plots show the overall distribution of durations of each n-gram, with the white dot showing the mean across participants and the line showing the median. The grey bars show the frequency of each n-gram. Violin plots are coloured based on the first microstate in the sequence. In general, we can see how the shorter coverage and frequency of C carries up to occurrences of n-grams that contain microstate C. Some 3-grams that show abnormally high frequencies appear to be repeating states, such as DAD and its cycle ADA, as well as DFD and BFB.

The microstate time series was then aligned with the fMRI time series, and n-gram parameters were calculated during each CAP label. Figure 20 shows the 2-

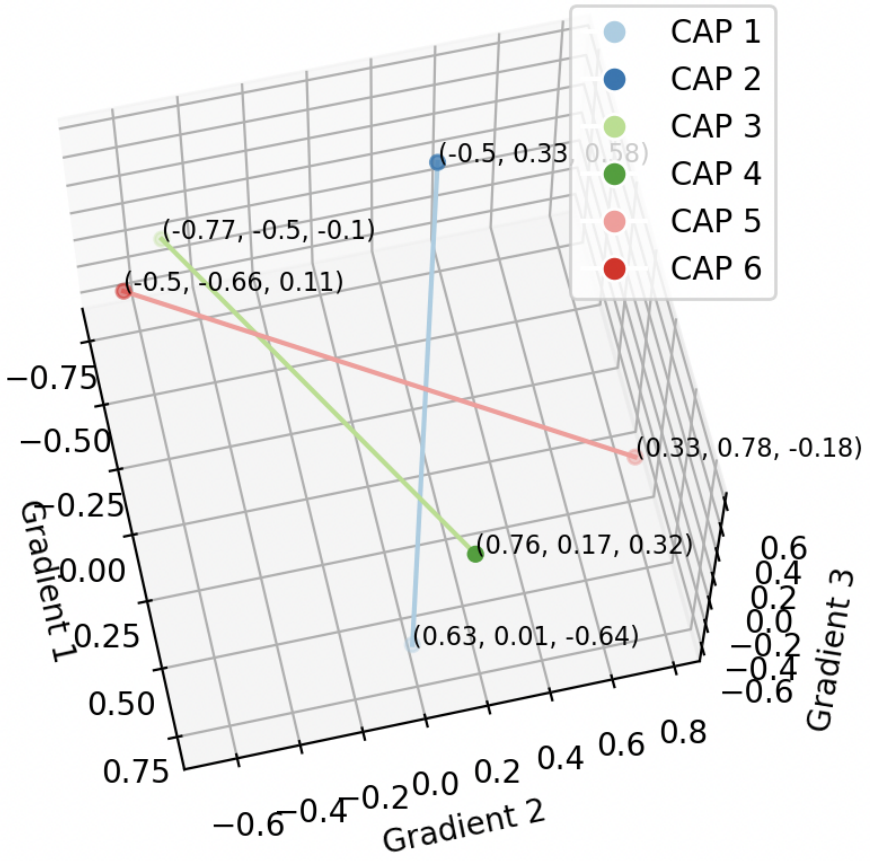

Figure 17: CAPs projected into gradient space. Each CAP was correlated with each of the gradients derived by Margulies et al. (2016). The coordinate on the given axis denotes the CAPs correlation with the given gradient. The CAP pairs are coloured similarly and connected with a straight line.

gram distributions of mean durations during different CAPs across participants. L1 distances were calculated between these distributions for each participant individually.

5.3.7 L1 Distance Null Distributions Compared to Observed CAP Differences

Since many calculations were made for each combination of parameters and n's for each participant, visualisation of all L_1 distances calculated and p-values generated can be found here.

After calculating L_1 distances, the process of creating sparse graphs was

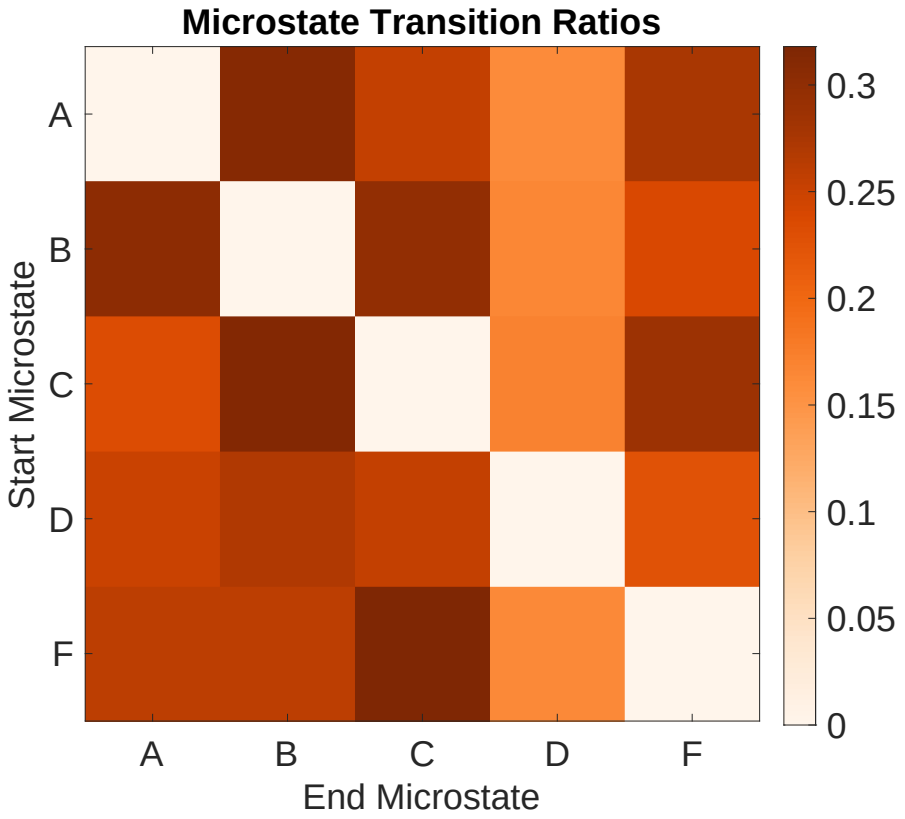

Figure 18: Matrix showing transition ratios between microstates across the group. Y-axis shows the starting microstate, with x- axis shows the ending microstate. Note that the matrix is not symmetrical. Colour bar indicates ratio.

carried out for both observed and shuffled values. Figure 21A shows the L_1 distances between CAPs for $n = 1$ coverage in participant $RS007$. The all-to-all graph between CAPs, with the L_1 distance values as edge weights, is given if Figure 21B. After thresholding of edge weights, those over the threshold are shown in Figure 21C.

The total number of degrees for the graph is then summed across all nodes. Note that the number of degrees in the sparse graph counts a single edge twice since that edge is counted in the degrees of both nodes. Figure 21D shows the total number of observed degrees at the red line - 10 degrees (5 connections on 21C, each counted twice). The histogram indicates the number of shuffles for each number of degrees. In the example shown, all shuffles showed no connections between CAP nodes over the threshold.

For all participant coverages, no edge weights were over the threshold during

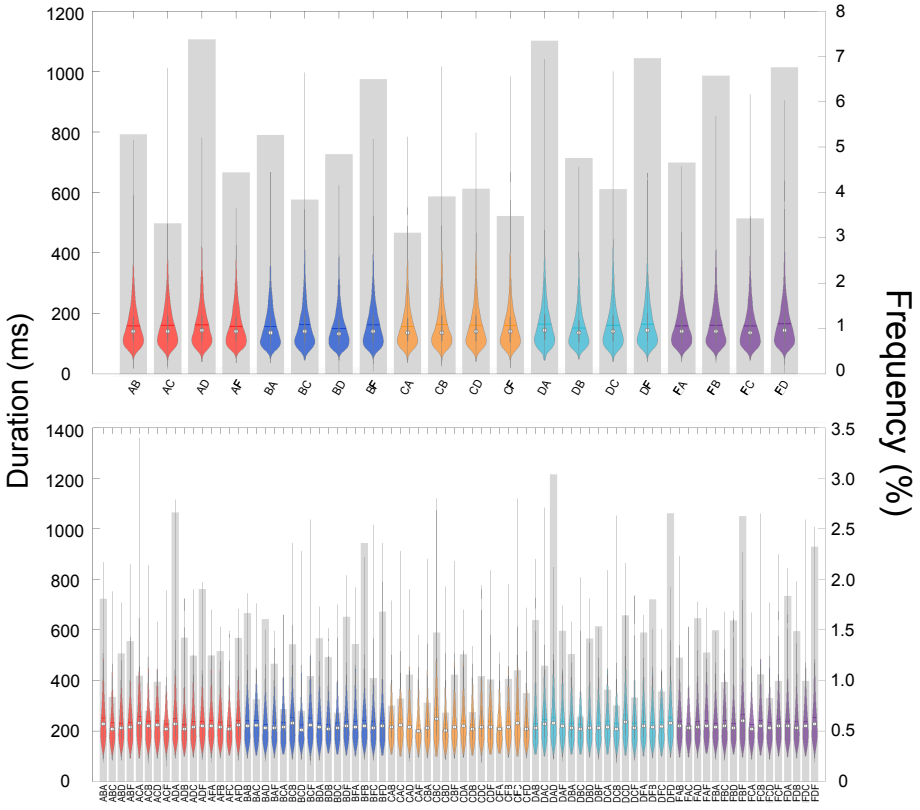

Figure 19: Visualisation of microstate 2- and 3-grams across the group. Coloured violin plots represent duration distributions of each n-gram across the group. Colours indicate the starting microstate of each n-gram. White dot on each violin indicates median, with line indicating mean. The duration distributions correspond to the left. Right axis gives frequency of each n-gram across participants, with grey bars corresponding to this axis.

shuffles, and all observed data showed at least one connection retained for each participant (equivalent plots for all observed and shuffled values, for frequency, coverage and mean duration, across all n's can be found here, in the Study 1 results section).

The distributions identified for shuffles were used as null distributions versus the observed value in each case to calculate a p-value. Matrices that show the p-values for each parameter at each n-gram length are shown in Figure 22.

Mean Duration exhibited three participants, which showed a meaningful difference to the shuffled CAP total number of connections: participants 5 and 18 for $n = 4$, and participant 8 for $n = 3$. In frequency, participant 17 showed abnormally different CAP labels compared to the group, with differences from the null distribution observed for $n = 1$, 2, 3 and 4. Participant 13 differed

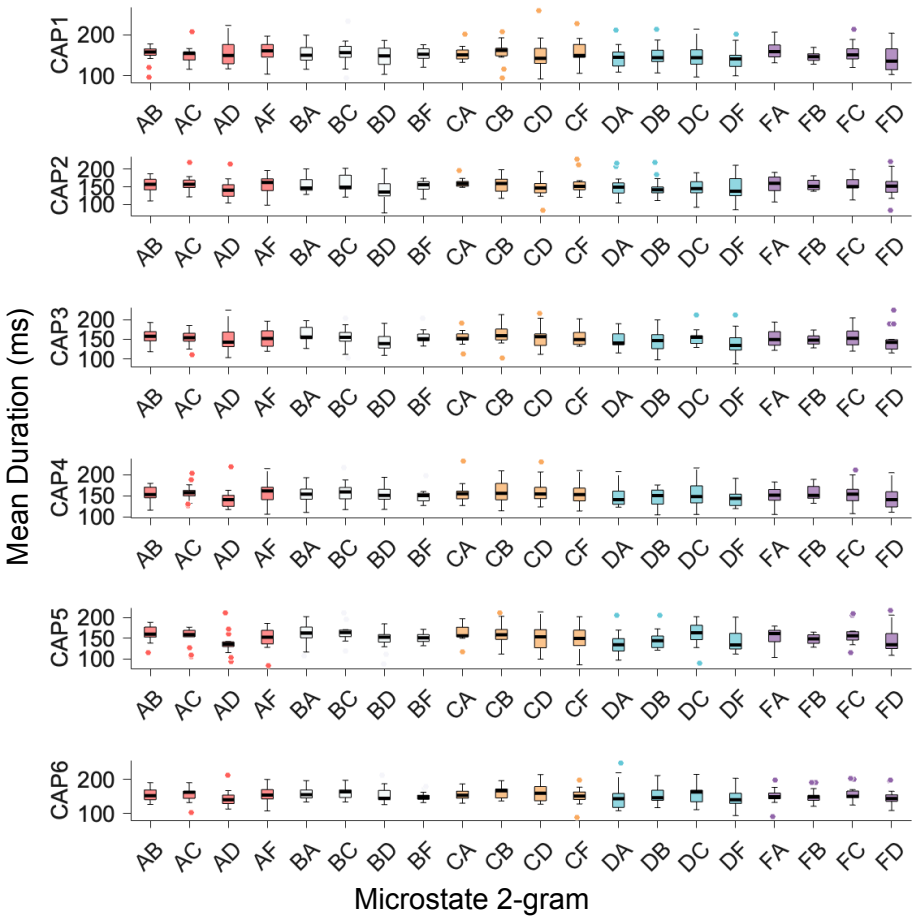

Figure 20: Mean duration of microstate 2-grams during each of the six observed CAPs across participants. Box plots denote each 2-gram mean duration distribution across participants. Subplots top to bottom show mean duration of each 2-gram for CAPs 1 to 6. Colour of each box denotes the starting microstate in the n-gram.

from the null distribution in $n = 3$. Since there were no coverage connections in shuffled CAP labels as mentioned previously, and all observed participants showed at least one connection, all p-values calculated for coverage were the minimum value of 0.001 ($P = 1000$).

In each case where any p-value was less than 0.05 for a single participant, the calculated p-values across all n's for the given participant and parameter were subject to FDR multiple comparison correction. Figure 23 shows an example correction across multiple p-values for participant 17 frequency.

All observed p-values which were less than 0.05 survived the correction. It

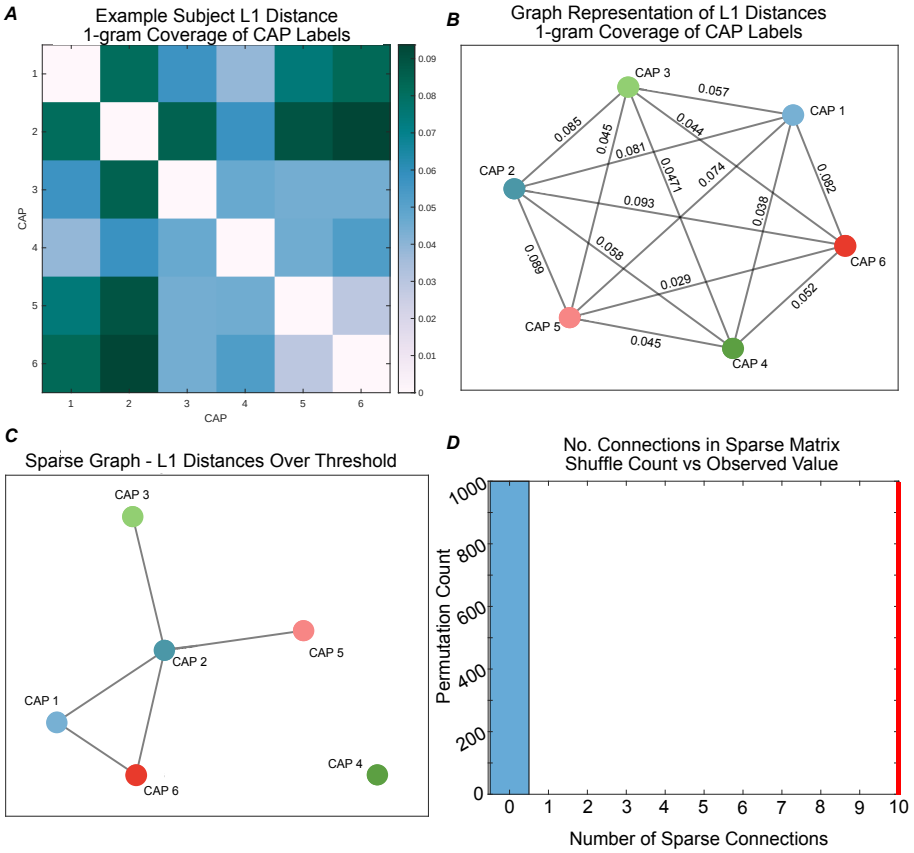

Figure 21: Plot showing example observed L_1 distances between observed EEG n-gram coverage distributions between CAPs at $n = 1$ for individual participant 7. (A) Pairwise L_1 distances between CAP labels. (B) All-to-all graph representation of L_1 distances between CAP labels. Nodes are CAP labels and edge weights are observed L_1 distances between the microstate n-gram parameter distributions of connected CAP labels. (C) Sparse graph of connections, removing all connections from (B) that are less than the threshold defined in methods. (D) Comparison of total number of degrees of observed graph (in red) and the same value in shuffled CAP labels (blue histogram). Y-axis shows the number of shuffles out of 1000 that were observed at the given number of connections. X-axis shows the number of connections. Equivalent plots for all other participants, parameters and n's can be found here.

was the case that all p-values were under the threshold for all those in the bottom section of Figure 22. The coverage p-values were not subject to correction since there was only one observation per participant.

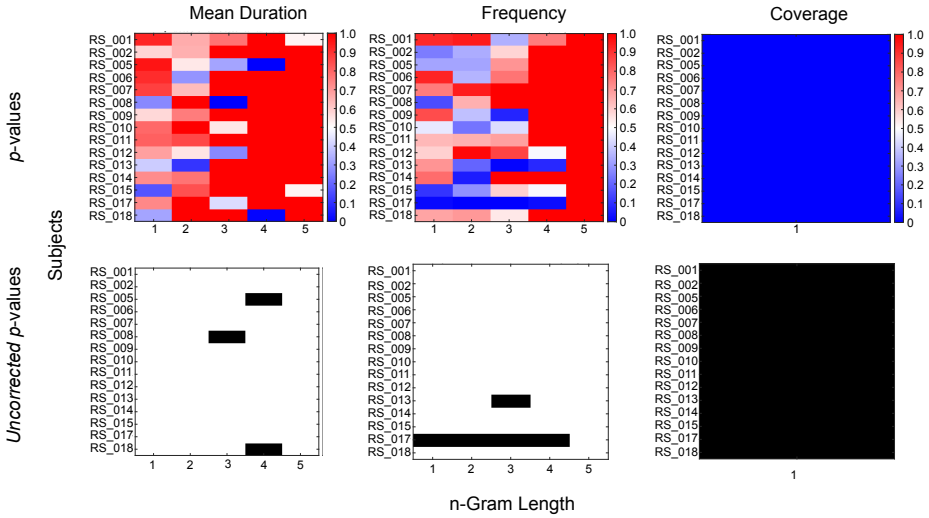

Figure 22: Output p-values for each participant across all lengths, for each parameter. Left column shows mean duration, middle column shows frequency and right column shows coverage. Top row shows p-values where most red is 1 and most blue is 0. Bottom row marks black where p-value is less than 0.05. In each matrix, the y-axis shows the participant and the x-axis shows the n-gram length investigated. Note that for coverage there is only 1 n-gram length, since coverage cannot be calculated for $n > 1$.

5.3.8 Observed Differences in Microstate n-gram Parameters between CAP Labels

The parameters at each n-gram length, which were significantly different, were subject to further investigation. Since all CAP node degree counts for coverage were significantly different from shuffles for all participants, the connections retained between CAPs were investigated across participants. The retained connections in the sparse matrix of each participant were first isolated. Figure 24 shows the number of degrees of each node for each participant.

Figure 24A shows the distribution of the number of degrees for each CAP across participants. All but CAP 4 show a median degree of 1. CAP 6 shows all but 4 participants with a degree of 1. Hence, they are identified as outliers. CAP 2 is the only CAP which observed a participant with 4 degrees.

Figure 24B then shows the most to least common connections between CAPs across participants, with the y-axis indicating the number of participants out of 15 where the observed distance was over the threshold.

The most common difference was between CAPs 2 and 5, with 5 of the 15 participants observed with this difference retained over the threshold. Reviewing differences in microstate coverage distribution between CAPs 2 and 5, the most common factor for differences between these CAPs is not consistent across participants(Figure 24C).

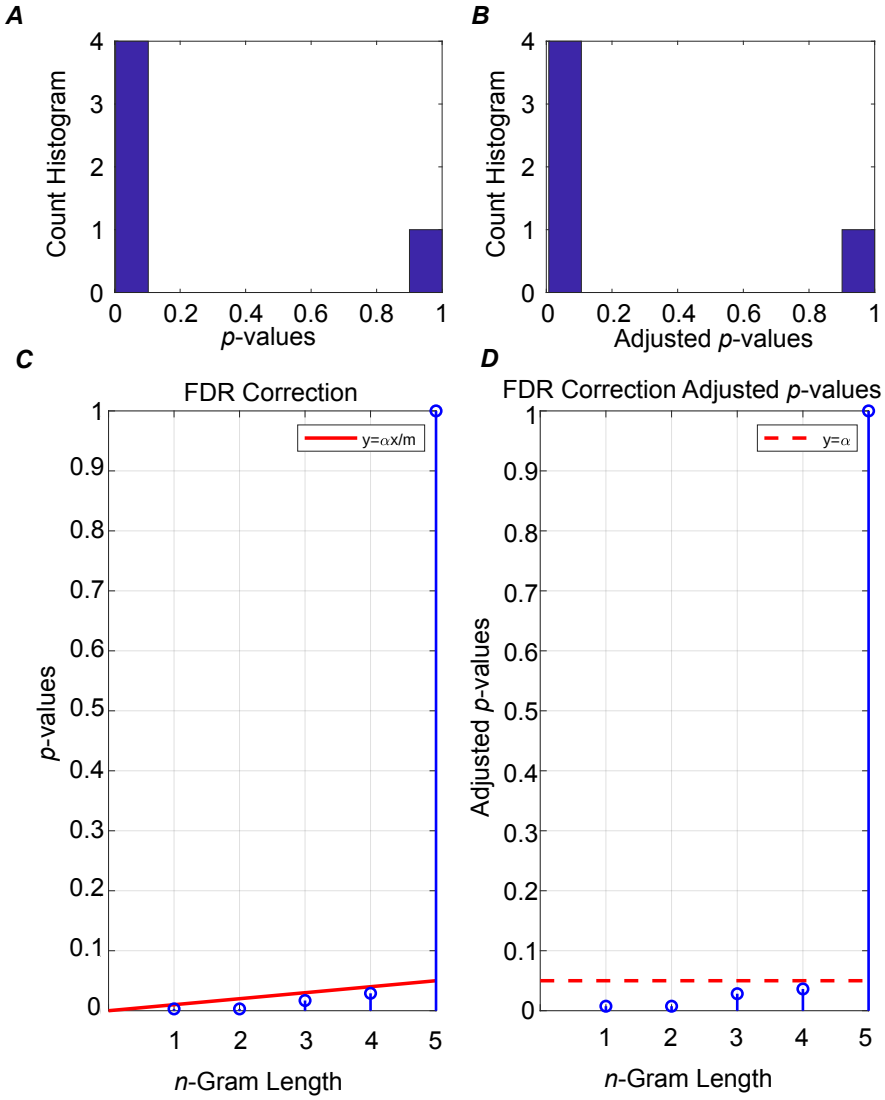

Figure 23: Example False Discovery Rate across p-values for participant 7 frequency. (A) Shows observed p-values. (B) shows p-values after Benjamini-Hochberg adjustment of threshold p-value. (C) and (D) show the value of each observed p in relation to the uncorrected and corrected thresholds respectfully.

5.4 Discussion

5.4.1 Microstates and Parameters

Included in the five data-driven microstates are the canonical four microstates presented in past studies (Khanna et al., 2014; Michel & Koenig, 2018; Milz

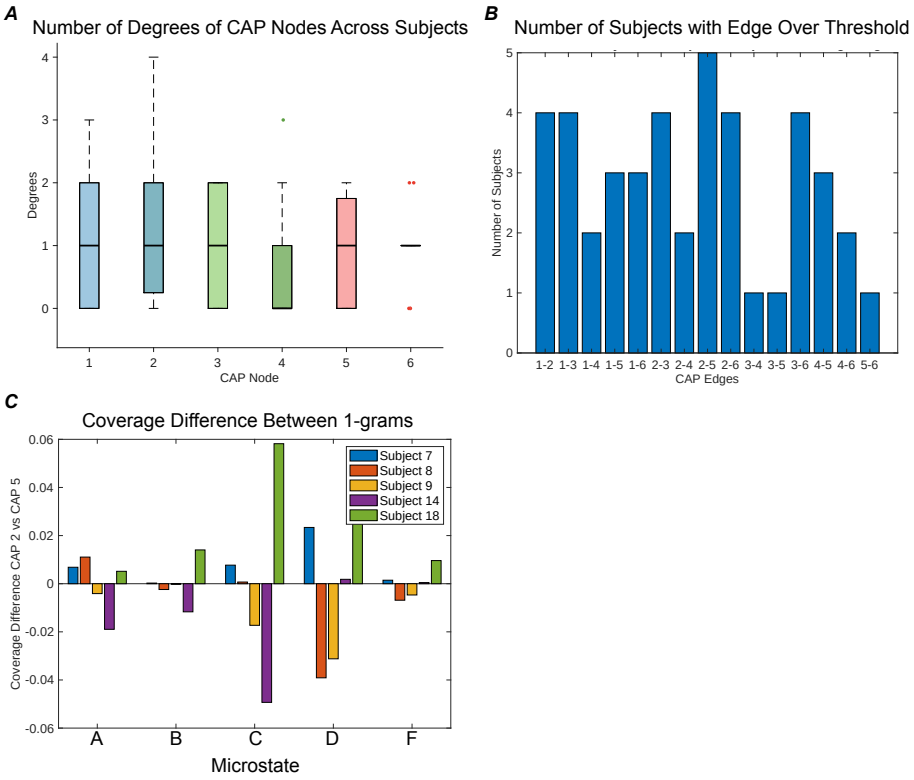

Figure 24: Observed Differences between microstate Coverage During Different CAPs across participants. (A) Degree of each CAP node across participants. Degree is the number of connections a node has to other nodes. Box plots show the distribution of each CAPs degree across participants. Outliers are defined as the median plus 1.5 times the standard deviation. CAPs are coloured by their pairs. (B) The number of participants out the 15 observed that retained an edge between CAPs after thresholding. Each bar shows an edge. Since graphs were not directed, 1-2 is equivalent to 2-1. (C) The difference between the coverage of each microstate between CAPs 2 and 5 across the five participants that observed a difference higher than the threshold in the graph analysis. Each colour indicates the participant in question. Positive numbers indicate coverage is higher in CAP 2, negative in CAP 5.

et al., 2017) as well as meta-microstate F (Koenig et al., 2023).

In general, the range of mean durations of the generated microstates agrees with previous studies, which state a general duration of $60 - 120$ms (Khanna et al., 2014; Michel & Koenig, 2018). Frequency is proposed here as a novel parameter independent from duration and can be considered a normalised occurrence (see Section 1.1.3).

Results reflecting a general lack of preference for microstate C is unusual

when related to the rest of the literature, as in general C has shown a preference for eyes open studies (Seitzman et al., 2017). However, the parameters observed for microstate B were in line with these findings, which showed a higher preference of B for eyes-open versus closed.

Microstate D consistently demonstrated the highest measures on all parameters on group averages. Microstate D is speculated as being associated with resting state by multiple studies (Antonova et al., 2022; Milz et al., 2016), which this study agrees with despite lack of comparison to other task modalities in the same group.

The meta-microstate F requires further investigation, as it rarely appears in analysis. Custo et al. (2017) conducted source localisation, which associated the microstate with regions of the anterior DMN. Bréchet et al. (2019) found associations with prefrontal regions of the DMN. The DMN has been associated with the resting state (Damoiseaux et al., 2006). Microstate F here showed the highest GEV across the group and high-frequency values, which may reflect this. Further investigation would, of course, be required.

5.4.2 Microstates fMRI General Linear Model

The standard approach for comparing microstates and fMRI was applied here to compare to the developed methodologies (Figure 13). In general, past simultaneous EEG-fMRI studies have not agreed upon the fMRI activation patterns associated with each microstate. A review of these associations can be found in Section 1.1.5.

Microstate A has, in general, been associated with a deactivation of the phonological network (Antonova et al., 2022; Britz et al., 2010), as well as visual processing (Milz et al., 2016). Here, microstate A showed activation in the right cingulate gyrus and the right supplementary motor area. Correlation is also found on the left and right superior temporal gyrus and the left and right frontal pole. Correlations with the temporal gyrus suggest some overlap with reports of the phonological network.

Microstate B has generally been associated with activity in visual regions (Antonova et al., 2022; Britz et al., 2010). Microstate B here, however, shows a very similar activation pattern to microstate A, which is unusual.

Britz et al. (2010) reported microstate C as being associated with the posterior of the Anterior Cingulate Cortex (ACC) as well as bilateral Inferior Frontal Gyri (IFG), the right Anterior Insula (AI) and the left claustrum. Here, microstate C showed strong correlations across the group with a small region on the border of the right inferior and middle frontal gyrus and the frontal pole.

Microstate D shows a group correlation with the right amygdala and the left occipital pole, which has not been reported previously, as microstate D is most commonly associated with the DMN (Antonova et al., 2022; Britz et al., 2010; Michel & Koenig, 2018).

Finally, microstate F showed group correlations with the right temporal occipital fusiform cortex and the left occipital pole. To my knowledge, microstate F has yet to be compared in this way to fMRI. Past reports of microstate F

have associated its activity to DMN (Custo et al., 2017), and the Prefrontal Cortex (PFC). The results here do not align with these suggestions.

Whilst this approach does assist in identifying potential generators of each microstate, as outlined in Chapter 2, this approach is unsuitable for microstate syntax analysis. The down-sampling of the EEG time series required to implement the approach destroys the syntactic structure of the microstate sequence.

Additionally, an equivalent set of custom regressors that used n-grams over microstates would not be possible, as the correlation is done on a time-point-by-time-point basis. A sliding window approach to correlation could be implemented, but an arbitrary window size would have to be used to implement this approach. As such, whilst the findings here contribute to previous findings, the methodology implemented here does not utilise the microstate syntax.

5.4.3 CAPs and Parameters

CAPs are an underutilised means of investigating dFC. Some studies have used CAP analysis to identify common patterns of activity that occur during the peaks of activity in a specific seed region by clustering the TRs where the peaks occur (Amico et al., 2014; Li et al., 2021; Liu & Duyn, 2013). Both seed-based and whole brain analysis are possible; however, (Liu et al., 2018), with the use of CAPs motivated by the possibility that spontaneous BOLD signal results from brief and temporally isolated (that is, on the fMRI times scale) neural activity.

Whilst past studies have investigated how CAPs differ between tasks (Li et al., 2021) and groups (Amico et al., 2014), I am not aware of any existing studies which attempt to investigate CAPs as a set of overall states that summarise the resting state time series as a sequence of these states.

The phenomenon of observed CAP pairs identified here was described in the past by Li et al. (2021), who reported anti-correlations in brain activity during task performance and in resting state data, pointing to potential fluctuations between anti-correlated states.

These pairs highlight that they are representative global activity patterns when investigating CAPs. The activation of a network is also observed with the attenuation of all other regions. A CAPs pair shows its opposite, with the attenuation of its activated regions and activation of its attenuated regions. These pairs may be explained by the observations made in gradient analysis studies. Each of the six CAPs observed in this study were found to be well anti-correlated pairs (Figure 17), and the three pairs correlated strongly with different gradients (Figure 17).

The first gradient (sometimes referred to as the principal gradient or axis) has been identified as a spectrum between sensorimotor and transmodal regions (Huntenburg et al., 2018; Margulies et al., 2016) and supports an existing theory that proposes that human brain activity is organised along a "unimodal-transmodal" axis (Mesulam, 1998). CAPs 3 and 4 showed a strong correlation with either end of this axis, with the other pairs also showing weaker correlations.

Furthermore, gradient two has been referred to as a sensory axis, distinguishing visual regions from other sensory regions (Margulies et al., 2016), and also corroborates the pre-existing theory proposed by Mesulam (1998) which states that sensory activity in the human brain is organised along this axis. CAPs 5 and 6 appeared to be strongest at the poles of this axis, with CAP 5 also showing some preference for the opposing end of gradient 1.

Whilst the third gradient has been investigated to a lesser degree, it has been speculated as the "task axis", with solid correlations being associated with goal-orientated cognition and negative correlations being associated with resting state and mind wandering (Margulies et al., 2016; Smallwood et al., 2021). CAPs 1 and 2 showed the most substantial relationship with either pole of this axis. The meta-analysis conducted using Neurosynth agrees with the task-axis suggestion (Figure 16).

The observations made here indicate that the use of clustering analysis in the time domain on fMRI data may derive the poles of each of these gradients organically, and both methodologies may identify similar functional patterns. It should be pointed out, however, that since each was derived on a different dataset (the gradients used here were derived by Margulies et al. (2016)), further investigation into these observations is required.

Chen et al. (2015) created a set of parameters to describe CAPs and proposed these be used to understand better the temporal dynamics of said states, as well as understand which CAP was the most dominant. Some parameters defined there are equivalent to those proposed here. The coverage of each CAP demonstrates overall dominance on the time series, and dwell time shows, on average, how long each CAP was occurring for, similar to those suggested by Chen et al. (2015).

The CAP pairs 3 and 4, and 5 and 6 showed high median coverage across the group versus 1 and 2, whereas the average dwell time across CAPs was relatively consistent, with few differences observed between CAPs.

The transition ratios between CAPs show a low-resolution representation of CAP dynamics, with the most common transitions from 3 to 5, 2 to 5, and 6 to 3. All these common transitions are between generally more unimodal sensory CAPs (Liu et al., 2018; Margulies et al., 2016). This, combined with relatively low coverage of CAP 2, suggests a less stable occurrence of sensory CAPs during this observed resting state data. However, more investigation would be required to confirm this.

5.4.4 Statistically Significant Differences between n-Gram Parameters Occurring during Different CAPs

The aim of identifying differences in microstate n-gram parameters that are occurring during simultaneously occurring fMRI CAPs is to show that microstate n-grams and their syntax differ between brain states. Whilst this was demonstrated in many participants in coverage, there were fewer observations in mean duration and frequency, which identified this, with observed differences in the

latter parameters being few. Additionally, there was a lack of consistency between the participant microstate coverages during different CAPs.

A comparison between underlying fMRI states has yet to be attempted between microstate or n-gram parameters. This is in part likely due to existing comparisons between microstates and fMRI not considering the temporal domain sufficiently (See Chapter 2), but may also be due to the difficulty of comparing these two domains that are active in different temporal resolutions.

It was shown that differences in coverage between CAP labels were significantly different in the observed data than in shuffles in all participants, indicating that the observed microstate coverages during each of the CAPs may be related to the CAP labels themselves, and not due to arbitrary labels. When investigating the driving label differences, the distance between CAPs 2 and 5 was most commonly observed in participants, with 5 of the 15 participants showing a difference greater than the threshold.

Despite this difference being the most common, when reviewing the difference across participants, no consistent difference in distributions was observed. For example, in 2 of the 5 participants, there was a higher coverage in CAP 5, whereas in another 2, there was higher coverage in CAP 2. Whilst the present study does suggest some connection between microstate n-gram parameters and underlying fMRI CAP states, this lack of consistency makes it difficult to tie down any solid comparisons, suggesting further study is needed using this methodology.

Whilst the methodology proposed here demonstrates a means of comparing the two modes of recording without reducing the dimensionality of either of them, some shortcomings are inherent in the approach. When comparing n-gram parameters between CAPs, it is assumed that each of the n-grams is independent of one another when they are most definitely not. Each n-gram in an observed sequence of microstates overlaps with one another, meaning that the calculation of the parameters of one n-gram includes the constituent n-grams that make up that n-gram at smaller ns, as well as neighbouring n-grams on the time series.

It should also be highlighted that the method implemented here; using $L1$ distances between CAPs across all n-grams, results in loss of information about differences in individual n-grams. Distances between individual n-grams may cancel out between two CAPs for example, leading to a false negatives. Since identifying difference in the parameters of specific n-grams between CAPs is the objective of the methodology, this shortcoming must be considered in earnest. This approach was chosen however due to the number of n-grams that must be compared between CAPs. Comparing every n-gram between CAPs would make identification of statistically significant differences impossible after multiple comparison correction. The extra steps taken to identify specific differences were implemented to attempt working around the issue of loss of information in comparing whole n-gram parameter distributions between CAPs. Following identification of a general difference between CAPs, investigation of the individual differences between n-grams was conducted, as shown in the example of Figure 24. It is accepted that this approach may miss out on false negatives, as

previously mentioned. Further refinement of this approach would be required to address such issues.

At longer lengths of n, it is also the case that the parameter distributions are stretched thin. More data would be required to increase n. Many n-grams at length five did not occur in multiple participants, meaning that when the observed participant time series was split into six classes for each of the CAPs, the possibility for insufficient data also increased.

Finally, the lack of consistency across participants when comparing microstate n-gram parameters between fMRI states suggests that the relationship may be more complex than the observed parameters of n-grams can identify. Comparison between categories of fMRI states may limit the differences that can be identified. Therefore, it may be more beneficial to move away from fMRI CAPs as a means of analysis and use a continuous representation of the fMRI signal to compare against the microstate parameters to allow more specific associations to fMRI networks. The following study develops a methodology which attempts to do this.

6 Study 2 - EEG Microstate n-gram Parameters in the fMRI Gradient Space

Previous studies have attempted to understand Electroencephalography (EEG) microstate function through association with simultaneously recorded Functional Magnetic Resonance Imaging (fMRI) Blood-Oxygenation Level Dependency (BOLD) signal. Such studies have most often used fMRI Intrinsic Connectivity Network (ICN)s as common spatial patterns of fMRI activity and have attempted to correlate such patterns with the dominance of each EEG microstate. More recently, an investigation utilising global patterns of fMRI activity has taken place, utilising dynamic functional connectivity states to understand how microstate activity may be associated with global fMRI activity rather than with localised networks. The previous chapter developed on this by associating EEG microstate n-grams with global fMRI activation patterns rather than individual microstates. Here, I develop the approach by considering the fMRI time series as a continuous signal rather than a set of dynamic states using gradient correlations. This continuous signal allows for an overview of activity across the time series rather than temporally localised associations.

6.1 Introduction

Past studies have attempted to investigate the relationship between EEG microstate n-grams and mental states using cognitive manipulation studies (Lehmann et al., 2005; Schlegel et al., 2012). Others have also identified that information is retained in the syntax of observed EEG microstate sequences (von Wegner et al., 2017). These short sequences of microstates were observed as non-Markovian, as the microstate observed at a given time depended on more than the previous state (first order) or the previous two states (second order).

Despite these observations regarding the functional significance of EEG microstate sequences, whilst there have been attempts to localise the source of EEG generators in the brain (e.g., Pascual-Marqui (1999), to my knowledge, there has been no attempt to investigate the relationship between EEG microstate syntax and fMRI. Studies which have investigated the relationships between microstates and fMRI have only done so using static states and have not investigated the dynamics of fMRI states or microstate sequencing (Abreu et al., 2021; Britz et al., 2010; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012). The difference between EEG and fMRI sample frequencies are seldom considered in these investigations (see Chapter 2 for a review). A method which uses sequences of microstates that occur for a period of time equal to that of a single fMRI Repetition Time (TR), or the dwell time of an fMRI state, has yet to be utilised.

Preliminary studies suggest there are differences between microstate parameters during different fMRI dynamic Functional Connectivity (dFC) states (Chapter 5), but the differences identified were difficult to pin down. Although differences in microstate n-gram parameter distributions could be observed between fMRI Co-Activation Pattern (CAP) states, those differences were either difficult to generalise across participants or differences identified were isolated to individual participants.

It may be the case that these difficulties were due to the finite number of fMRI states that were derived. Categorising the whole fMRI time series into a set of six states may be an oversimplification of the BOLD signal and may have made identifying associations to microstate syntax difficult.

Recently, fMRI gradients have been utilised to relate the activity profiles of brain regions across parcellation boundaries(Margulies et al., 2016). They are a set of functional maps derived by applying diffusion embedding (Coifman et al., 2005) to the overall connectivity matrix of recorded fMRI. A set of gradients has been derived on the Human Connectome Project (HCP) dataset (Margulies et al., 2016), which have since been identified as functionally significant axes (see Figure 5 and Section 2.1.2 for an in-depth description).

As such, I suggest a novel methodology that utilises fMRI gradients in conjunction with microstate sequence analyses. I describe an fMRI gradient space, where the observed fMRI time series is correlated with each of the gradients, taking said gradients as axes and creating a multi-dimensional space. This dimensionality reduction reduces the observed continuous BOLD signal from thousands of voxels down to a few relatively orthogonal representations of the

original signal.

The signals are correlated with each of the first three gradient axes proposed by Margulies et al. (2016). Each of these first three axes has been proposed as a "unimodal-transmodal" axis, "visual-somatosensory/motor" axis, and "rest/-task" axis, respectfully (Margulies et al., 2016; Mesulam, 1998).

The proposed method attempts to predict the dimensionally reduced BOLD signal using co-occurring EEG microstate n-gram parameters as features in an ensemble regression model. Preliminary results using this approach are reported here using peak-mode n-grams, with the goal of the method being to identify how EEG microstate n-gram parameters may change with the simultaneously changing coordinate of the fMRI gradient space.

6.2 Methodology

All EEG and fMRI pre-processing that was outlined in Study 1 (Chapter 5) was also used for analysis here. The microstate classes generated in that study were also used here (see Chapter 4 for an overview of the dataset used). First, The difference is in the derivation of microstate durations and temporal modes. Here, the microstate sequences were derived using the peak-mode, meaning interpolation was used to define durations instead of back-fitting (see Chapter 3 for a detailed overview of these differences). The other difference in the pipeline before analysis was using fMRI gradient coordinates. These differences are described in the following two sections of this chapter.

6.2.1 Derivation of Microstate Parameters for Peak-Mode n-Grams

As discussed in Section 3.2.3, there are different temporal modes for analysing microstate sequences. Here, peak-mode was used. Briefly, peak-mode sequences consider the boundaries between microstate readings as the centre points on the time series between consecutive Global Field Power (GFP) peaks. In the peak mode, microstate sequences can have multiple repetitions of a single class next to one another if consecutive GFP peaks were labelled with the same class. This cannot take place using the more standard event-mode.

Microstate n-gram parameters were calculated for peak-mode n-grams. These are equivalent to the parameters calculated in the previous Chapter 5 where event-mode was used: mean duration, frequency and coverage for $n = 1$ to 5. Note, however, that comparison of parameters between modes should be done cautiously. Whilst comparison between peak-mode and event-mode durations and coverages can provide a direct comparison between the two methods of duration definition (back-fitting and interpolation, see Section 3.2.2 for further discussion), frequency of a given n-gram depends specifically on how many n-grams there are at the given n. More peak-mode n-grams can occur versus event-mode due to the possibility of repeating classes.

6.2.2 fMRI Gradient Coordinate Calculation

The cortex parcellations of the observed data of each participant were correlated with each of the first three gradient axes derived by Margulies et al. (2016). This approach aimed to create a coordinate space where all participants' fMRI data could be located. Since each of the three axes has also been associated with different cognitive functions, placing each participant's time series within a space categorised by these axes could potentially point to changes in participant functional association throughout the recording.

6.2.3 Alignment of Microstate n-Gram Parameters with fMRI Time-Series

Similar to the alignment process described in Chapter 5, the recorded EEG and fMRI time series were aligned with recording offset, as well as the Haemodynamic Response Function (HRF) time to peak of 6 seconds. Following this alignment, the microstate sequence was cut in windowed blocks that corresponded to fMRI TRs in time (2 seconds). If a microstate (a single peak mode microstate) occurred across the boundary between TRs, it was removed from the analysis.

The constituent n-grams of each TR window of EEG sub-sequence were then used to calculate parameters, using n-grams *within the TR* only. The mean duration and frequency of n-grams with $n = 1$ to 5 were calculated for each TR. It should be highlighted that these parameters cannot be compared directly to those derived in Chapter 5, since the whole time series of a participant was used to calculate parameters in the previous chapter.

The goal was to have a distribution of EEG microstate n-gram parameters, which could predict the observed fMRI gradient coordinates at each TR.

6.2.4 Application of Random Forest Regressors

A random forest serves as a meta-estimator (Breiman, 2001), constructing multiple classification decision trees on different subsets of the dataset. A decision tree is a flowchart-like structure used for making decisions in a tree-like model. The tree is constructed by recursively partitioning the data based on features that best predict a continuous target variable. Mean Squared Error (MSE) was used in this case as a measure of "impurity" to minimise.

The random forest employs averaging across the decision trees to enhance predictive accuracy and manage the risk of over-fitting. In this case, the features used were the distribution of microstate n-gram parameters at a given n, and the target continuous variable was the fMRI gradient coordinates of the given TR. This was repeated for each participant, at each n-gram length, for each parameter separately.

Here, *sklearn* was used to apply the random forest regressor (Abraham et al., 2014). For each model, a grid of hyper-parameters was searched for the best combination for the particular data, and the combination of parameters that best fit the data was used in each case. The hyper-parameters were: number of estimators as 50, 100, 150 or 200; maximum tree depth as none, 10 or 20;

minimum samples to split as 2, 5 or 10; and the minimum number of samples per leaf, at 1, 4 or 8. The data were split into train, validation and test segments. 60% of the TRs of the given participant were randomly selected as the training segment, 30% were used for validation, and 10% were selected for testing. TRs were selected randomly to avoid temporal dependencies.

6.2.5 Analysis of Random Forest Models

Upon fitting, each model which returned an R^2 value over 0.3 was subject to feature importance analysis. The overall importance of a feature is the average of the MSE decrease across all decision trees. The value is normalised across features so that all sum to 1 to quantify how much of the fit was attributed to each feature. In context, this quantifies the parameters of which microstate n-gram contributed the most to predicting the fMRI gradient coordinates.

The permutation feature importance was also quantified. This measures how much of a drop in model performance is observed when the given feature is permuted in the training data. In context, parameters of each n-gram are permuted within a given model, and a more significant drop-off in performance for the prediction of fMRI gradient space coordinates indicates the n-gram has high permutation feature importance.

Once the n-grams with the highest contribution were identified for the parameter, a partial dependency plot was used for these n-grams to identify how the given parameter changed with the gradient coordinate. For example, if the frequency of n-gram ABC contributed highly to the prediction of fMRI gradient 1 coordinates, the coordinates where that dependency was highest would be identified using partial dependency. Following this, the coordinates were projected into the original fMRI space to visualise the global fMRI activation pattern they represented.

6.3 Results

6.3.1 fMRI Gradient Coordinates

Each participant TR was correlated with each of the three gradients derived by Margulies et al. (2016). The gradients were conceptualised as axes, and correlation values were taken as coordinates in space to define an "fMRI gradient space".

Figure 25 shows the gradient coordinates across participants, with the global signal removed.

Figure 26A shows kernel density estimate of TRs across participants within the gradient space visualised as cross sections. The distributions appear to be relatively normal on each axis across the group. There is a general preference for the positive gradient 1 and 2 to negative 1 and 2 in the left panel rather than the more general distributions shown on the other cross sections.

Figure 26B shows how the density estimate differs between participants. In general, the normal distribution from the 0 point is retained, but slight

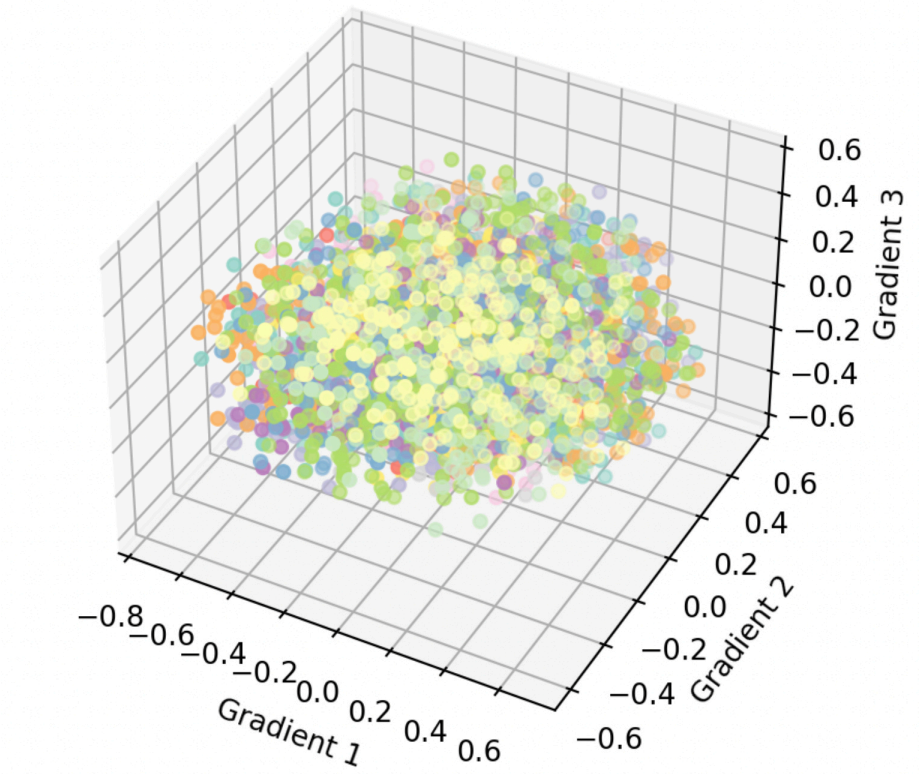

Figure 25: Participant fMRI gradient coordinates in 3D space. Each of the three axes denotes a correlation with one of the first three gradients derived by Margulies et al. (2016). Each colour indicates a different participant. Correlations are with global signal removed, since the gradients that were correlated with also have this signal removed. TRs are visualised as a scatter plot but each participant's time series of TRs can be conceptualised as a trajectory through the space.

differences in the amplitude of peaks and the width of the distribution can be seen. A few participants also appear to have slight preferences for extremes of gradients 1 and 2, which likely explains the offset in the distribution in the left panel of Figure 26A.

6.3.2 EEG Microstate Sequences in fMRI Gradient Space

Figure 27 shows a plot of the density distribution of each of the microstates within the fMRI gradient space.

Distributions of all microstates appear relatively normal, with specific density maxima appearing in slightly different locations. Microstates *A* *B* and *D*

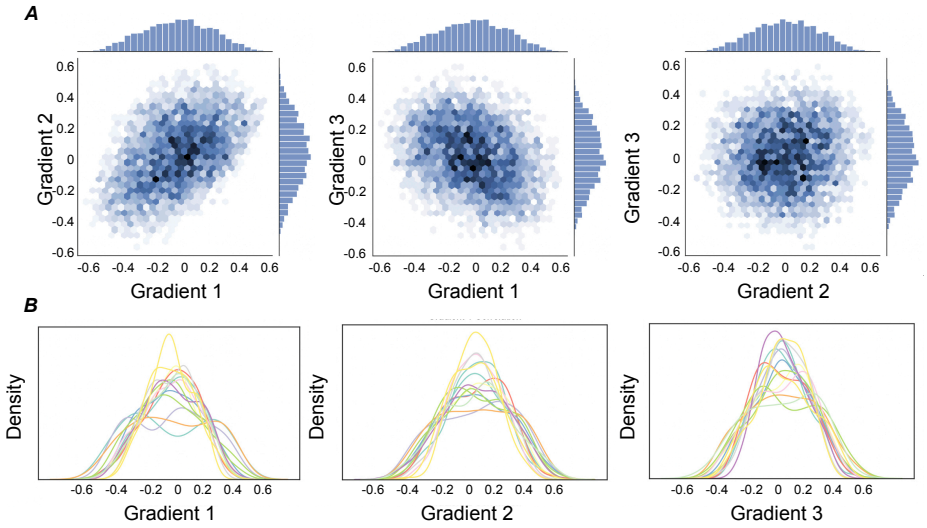

Figure 26: Kernel Density Estimate of fMRI coordinates in the three dimensional gradient space. (A) Cross sections of density estimates across all participants, 1 to 2, 1 to 3 and 2 to 3 left to right. (B) Kernel Density Estimate distributions of individual participants. Gradients 1 to 3 left to right. Individual colours indicate density estimates of different recorded participants.

show a pocket of density maxima at the null point, whereas microstate C is more evenly distributed. This may be due to the number of microstates in each case, which should be considered. Microstate F appears to have a more evenly distributed density, with no one particular point showing a pocket of high density.

6.3.3 Peak-Mode Microstate n-Grams

Figure 28 shows peak-mode microstate parameters across participants. The mean duration (Figure 28A) of peak-mode microstates is around $50 - 60$ms consistently across participants and microstates. There are differences in the distribution of microstates due to the difference in derivation methods when comparing these to the mean duration of event-mode microstates shown in the previous Chapter (Figure 12). Peak-mode microstates appear to be shorter, likely due to their derivation method. It is also the case that the underrepresented event-mode microstate C is not found in peak-mode.

Each microstate's frequency (Figure 28B) is also relatively consistent, with microstate D showing a group average that is lower than the other microstates, which generally showed frequencies of around 20%. This is in contrast to the event-mode microstates, where, in general, microstate C demonstrated lower parameter values. A similar distribution to frequency is found in the coverage distribution (Figure 28C), likely due to the relatively uniform distribution of

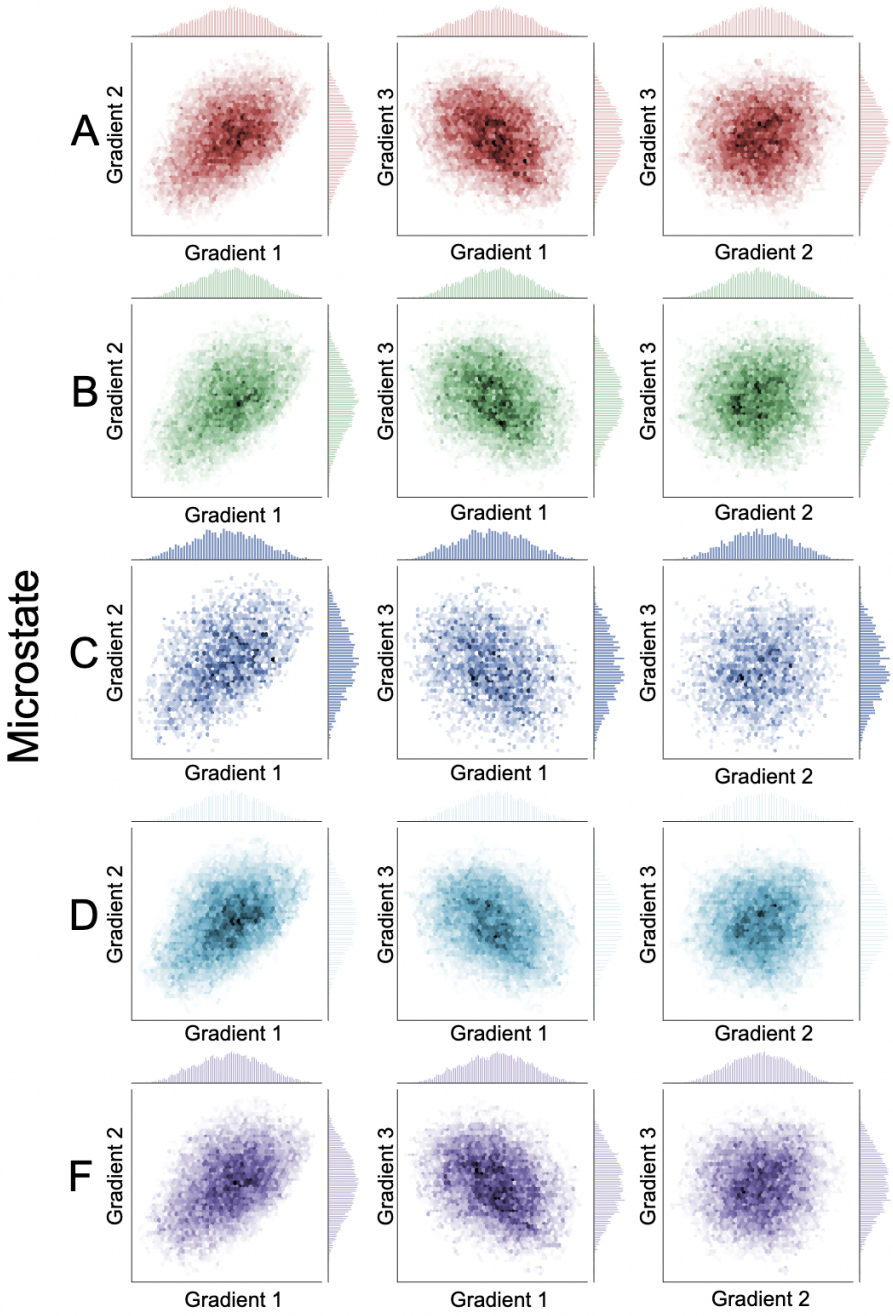

Figure 27: Kernel Density Estimate of each microstate label projected into the fMRI gradient space via interpolation across the group. Rows of subplots show microstates, columns show the three cross sections of the 3 dimensional space. Colour indicates microstate in question. Histogram shows the given microstates' density distribution across the axis.

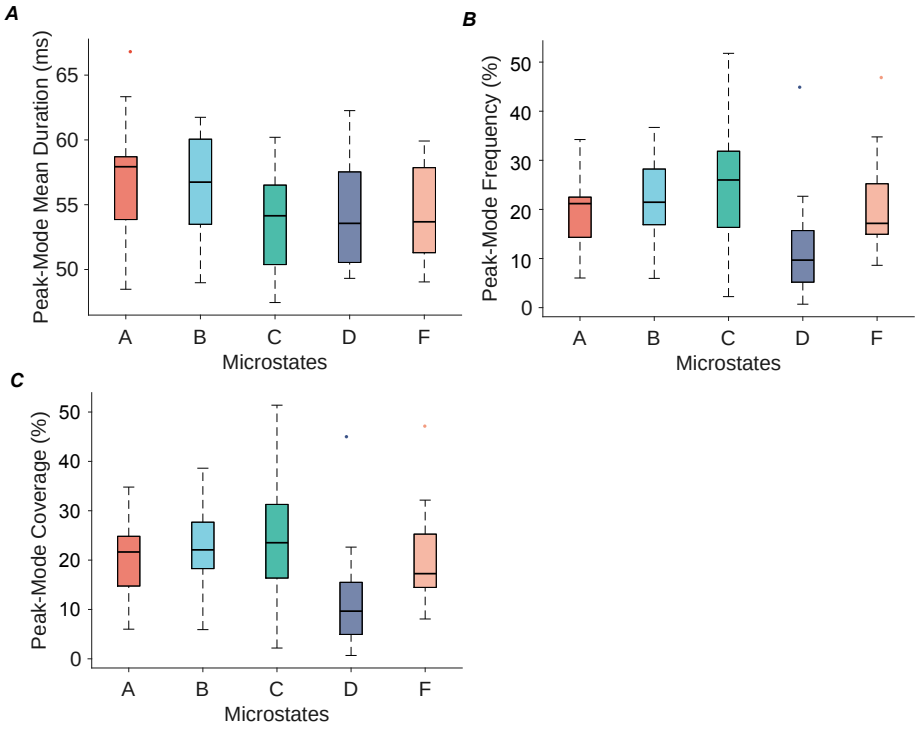

Figure 28: Microstate parameters for microstates derived using the peak-mode. (A) shows mean duration, (B) shows frequency, and (C) shows coverage. X axes indicate the microstate in question, with different coloured box plots indicating different microstates. Box plots show the parameter measure across participants. Dots indicate outlier participants. An outlier is defined as a value over the defined threshold of the upper quartile plus 1.5 times the interquartile range.

mean durations across microstates.

Peak-mode n-gram parameters were also calculated for $n > 1$ up to $n = 5$. Figure 29 shows the duration and frequency distributions of microstate 2-grams across the group. In general, the parameters observed in each microstate (Figure 28) translated to the longer n-gram lengths. The lower coverage of microstate D likely resulted in lower durations of 2-grams, which included microstate D due, perhaps due to a smaller sample of these 2-grams occurring versus others. The observation of multiple outliers in each box plot indicates a high variation of these parameters between participants, which is reflected more at higher n-gram lengths.

Figure 30 shows how the frequency of peak-mode 2-grams varies within each TR across the time series of an example participant, within a 40-second window (20 TRs). Each line shows the change in frequency of each 2-gram across TRs.

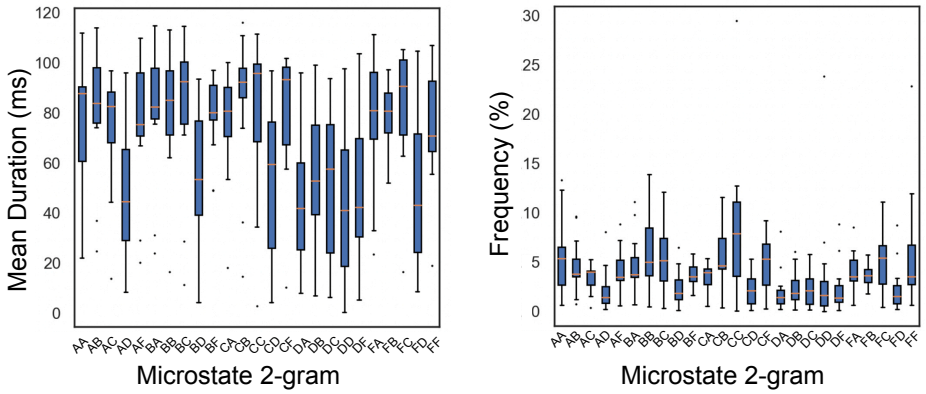

Figure 29: Peak-mode microstate 2-gram parameters observed across the group. Left shows mean duration, right shows frequency. Box plots show the distribution across participants. Dots indicate outliers. An outlier is defined as such if they were greater (or less than) the median plus (or minus) 1.5 times the interquartile range. X-axis shows the 2-gram in question. Note the inclusion of repeated microstates due to peak-mode.

In many cases, the frequency drops to zero for a given TR since the n-gram in question did not occur within the 2-second window of the given TR. In this case, the green line with high peaks is the 2-gram DD, which in this participant was more prominent than other n-grams across the time series.

Similarly, the frequency across the whole time series of TRs for an example participant is shown in Figure 31. The colour indicates the frequency of the given 3-gram for each TR. This visualisation demonstrates that in many TRs, there is no occurrence of some 3-grams and that some are rarer than others. In some cases, there are also peaks in the frequency of specific n-grams during specific TRs, which the following models aim to predict.

6.3.4 Random Forest Regressors Three Coordinates

A random forest regressor model was run for each n-gram length for each participant, using mean durations or frequencies of those n-grams as features and three-dimensional fMRI coordinates as targets. Figure 32 shows the output R^2 value for each model fit.

Most models failed to produce a strong accuracy, with any fit between the feature n-gram parameters and the fMRI coordinates, with only one instance of an R^2 value over 0.25 (frequency $n = 3$ participant 14).

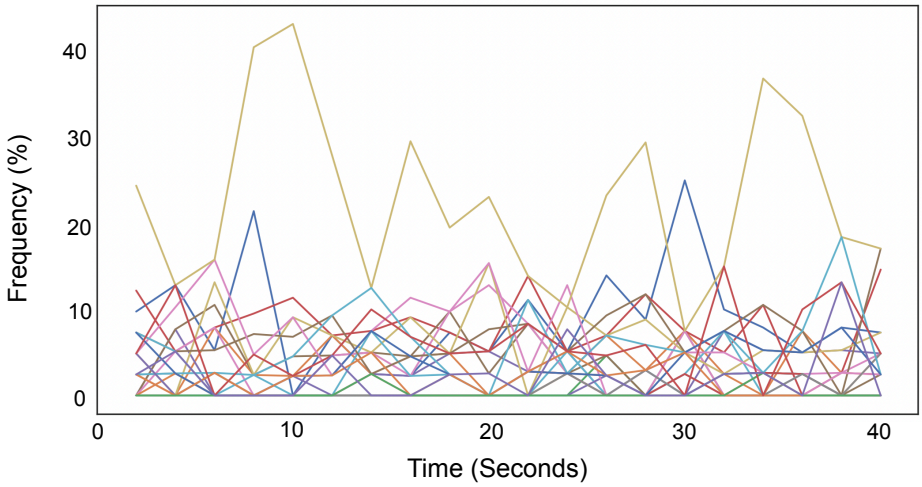

Figure 30: Example 40 second window observed in participant 18 showing the frequency of peak mode 2-grams that are occurring simultaneously during each TR. Each coloured line indicates the frequency of each peak-mode 2-gram within the example participant at the given TR.

6.3.5 Random Forest Regressors for Separate fMRI Gradient Coordinate Axes

Due to the lack of good fits using all three coordinates as targets at once, models were instead applied using one gradient axis as a target at a time. Figure 33 shows the R^2 values for each model.

More success was found using this method versus using all three coordinates together. Participant 5 showed a stronger fit for the first gradient at $n = 3$ for mean duration and frequency, with R^2 over 0.3 in each case. Interestingly, this relationship was only found for length 3. Participant 6 showed similar correlations for $n = 5$ in mean duration and frequency. Participant 14 also showed R^2 values over 0.3 for both $n = 2$ and $n = 3$.

For gradient 2, participant 7 showed some weak R^2 values for $n = 3$ and $n = 5$ for mean duration and frequency. Participant 14 showed values over 0.3 for $n = 4$ and 5. Gradient 3 showed the fewest higher value correlations (relative to the mean R^2 value across models).

There was no n-gram length for any participant with an R^2 of 0.3 for more than one gradient axis. In each case where a participant showed more than one gradient axis showed a relatively high R^2 value across different n-gram lengths. Interestingly, all models at $n = 1$ failed to find a strong fit for any participant.

6.3.6 Random Forest Post-hoc Analysis

Figure 34 shows the feature importance and permutation feature importance plots of models from participant 14 at n-gram lengths where R^2 was more sig-

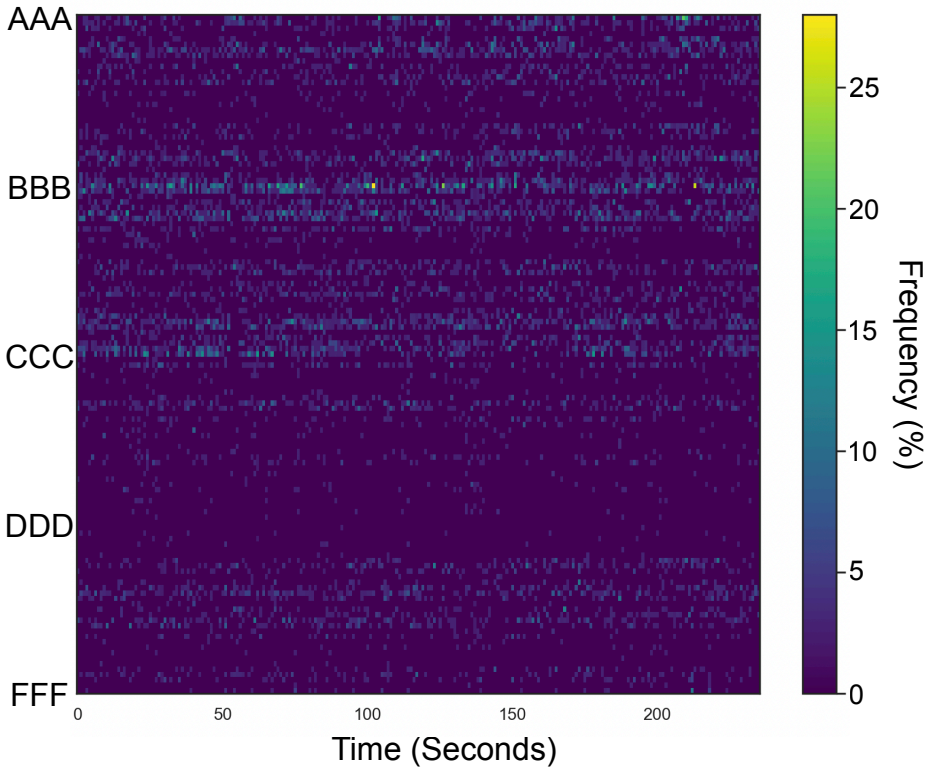

Figure 31: Matrix visualising participant 10's peak-mode 3-gram frequency per TR for all 3-grams. Each row of the matrix shows the frequency of each 3-gram within each TR column. Colour bar indicates the frequency of each 3-gram across the time series per TR.

nificant than 0.3 for Gradient 1. Feature importances sum to 1. Hence, each feature's importance score should be considered as contributing to the resulting R^2 value. Figures 34A and 34C show said values for $n = 2$ frequency, with Figures 34A and 34C showing $n = 3$ frequency. Figures of feature importance and permutation importance for all subjects at all n-gram lengths can be found here.

In the example of Participant 14, models showed a higher than average R^2 score for both $n = 2$ and 3 (Figure 32). Interestingly, the 2-gram CA could contribute to differences at the $3-gram$ level, with CAD showing higher than average feature importance. It can also be seen that, generally, n-grams containing microstate D have a lower importance. This may be due to the group's lower coverage and frequency of microstate D in general (Figure 28).

Table 1 shows the n-grams in each participant that contributed the highest feature importance for all models where R^2 was more significant than 0.3 for each gradient axis.

Highest Contributing Importance	Gradient 1		Gradient 2		Gradient 3	
	Mean Duration	**Frequency**	**Mean Duration**	**Frequency**	**Mean Duration**	**Frequency**
RS_005	CCC FBC	FBC	-	-	-	-
RS_006	CBAAD CDDBD	CBAAD CDDBD	-	-	-	-
RS_007	-	-	-	BFCFC	-	-
RS_008	-	ABCA BCAF	-	-	-	-
RS_013	-	-	-	-	AFFDD	-
RS_014	-	BB FF CBA CCD	CCCBF FFDFC	CFBB FFFA CCCBF FFDFC	-	-
RS_015	-	CBCBF CCCDA	-	-	DBCB	CCFB DBCB
RS_018	-	-	-	AAAAB	-	-

Table 1: Table showing the n-grams with the highest feature importance for each model where $R^2 > 0.3$. Each row gives a participant, with each column shows the n-grams for each gradient individually, either mean duration or frequency models. The n-grams specified are the top two contributing in each case for the given length. If there was only 1 n-gram with high contribution, then that is report.

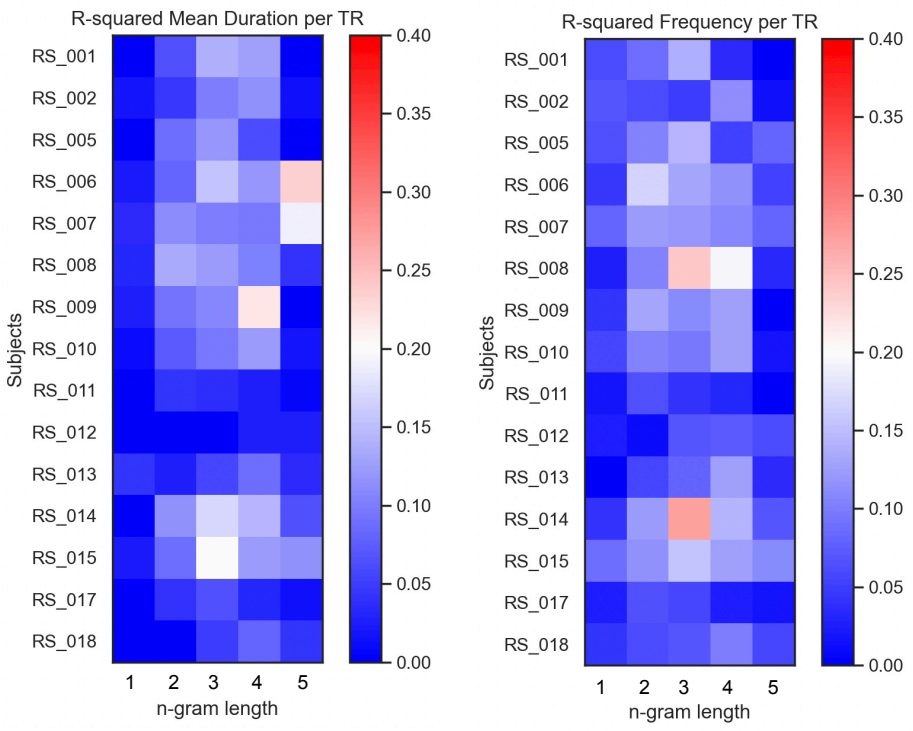

Figure 32: R^2 values calculated for each random forest regression model using best fitting hyper-parameters in each case. Left matrix shows tests which used n-gram mean durations as features, right matrix shows tests which used n-gram frequencies as features. Y-axes of each matrix shows the participant tested, and x-axes show the n-gram length of the features used. For example, at $n = 3$ for mean duration of participant 7, the features used in the regression model were mean durations of 3-grams per TR. All three fMRI gradient coordinates were used together as the target in these models.

Eight out of the fifteen observed participants showed R^2 values greater than 0.3 across the gradient axes. Two of those 8 participants had models over the 0.3 threshold on more than one gradient axis. Multiple instances of n-grams containing BC are apparent in three of the eight participants on gradient 1. Other than this, there is a lack of consistency between participants for each n-gram length.

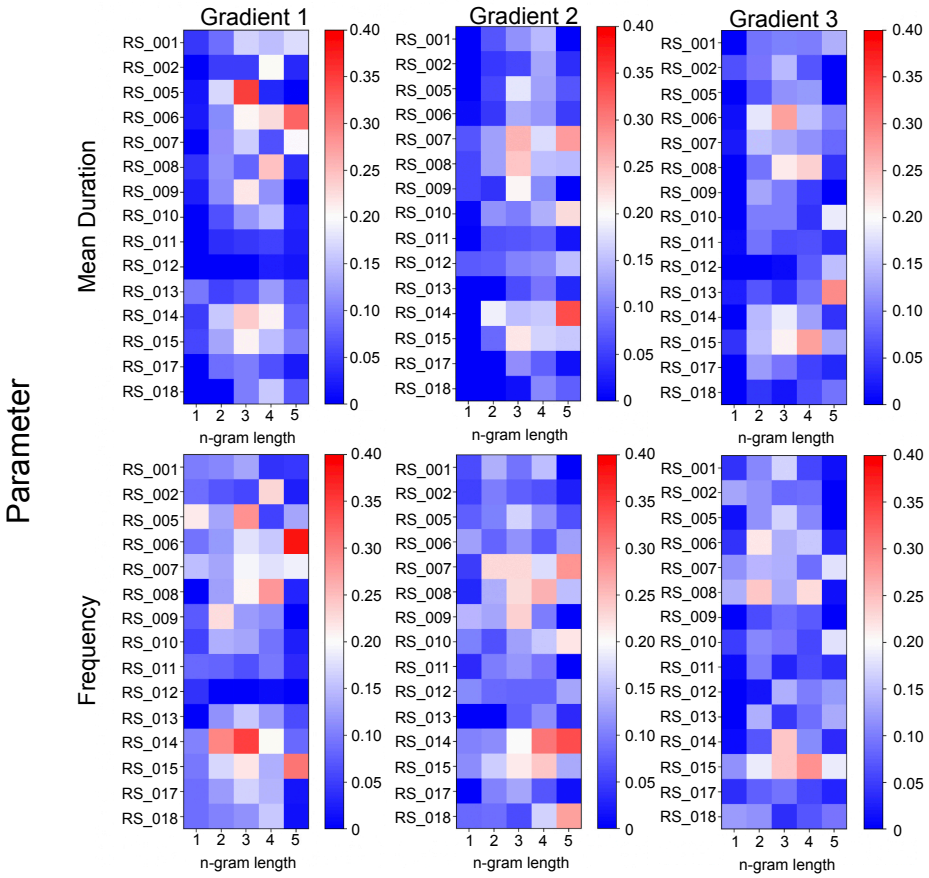

Figure 33: R^2 values calculated for each random forest regression model using best fitting hyper-parameters in each case. Top row shows tests which used n-gram mean durations as features, bottom row used n-gram frequencies as features. Y-axes of each matrix shows the participant tested, and x-axes show the n-gram length of the features used. For example, at $n = 3$ for mean duration of participant 7, the features used in the regression model were mean durations of 3-grams per TR. The columns show the target coordinate for each model. In these cases, the gradient coordinate axes were considered individually to simplify the fitting process for the model.

6.4 Discussion

6.4.1 fMRI Gradient Space

This study used an fMRI gradient space, which takes correlations with each of the derived gradients of Margulies et al. (2016) as coordinates in a multidimensional space. A multidimensional coordinate in the low-dimensional space can be projected back up into the original space post hoc, allowing for a drastic

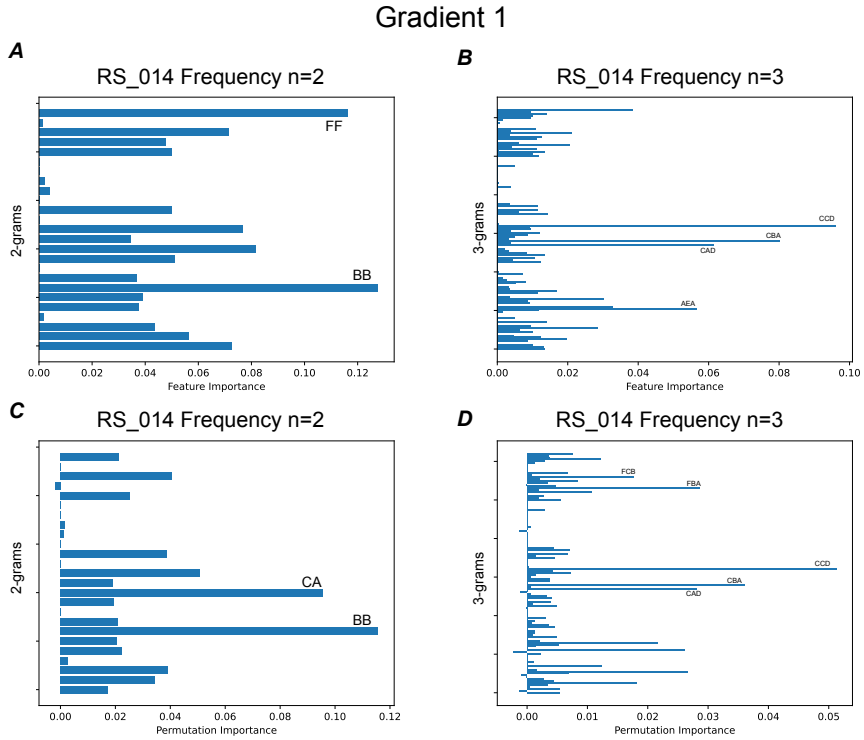

Figure 34: Feature importance and permutation feature importance of EEG microstate n-gram frequency for individual participant at $n = 2$ and 3 across gradient 1. (A) and (B) show the feature importance of 2-grams and 3-grams for the participant respectively. (C) and (D) show the permutation importance of 2-grams and 3-grams respectively. The n-grams with the highest scores in each plot are indicated with a label.

reduction in the spatial dimensionality of the fMRI whilst still retaining the complexity of global activation patterns, as well as temporal dynamics (Brown et al., 2022).

Figure 26 shows the density distribution of fMRI TRs across the group within the gradient space. In general, the group appears to prefer the negative to positive $x = y$ plane between gradients 1 and 2. This is similar to the previous study in Figure 17, where two of the three CAP pairs were at either end of this plane. The lack of TR coverage along the $x = -y$ line of this slice points to a lack of visual cortex and Default Mode Network (DMN) co-activation in general in the group (Margulies et al., 2016).

The density distribution along each axis tended to be relatively normal for each participant, with some participants showing a preference for the poles of the first axis. The principle axis has been called the "sensory association axis" to highlight its separation of regions associated with sensory processing and

cognition (Margulies et al., 2016). The few participants showing this polarisation of density at either end of the axis suggest a more robust separation of the sensory and cognitive regions. Such polarisation is less common in the second and third gradients, potentially pointing to less separation of activities on the sensory axis (gradient 2) and task axis (gradient 3) (Mesulam, 1998).

It is worth highlighting that since the gradient axes used for correlation values here are taken from a past study (Margulies et al., 2016), it may be more beneficial in future studies to derive data-driven gradient axes so that the axes that the time series are related to are more directly related to the dataset, in a similar vein to how recent studies suggest use of data-driven EEG microstates (Michel & Koenig, 2018).

Further investigation into the fMRI gradient space is advised in general. For example, other measures, such as the directions moved in the space from TR-to-TR, may uncover general patterns of global state transitions. These movement patterns could also be related to EEG microstates.

6.4.2 Microstates in fMRI Gradient Space

Alignment of the EEG time series with the fMRI time series via the use of the HRF allows for a visualisation of the relationship between microstates and fMRI that has not yet been achieved. Figure 27 shows the density of microstate occurrences within this fMRI derived space.

The distribution of microstate occurrences appears to follow the distribution of fMRI coordinates. This suggests a uniformity of occurrences of each microstate when comparing their activity to global patterns of activity in fMRI recordings, something touched on in a previous Chapter 2.

Past studies that have investigated the relationship between microstates and fMRI have generally identified specific networks which show correlations with the microstates (Britz et al., 2010; Xu et al., 2020; Yuan et al., 2012). The methods used in these studies are equivalent to the General Linear Model (GLM) reported in Figure 13. Such an approach does not attempt to draw a relationship to global activation patterns in fMRI, but rather more static parcellations of states.

The normal density distribution of each microstate within the fMRI gradient space (Figure 27) suggests that these patterns do not appear when considering whole brain activities. A recent study which did investigate global fMRI patterns and was able to predict the dFC state (Abreu et al., 2021). This approach, however, generated unique sets of microstates per TR, rather than using the conventional approach of deriving microstates across the group, making comparison to these results difficult.

The results shown here highlight the difficulty inherent in drawing a relationship between EEG microstate occurrence and fMRI global activation pattern, as a simple increase in the occurrences of specific microstates in particular regions of the gradient space does not occur.

6.4.3 Peak-Mode EEG Microstates

As discussed in Chapter 3, the lack of definition of sequence modes in the EEG microstate literature highlights a need for solid definition. Here, peak-mode microstates were derived and used to highlight the validity of their derivation and show the differences between the conventional event-mode and this peak-mode.

First, worth highlighting is the difference between parameters between peak-mode (Figure 28) and the event-mode microstates (Figure 12) derived in the previous study. Although the same participants and the same data were used in both studies, differences in parameters are apparent. Microstate C had low occurrences in the event-mode relative to other microstates. In the peak-mode, however, microstate C showed the highest average frequency and coverage amongst all microstates, and microstate D generally showed relatively low parameters. Coverage is different between the event-mode and peak-mode values in Chapters 5 and 6 respectively due to event-mode being derived using back-fitting, and peak-mode being derived using interpolation (see Chapter 3 for a discussion).

It is the case that peak-mode is a more accurate representation of GFP peak dynamics than event-mode methods. GFP peaks are used to derive microstate classes (Pascual-Marqui et al., 1995) due to their high Signal-to-Noise Ratio (SNR) (Lehmann et al., 1987). Peak-mode sequencing allows for investigating within-cluster transitions between GFP peaks, which is impossible in event-mode.

When considering that microstate parameters differ significantly between peak-mode and event-mode, as well as accepting that peak-mode may more accurately represent microstate dynamics and therefore may be more suited for syntax investigation, highlights that adoption of interpolation over back-fitting (see Section 3, and hence the use of peak-mode over event-mode, may be beneficial for microstate syntax investigations.

6.4.4 Predicting fMRI Gradient Space Coordinates Using EEG Microstate n-Gram Parameters

Of the seventy-five models for mean duration and seventy-five for frequency (fifteen participants by five n-gram lengths), sixteen had R^2 values greater than 0.3. These are the best-performing models; in each case, the n-grams with the highest feature importance were reported. The highest contributing n-grams in each case were not shared between participants. This may be due to individual differences, but further investigation is required to understand the significance of the associations between gradient axes coordinates and EEG microstate n-gram parameters.

A shortcoming of this methodology is the difficulty of reporting the influencing n-grams. The number of feature n-grams in each model is high, meaning identification of the influence of each is difficult. The use of n-grams in this context is also problematic because each n-gram length is assumed to be inde-

pendent of every other n-gram, which is not the case. It may be beneficial in future analysis to use n-gram parameters across n within a single model, but this will require careful consideration.

Additionally, future use of this methodology should apply a permutation step where models attempt to fit the same models using shuffled n-gram labels to establish a baseline R^2 value to understand the model results' significance better. Using multiple n-gram parameters as a set of features in the same model will also be investigated as a potential method.

It should also be highlighted that fMRI gradient space allows for the opportunity to investigate BOLD signal dynamics. Future work which applies this method to understand the functional significance of EEG microstates may simultaneously investigate the dynamics of both the EEG and fMRI activity.

Since there were no n-grams which showed a good fit across gradients, there was no need to pinpoint a coordinate where the peak of the given n-grams parameter was. A back-fitting process would take place in future uses of this methodology should a well-fitting model consider more than one fMRI gradient at a time. This particular advantage of this approach could not be utilised in this instance.

The number of models needed to report results across n-gram lengths and the need to apply models for each parameter individually point to an issue of multiple comparisons. It may be more beneficial to use a model that initially considers the continuous EEG signal and then considers the syntax of microstates post hoc to avoid these pitfalls. The following study attempts to implement such a method.

7 Study 3 - EEG Gradient Space as a Method for the Investigation of Microstate Sequence Patterns between fMRI States

Past studies that investigated the relationship between Electroencephalography (EEG) microstates and Functional Magnetic Resonance Imaging (fMRI) Blood-Oxygenation Level Dependency (BOLD) signal, including those here (Chapters 5 and 6), have done so by drawing relationships directly, either through application of modelling techniques such as voxel-wise General Linear Model (GLM)s, or by using microstate and microstate n-gram parameters as predictors of fMRI signal. Here, a novel method is introduced, which uses microstates as post hoc labels for the investigation of the measured EEG sequence, rather than microstates being the subject of analysis. This pivot allows for a more robust modelling of EEG signal dynamics that can be used in future studies to understand EEG microstate syntax.

7.1 Introduction

Past investigations of microstate syntax generally identified differences in transition probabilities between microstates (Michel & Koenig, 2018), or differences in occurrence rates of microstate n-grams; short sequences of microstates of length n (e.g., von Wegner et al. (2017) or Chapters 5 and 6). The investigation of microstate syntax has demonstrated differences in clinical populations (Lehmann et al., 2005), as well as between specific groups (Schlegel et al., 2012), but the underlying mechanisms of microstate syntax are yet to be understood.

Methodologies previously applied which attempt to understand microstate syntax do so using the microstate sequence as an input into analyses. Methodologies will generally apply analysis which accounts for the discrete sequence of states, limiting investigation to either parameters which describe microstates, or which describe short sequences of microstates. Whilst there have been identified differences in n-grams and their parameters between populations (Schlegel et al., 2012; von Wegner et al., 2017), as well as between simultaneously occurring fMRI states (Chapter 5) the number of possible n-grams that can be investigated as n increases leads to a combinatorial explosion, which limits the length of the microstate sequence that can be investigated.

von Wegner et al. (2017) demonstrated that information is retained within clock-based microstate sequences up to a length of $1000ms$, and that microstate sequences are therefore non-Markovian. In the context of the conventional event-mode microstate being approximately $100ms$ long, a microstate 10-gram would be required to account for this length. The number of possible microstate 10-grams that could occur assuming event-mode sequencing (see Chapter 3) and $k = 4$ (canonical microstates; Lehmann et al. (1987)) would be $4 \times (4-1)^{(10-1)} = 262,144$. This is referred to by von Wegner et al. (2017) as the "dictionary of possible words" of the given length. This does not even consider the possibility of information transmitted *between* n-gram lengths. It may be the case that information which is gained from investigating previous microstates in the sequence, that for example a 4-gram may inform the prediction of the following 5-gram. Previous studies which have investigated microstate sequences have not accounted for this (including the previous two chapters). The number of possible n-gram combinations as well as the number of microstate combinations, makes understanding syntax through n-gram analysis difficult.

For this reason I propose an adjustment regarding the target of analysis. Microstate sequences are derived primarily from the occurrence of EEG Global Field Power (GFP) peaks. These peaks of activity are points of spontaneous activity on the time series of EEG recordings. The EEG time series is initially a continuous recording of electrical activity before the conversion into a discretised sequence of microstates, where the difficulties of n-gram combinatorics arise. Hence, I suggest that the continuous time series of EEG first be investigated in relation to fMRI, with a reduction of the EEG time series to a discretised sequence of microstates being applied post-hoc. Such a method bypasses problems regarding associations between microstates, and instead conducts the association between EEG and fMRI initially, and categorises those associations

using the microstates as labels, rather than making the labels the subject of analysis.

Here, as a novel means of investigating the EEG signal, I propose the EEG gradient space. I derive a set of EEG gradient axes, using the methodology used to derive fMRI gradients in previous studies (Margulies et al., 2016) and hence reduce the EEG time series into a trajectory which occupies the "EEG gradient space". The coordinates of the continuous trajectory of EEG are used as features in a classifier, where the EEG time series is segmented by simultaneously occurring fMRI Repetition Time (TR)s, and the classifier attempts to predict the corresponding fMRI Co-Activation Pattern (CAP) label (see Chapter 5 for an explanation of CAPs). This method is a preliminary means of bypassing the n-gram problem, and proposals for improvements and developments are discussed.

7.2 Methodology

7.2.1 EEG Gradient Space and fMRI Alignment

EEG gradient space "sub-trajectories" that occurred per TR were then used as features to predict an fMRI CAP label (defined in Chapter 5) that was associated with each TR. A visualisation detailing the process to predict the fMRI CAPs is shown in Figure 35.

A dimensionality reduction of the EEG time series was applied that retained the continuous signal observed. The pre-processed EEG time series of the fifteen participants was concatenated, and each channel was correlated with every other channel (Figure 35A), resulting in an EEG channel connectivity matrix, $30x30$ (30 EEG channels). Principle Component Analysis (PCA) was then applied to this connectivity matrix, and the first three components were used as prototype gradient axes for the dataset (Figure 35B). This process was applied using the concatenated time series of all participants, as well as the GFP peaks alone.

Following this, each individual time point of each participant (or GFP peak) was correlated with each of the gradient axes, and the correlation value for each time point was used to generate a correlation time series. The time series for each of the gradient axes resulted in what is referred to here as a "gradient trajectory", where trajectory refers to the movement of the observed EEG through a space defined by the derived gradients. Each participant time series was placed in the space individually (Figure 35C).

The EEG signal in this space was then aligned with the observed fMRI using an Haemodynamic Response Function (HRF) with a peak at 6 seconds. After cleaning, the EEG was observed at 250Hz, and the fMRI had a TR of 2 seconds. Hence, for each fMRI TR, there are 500 time points of EEG when considering the whole time series. Following the alignment, the EEG gradient trajectory was segmented into these 500 time point long sub-sections. A trajectory of EEG within the space was therefore isolated for each individual fMRI TR.

GFP peaks were isolated during each simultaneously occurring TR, and a trajectory between these EEG peaks (rather than all time points) was also used.

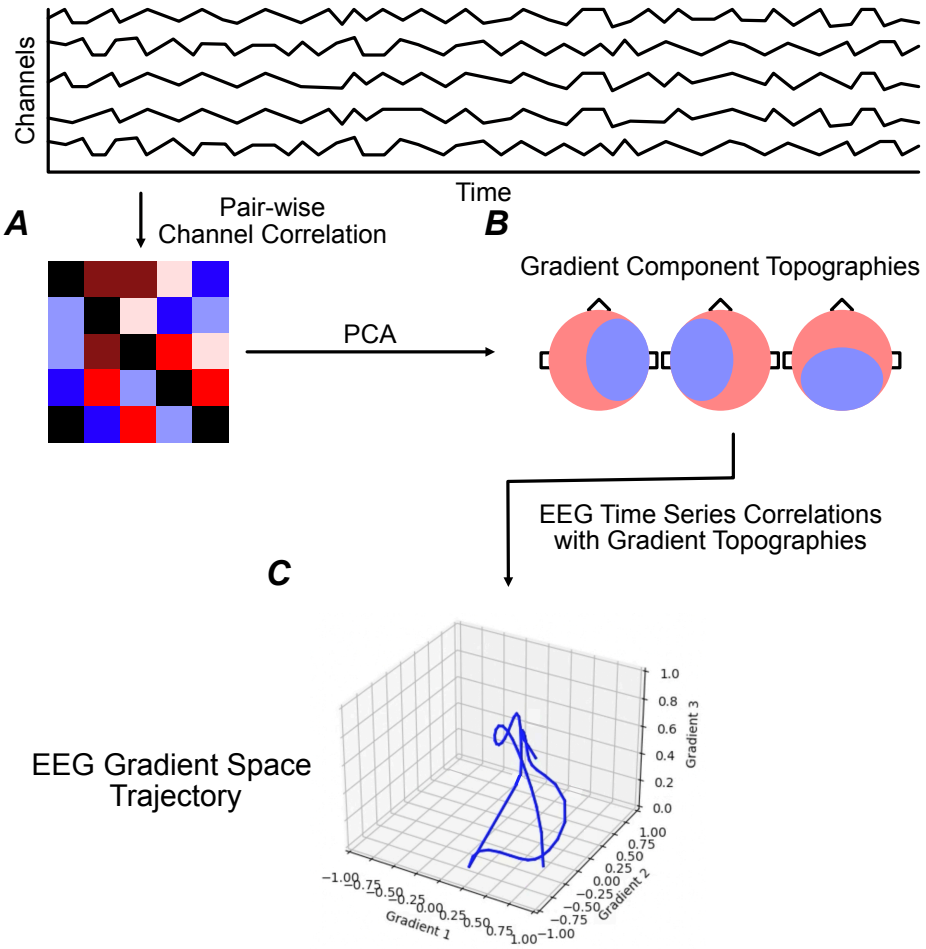

Figure 35: Pipeline of data preparation to derive EEG gradient space. The whole time series of each EEG channel (top) is correlated with every other channel (A). A PCA is applied to this "connectivity matrix" and the highest contributing components are defined as the EEG gradients (B). The time series of EEG of an individual is then correlated with each of the gradients (in this toy example, three) in order to derive an EEG gradient space (C). Each axis is between -1 and 1 based on correlation values.

In this case, a different number of EEG coordinates were isolated for each fMRI TR.

7.2.2 Predicting fMRI CAPs using EEG Gradient Space Trajectories

The models used to predict fMRI CAP labels were Recurrent Neural Network (RNN)s. An RNN is an artificial neural network designed for sequential data processing where the order of the input elements is important (Yu et al., 2019). In general, RNNs have the ability to maintain a hidden state that captures information about previous inputs in the sequence. For this application, Long Short-Term Memory (LSTM) layers were used in the RNNs. LSTMs can selectively remember or forget information over long sequences, making them well suited for tasks that require capturing dependencies over extended periods (Hochreiter & Schmidhuber, 1997).

Tensorflow was used to build an RNN (Abadi et al., n.d.), an open source Python package that can be used to build neural networks and machine learning algorithms. The model was a simple LSTM model. An input layer was followed by the LSTM layer. The LSTM layer returned full sequences, and included a kernel regulariser to prevent over-fitting. This was followed by a self-attention layer which was applied to the LSTM layer.

A self-attention layer enables the model to weigh the importance of different sections of a sequence with respect to each other, allowing it to focus on relevant parts of the input when making predictions. In this case, where the self-attention layer was used it was used in conjunction with the LSTM layer. LSTMs are effective handling sequential information over extended periods. Self-attention is good at capturing local dependencies. A combination of the two allows for the possibility of capturing relevant information at both the local and long-range levels.

The self-attention layer was followed by a flattened layer, resulting in an output dense layer with six units, each unit corresponding to a CAP label. The model was fitted to whole time series trajectories, as well as GFP peak trajectories. Thirty training epochs were used for each model fitting. Sixty percent of participant TRs were used for training, twenty percent for validation, and ten percent for testing. Models used categorical cross-entropy loss as a target to minimise. The coverage of each CAP across a participant was used to control for an offset in the proportion of categories.

In context, a single input to the RNN would be the 2 second *EEG* gradient space, with its associated CAP label. The use of an RNN aimed to capture patterns of activity within the trajectory at different ranges, with the target being to identify patterns of EEG activity that may be unique to each fMRI CAP.

A range of hyper-parameters were searched across for each model during fitting. These were, the number of LSTM units (8, 16, 32, 64, 128), the kernel regulariser weight (0.01, 0.05, 0.1) and the dropout layer weight (0.1, 0.25, 0.5). The precise architectures with the resulting hyper-parameters after fitting are reported in the following section.

7.3 Results

7.3.1 EEG Gradient Space Axes

The groups EEG time series' were concatenated and used to generate a set of data-driven EEG gradient axes by correlating each channel time series with every other channel time series. Figure 36A shows this output "connectivity" matrix. Figure 36B shows the output of the PCA as cumulative explained variance of each added gradient component. The first component explained 59% of the variance, the second explained an additional 35%, and the third explained 3%. All other components following this explained less than 1% variance each. For this reason, the first three components were used for further analysis, the three collectively explaining 97% of the variance.

Figure 36C shows the topographies of each of the first three gradients. Gradient 1 is a slightly skewed hemispheric split, with poles at the most lateral positions. Gradient 2 is a fronto-occiptal configuration and gradient 3 a dorso-ventral configuration (ventral here being the ventral part of the scalp, which are the most superficial regions of the scalp recorded).

7.3.2 EEG Gradient Space Trajectory and Density

Following derivation of the gradients, each of the gradients were conceptualised as axes of a multidimensional space. The time series of each participant was correlated with each of the axes at each time point, generating a trajectory EEG activity through the EEG gradient space. Figure 37 shows the kernel density estimate of each participant along each axis, with each colour distribution representing a participant.

Figure 38 shows same density estimate across participants, visualised as cross sections of the three dimensional space. Darker colours indicate higher density.

In general across the group there is preference for the poles of the first gradient, with a less profound preference also indicated across gradient 2. Interestingly, the EEG trajectories occupy an ellipsoid within the EEG gradient space.

7.3.3 LSTM Prediction of CAPs

Following hyperparameter tuning, the optimal RNN configuration across those tested for the models which used the whole EEG time series as input, is found in Table 2, using a regulariser value of 0.01 in the LSTM layer. Sixteen nodes were used. This was the case for both the whole time-series model, and GFP peak model.

The training and validation loss and accuracy are found in Figure 39 for application to the whole time series for each participant. Figure 39A shows the performance on the training segment. Loss was generally low in the training data. After thirty epochs, all participants showed an accuracy over 60% of CAPs during training. Validation did not perform as successfully (Figure 39B). Loss was substantially larger across epochs during validation versus training.

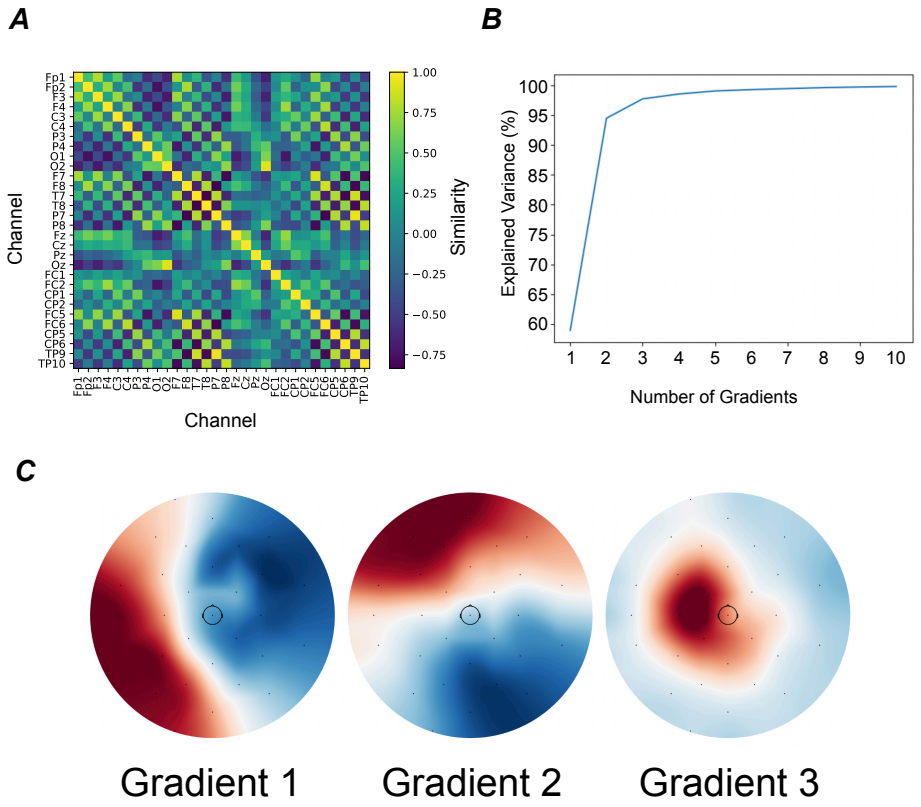

Figure 36: EEG channel time series similarity matrix and EEG gradient space axes. (A) shows the similarity matrix between recorded EEG channels across all participants. Both axes are labelled with channel names. Colour bar indicates similarity. (B) shows the cumulative explained variance of each of the gradients as the number of gradients is increased. (C) The EEG topographies of the first three gradients. Gradients are ordered by their explained variance. Gradient 1 shows a hemispheric split, gradient 2 a frontal-occipital split, and gradient 3 a dorsal-ventral split.

Accuracy did not improve during validation across epochs, with a maximum value across participants of 40%. The average test accuracy across participants was similar to the results found in the validation data, around 19.4%.

A similar result was observed in the trajectory which only used GFP peak coordinates. Figure 40 shows the loss and accuracy of this model. Trajectory of learning across epochs is similar to the whole time series, with the subtle difference of a slightly lower loss in both the training and validation steps. Average test accuracy was also similar, at 17.4%. Both test accuracy results are close to what would be considered guessing ($1/6 = 16.7\%$).

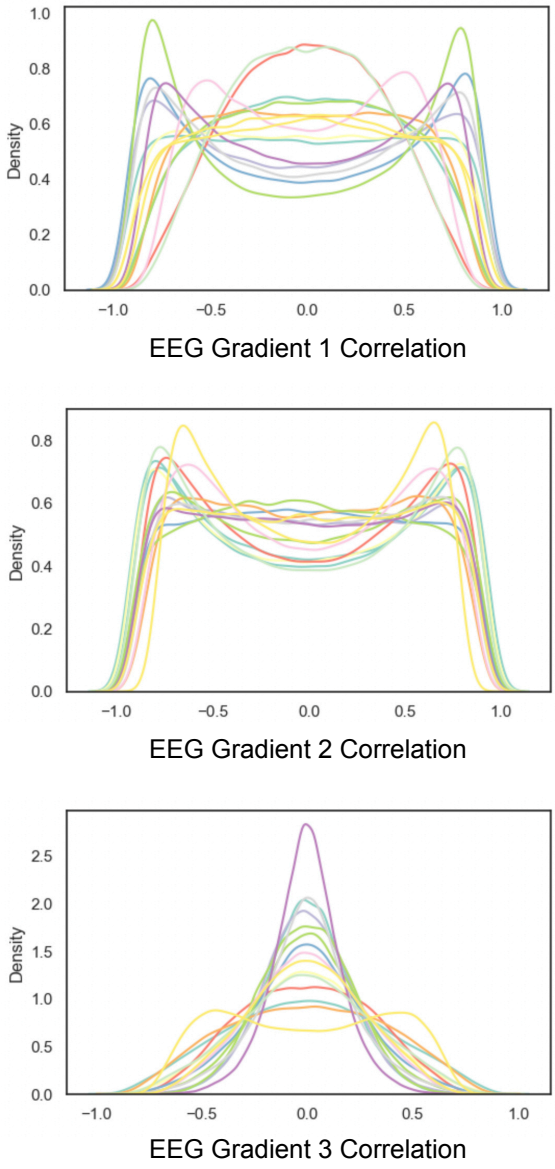

Figure 37: Kernel density estimate of individual participant EEG gradient trajectories within the three dimensional space. Each of the three subplots shows the kernel density estimate of the given gradient axis.

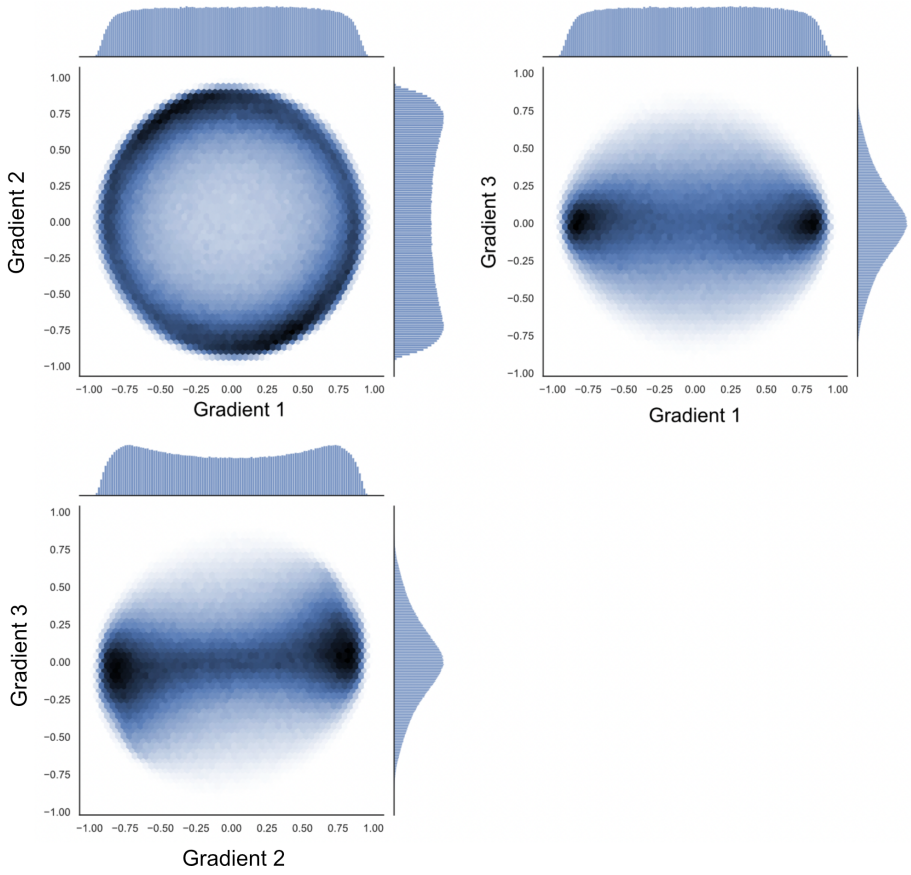

Figure 38: Kernel density estimate of all participant EEG gradient trajectories within the three dimensional space. Each of the three subplots show a plane of the three dimensional space.

7.3.4 EEG Microstates within the EEG Gradient Space

In order to understand how the EEG gradient space relates to EEG microstates, GFP peaks across participants were subject to modified k-means clustering. Figure 41A shows the output of clustering in the gradient space. Each colour indicates a cluster centre. There are two points for each microstate, since the polarity of the microstates is invariant, which are connected by a line.

Figure 41B shows the reconstruction of the coordinates of each cluster centre as scalp topographies. The topographies found appear to be equivalent to meta-microstates A, G, F, C and B respectfully, but the 5th cluster centre has a less anterior/posterior delineation between poles than microstate B. Microstate D does not seem to appear here, likely due to its gradient 3 weighted topography, which contributed a low explained variance to the derivation of axes.

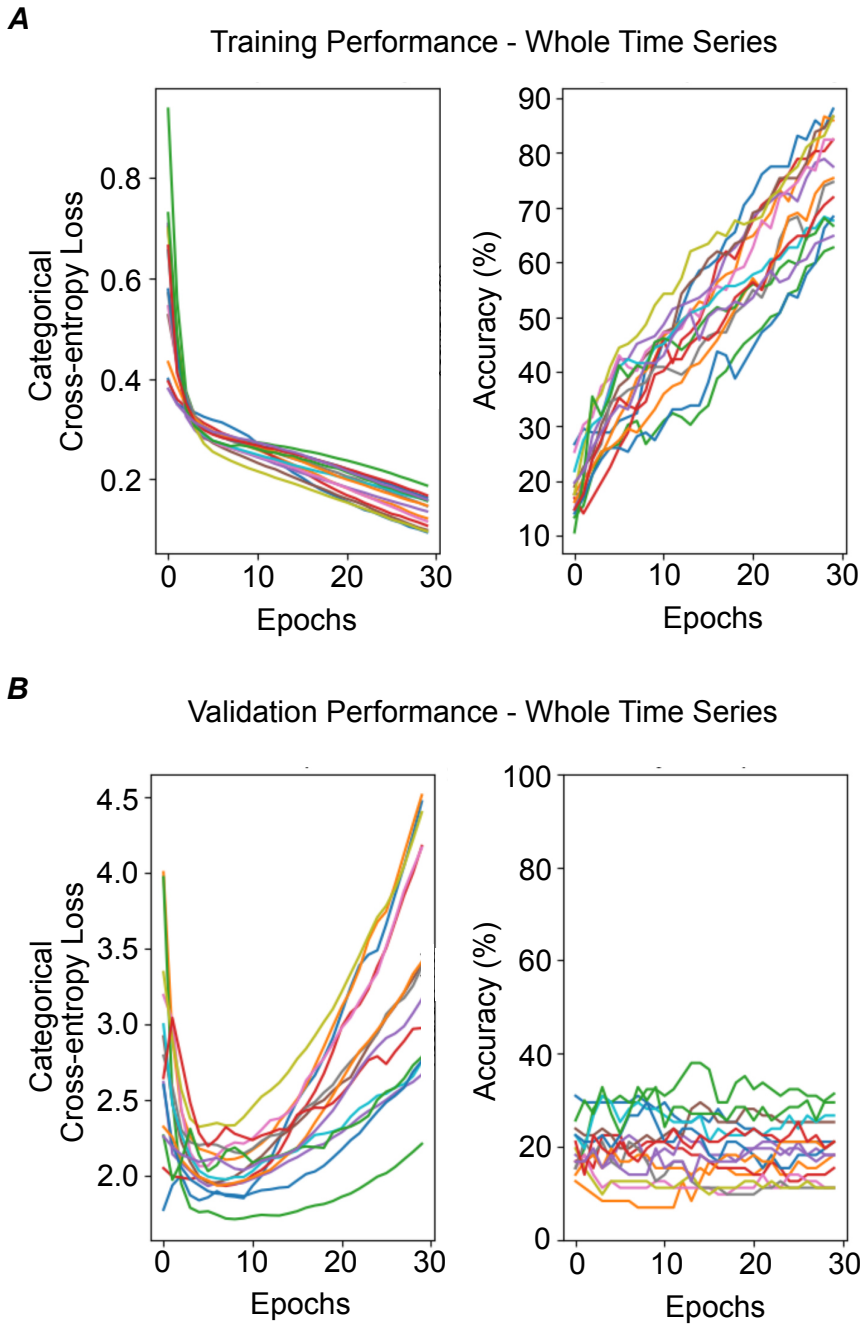

Figure 39: Performance of RNN model using whole EEG gradient trajectory time series as input. (A) Training categorical cross-entropy loss (left) and accuracy. (B) Validation categorical cross-entropy loss (left) and accuracy. Epochs on x-axis denote number of training epochs.

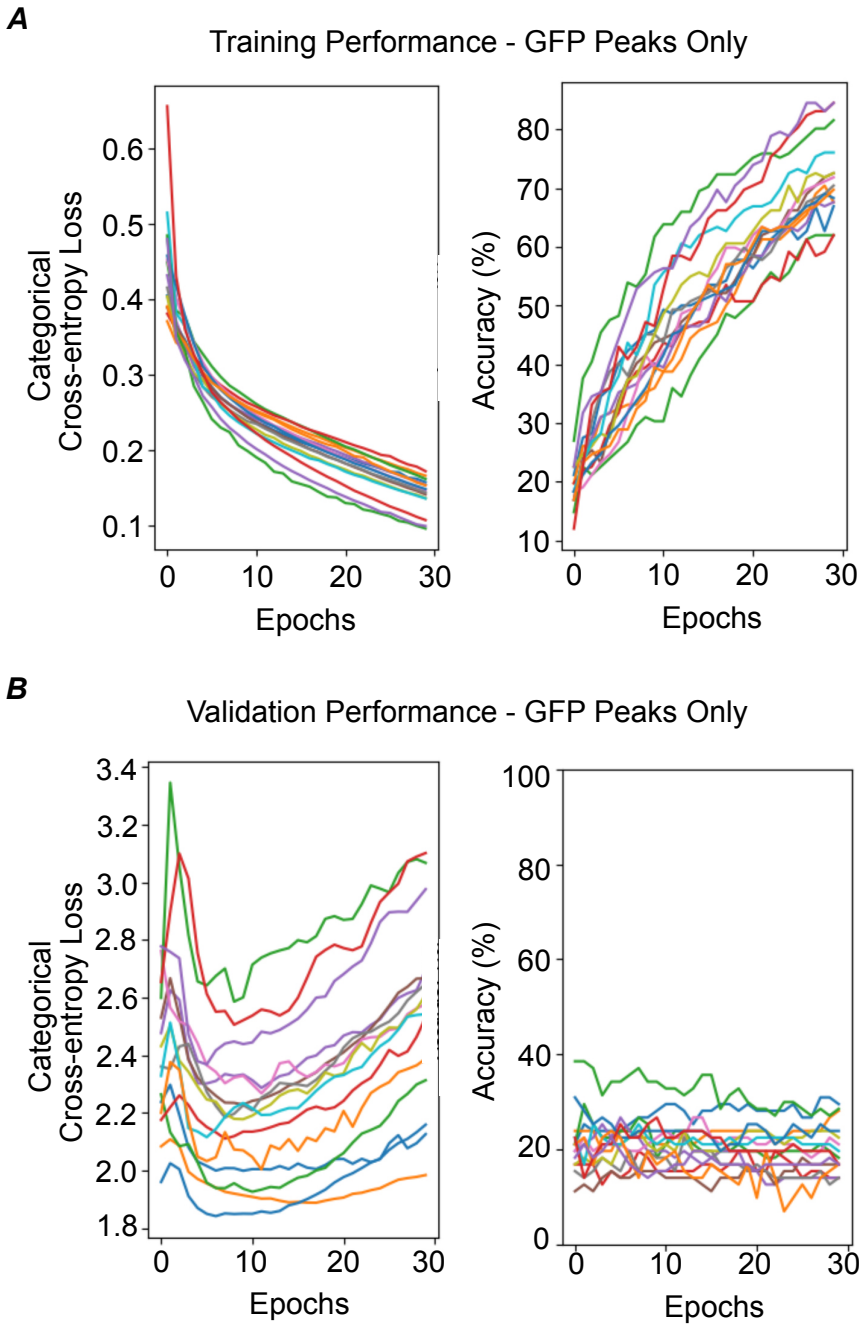

Figure 40: Performance of RNN model using GFP peaks of EEG gradient trajectory time series as input. (A) Training categorical cross-entropy loss (left) and accuracy. (B) Validation categorical cross-entropy loss (left) and accuracy. Epochs on x-axis denote number of training epochs.

Layer	Output Shape	No. Parameters	Connected to:
Input Layer	(None, 500, 3)	0	
LSTM	(None, 500, 16)	1280	Input Layer
Attention	(None, 500, 16)	0	LSTM
Flatten	(None, 8000)	0	Attention
Dense Output	(None, 6)	48006	Flatten

Table 2: Optimal configuration of the LSTM model applied to the whole EEG time series. Each layer's output shape and number of output parameters are shown in consequent columns. The pipeline of the RNN is also shown in the "connected to" column. Note the input shape in this case as (None,500,3), where "None" indicates the undefined number of TRs taken as input, and the (500, 3) denoting the 500 EEG time points per TR that are used in the whole EEG time series model, and the 3 gradient axes.

Comparing each of the derived cluster centres here each gradient axis makes clear that most cluster centres are within a ring that is at the boundary of the ellipsoid found in Figure 38.

With each of the GFP peaks within the EEG gradient space, each of their original microstate labels (see Chapter 5) were isolated within the gradient space to understand microstate distribution within the space. Figure 42 shows the kernel density estimate of each of the original microstates within this space.

Each microstate occupies an area of the space, with each microstate showing a cone-like distribution of occurrences out from the origin, and high density at the edge of the ellipsoid.

7.4 Discussion

7.4.1 EEG Gradient Axes

The EEG gradient axes derived in Figure 36 are a novel means of reducing the spatial dimensionality of the EEG connectivity matrix into a set of components. The three axes explain 97% of the variance within the EEG channel connectivity matrix. In general, the gradients demonstrate a preference for channels to activate with those closest to them. The pole locations on the topographies of each gradient could be considered similar to the three dimensions of space (i.e., left to right, front to back, top to bottom). Any pair of poles could reasonably be conceptualised within a space where each of these gradients are the axes, making visualisation intuitive.

The connectivity matrix calculated here using simple correlation coefficients between EEG channel time series' is preliminary means of investigating the EEG gradient space. The field of EEG connectivity analysis is a large one, with many options for deriving connectivity measures between channels (Bakhshayesh et al., 2019; Haufe et al., 2013). Future work will investigate the utility of each of these methods with application to gradient analysis, and compare the spaces derived therein.

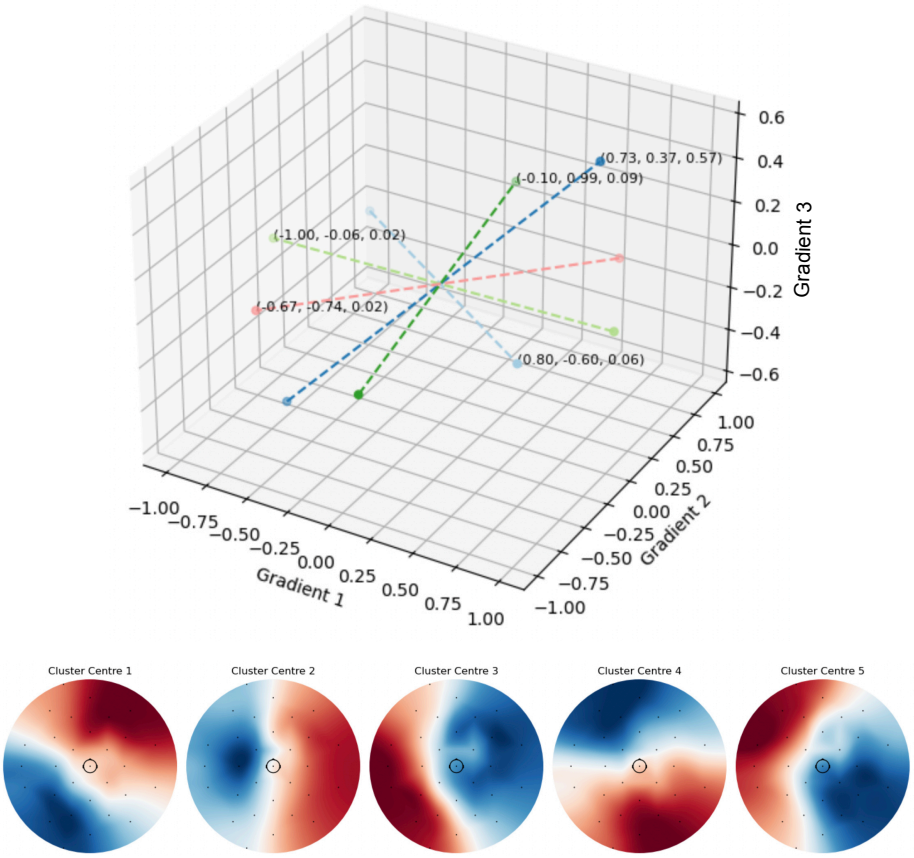

Figure 41: (A) Output of modified k-means clustering with k=5 in the EEG gradient space, using coordinates of EEG GFP peaks as input across participants. Each colour indicates a microstate, cluster centre, with each cluster showing two points to the polarity invariance of the clustering method(Pascual-Marqui et al., 1995). (B) Cluster centres from EEG gradient space reconstructed as scalp topographies. Note that polarity is invariant.

Additionally, past studies which have applied gradient analysis to BOLD signal have done so using diffusion embedding (Margulies et al., 2016) rather than PCA, which was applied here. Diffusion embedding is known to project long distance connections more effectively projected into a common space more effectively than linear dimensionality reductions such as PCA (Coifman et al., 2005). Future investigations of the EEG gradient space will utilise diffusion embedding over PCA for this reason.

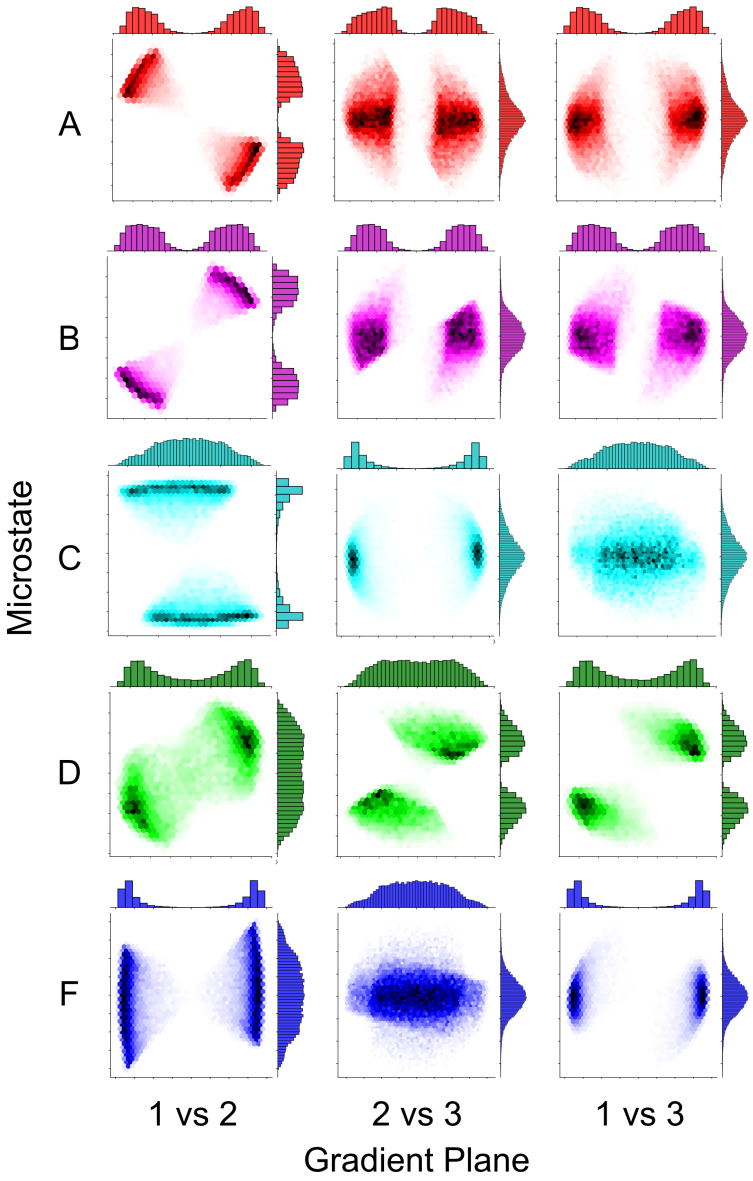

Figure 42: Planes of kernel density estimate of the originally derived microstate labels placed in the EEG gradient space. Each row denotes a microstate, also coloured accordingly, each column shows a single plane of the three dimensional space. X and Y axes of each column are gradient 1 to 2, 2 to 3 and 1 to 3 respectively. Histograms show the density distribution across the given axis, and darker colours within the plot indicate a higher density. All axes are between -1 and 1.

7.4.2 EEG Gradient Space

Placing the time series of each individual into the gradient space demonstrated a limitation and preference for the locations of coordinates within the space. Firstly, all participant time series' occupied an ellipsoid within the space. Whilst further investigation is required to understand this phenomena, most likely the outer boundaries of the shape are simply likely to the correlation calculation. For example, it is simply not possible for an EEG topography to have a correlation approaching 1 for both the first and second gradients.

The density distribution within the ellipsoid however requires further examination. There is a higher density of coordinates on the surface of the ellipsoid than within the body of it. Whilst it may be the case that this is due to the correlation with each gradient, there are still occurrences within the body of the ellipsoid, unlike the complete lack of density outside it. Further investigation into this phenomena is required.

Here, the density was much higher than average at either end of the gradient 1 axis. This was also the case for gradient 2 to a weaker extent. This indicates a preference for these topographies as well as combinations of the two in the recorded participants. Gradient 3 showed a relatively normal distribution with of density with a mean of zero, across participants. This is to be expected due to the low explained variance of gradient 3 relative to the other axes (Figure 36).

7.4.3 Prediction of fMRI CAPs

Application of LSTM models that use EEG gradient space trajectories to predict fMRI CAPs were not successful. Although model accuracy on the training data was high (between 70% and 90% across participants for both whole time series and GFP peak models), this accuracy did not generalise to the validation or test data, indicating over-fitting.

This over-fitting occurred even with implementation of regularisation on the LSTM layer. It may be the case that the model implemented was too simple, and therefore was incapable of capturing patterns of activity between features and targets.

Additionally, it may also be the case that the complexity of EEG gradient trajectories cannot be binned into six distinct categories that represent global fMRI activity patterns, and that associations are more particular. For this reason, future work will apply the same approach as is used here, only using a regressor with a target of the fMRI gradient space (as in Chapter 6, rather than a classifier which targets fMRI CAP states.

7.4.4 Microstates in EEG Gradient Space

The coordinates of GFP peaks in the gradient space were subject to the same modified k-means clustering (Pascual-Marqui et al., 1995) that is used to derive microstates using whole topographies (Figure 42). Interestingly, the cluster centres that were derived in the gradient space are not the same as those derived

in the topographical space. The derived clusters do reflected previously observed meta-microstates however (Koenig et al., 2023). This phenomena is most likely due to the lack of gradient 3 activity within the space, which as is seen on Figure 42, is where microstate D occurs.

The coordinates of GFP peaks were coloured using the microstate labels they were assigned in Chapter 6 (Figure 42). These distributions will be used in future studies to derive the microstate sequence which leads to the prediction of fMRI using the LSTM model. The microstates are likely attractor points in this space (Milz et al., 2017). Extraction of the microstate sequence post hoc allows for the use of complex models such as LSTMs that capture long-range dependencies, something that is more difficult to do when investigating sequences of symbols. Output attention weights can be used to compute heat maps per TR to identify preferences for each fMRI CAP label.

This preliminary investigation into the use of microstates as a post hoc analysis therefore proposes that further use of this approach is encouraged. The liberation of EEG microstate analysis from the investigation of microstate parameters and transition dynamics allows for a more robust understanding of microstates in the context of the EEG signal in general, and may lead to a better understanding of microstate syntax in the long run.

8 General Discussion

Multiple methodologies were proposed across three studies. Preprocessing and analysis principles were proposed in the literature review as necessary for the investigation of simultaneous Electroencephalography (EEG)-Functional Magnetic Resonance Imaging (fMRI) and EEG microstate syntax (Chapters 2 and 3). Whilst a consistent relationship between microstate sequences and fMRI signal could not be drawn across participants from the presented studies, each study proposes a novel methodological development based on the lessons learned from the previous one. Preliminary results in each instance point to some association between EEG microstate sequences and fMRI Blood-Oxygenation Level Dependency (BOLD) signal. Each has the aim of retaining the EEG microstate syntax and not reducing its dimensionality to fit the fMRI temporal resolution. Such attempts have not yet been made in the literature, other than some recent attempts that do not cover all bases (Abreu et al., 2021; Artoni et al., 2022; von Wegner et al., 2017). The following Chapter will review the methodologies proposed, relate findings across studies to the literature, and will suggest future directions.

8.1 Review of Proposed Methodologies

8.1.1 Pre-processing Considerations for EEG Microstate Syntax

The pre-processing methods proposed herein may provide a more robust means of investigating microstate syntax. The "interpolation method" proposed as a microstate derivation method in Chapter 3, along with the avoidance of the smoothing method has already been highlighted by others as ill advised due to its tendency to destroy the syntactic structure of microstate sequences (von Wegner et al., 2017). von Wegner et al. (2017) did apply the smoothing method however, only without the removal of microstates that were deemed too short. The peak-mode method of derivation applied in Chapter 6 takes the proposal of avoiding smoothing a step further by simply taking the centre point between Global Field Power (GFP) peaks as a transition between microstates, using the GFP peaks from which the microstates are derived as the focus of syntax analysis. The result was a set of parameters that more accurately described the patterns of activity between GFP peaks, the source of microstate derivation (Lehmann et al., 1987).

Minimisation or complete avoidance of epoch cutting of the EEG time series is also used here to avoid both the misalignment of EEG and fMRI and the damaging of syntactic structure. Conventional methods either down-sample the temporal resolution of EEG to that of fMRI such as in the voxel-wise General Linear Model (GLM) approach (Britz et al., 2010), or apply epoching. Whilst epoching does improve Signal-to-Noise Ratio (SNR)s, cutting periods of the time series from analysis may weight the calculation of n-gram parameters, or cause difficulties in the alignment process between EEG and fMRI data. The methods used in each study here would not be possible with cutting, as the amount of EEG data per fMRI state or Repetition Time (TR) would vary, likely skewing sample sizes in favour of particular Co-Activation Pattern (CAP)s (Chapters 5 and 7) or fMRI gradient coordinates (Chapter 6).

It is suggested that future work minimises the cutting of noise in the time series, instead opting for cleaning algorithms. It is also suggested that the interpolation method of microstate derivation is to be utilised on the time series when investigating microstate syntax to avoid loss of information.

8.1.2 Considerations for Simultaneous EEG-fMRI Analysis

As investigations of the relationship between EEG and fMRI become more concerned with dynamic, global states (Abreu et al., 2021), it is important to consider whether the analysis methods used are chosen with those dynamics in mind. In general, methods that draw relationships between EEG and fMRI data use GLMs (Britz et al., 2010; Xu et al., 2020; Yuan et al., 2012). Whilst these models do provide a general overview of EEG signal that correlate with the BOLD signal at each voxel, the temporal dimensionality reduction of the EEG signal completely destroys its dynamics, whether investigation is of microstates or EEG signal in general (Chapter 2).

Each study presented here demonstrates the complexity of the EEG signal, and each proposes a potential modelling method that could circumvent the issue of temporal dimensonality reduction of the EEG time series. Particularly the EEG gradient space, which conceptualises the EEG signal as a trajectory in a multidimensional space (Chapter 7). The whole of the trajectory within a window that corresponds to a simultaneously recorded BOLD image allows for potential similarities in trajectories to be drawn across TRs. This direct symmetrical comparison avoids any destruction of the EEG time series, retaining its complexity.

8.1.3 Sparse Distance Graphs

Study 1 (Chapter 5) used simultaneously occurring EEG microstates and fMRI CAPs. The methodology attempted to measure the L_1 distance between distributions of EEG microstate n-gram parameters during each of the fMRI CAPs. The distances were then used as edges to connect CAPs as nodes, and edges over a given threshold were retained in sparse graphs. The number of remaining edges from the observed CAPs' sequences were then compared against the number of remaining edges in the CAPs' sequences with shuffled labels.

This methodology allowed for direct comparison between the fMRI classes, which simplified the fMRI analysis. Although there was not a consistent result found across participants, the method showed statistically significant results at the subject-level, even after Benjamini-Hochberg multiple comparison correction (Benjamini & Yekutieli, 2005). The method shows that there is a clear difference between each n-grams parameters during simultaneously occurring CAPs, and that the observed differences between CAPs are not due to random chance.

The shortcomings of this approach however are, firstly, in the derivation of EEG microstate n-gram parameters during each fMRI CAP. Since the coverage of each CAP can vary, there is an immediate possibility for a difference in 'sample sizes' between the CAPs, which can skew the estimates of the EEG microstate n-grams' 'preferences' across the CAPs. Additionally, since comparisons have to be made between CAPs at each n-gram length individually, it is difficult to show comparisons across n-grams of lengths. It is not the case that n-grams of different lengths are independent of one another, but the process of conducting tests on each n individually assumes that they are. Applying multiple comparison correction across lengths was used as a potential control for this shortcoming, but the method itself does not allow for comparisons across n's, which is more desirable due to being true to the actual nature of neural dynamics that do not follow a fixed length of sequence iterations.

Finally, the use of the degree centrality of each node in the sparsely connected graphs may be developed in future investigations of differences between fMRI states. Future iterations of this model should utilise other measures to compare nodes or fMRI classes (CAPs) such as betweenness centrality and clustering coefficients to get a better understanding of how microstate n-gram parameters vary across fMRI dynamic Functional Connectivity (dFC) states (Dennis et al., 2012).

8.1.4 fMRI Gradient Coordinates and Ensemble Regressors

The regressor models used to identify a relationship between the change in microstate n-gram parameters and fMRI gradient space coordinates (Chapter 6) could fit the data at $R^2 = 0.3$ when considering each axis individually.The value of 0.3 indicates a weak fit for the data, but in drawing relationships between EEG and fMRI BOLD signal, it is not uncommon to report such relationships as meaningful (e.g. (Yuan et al., 2012)). Both EEG and fMRI data are subject to noise components of many types (see Chapter 2) which makes drawing any relationships difficult. Furthermore, a null R^2 should be calculated in future studies to confirm the predictive power of these models.

However, better fitting models would of course be desirable. The use of a single model per gradient axis, as well as per parameter and per participant, increases the number of models to a point of concern. Use of parameters of each individual n-gram, rather than the whole of a microstate sequence (as in Study 3) causes this combinatorial increase in the number of models.

If one was concerned with the investigation of n-gram parameters specifically, an ideal model would consider the parameters of microstate n-grams at multiple lengths at once. It is clear syntactically that relationships between microstates within a sequence could occur between lengths. An immediate shortcoming of this approach (as well as the approach used in Study 1) is the inability to consider multiple lengths simultaneously. It may be possible to use a regression model to do this, but the number of features in such a model would be very large, requiring an extremely large dataset, but even then, the regression approach will be compounded by collinearity across the predictors.

8.1.5 Using the EEG Gradient Space to understand EEG Microstate Syntax

It was proposed that the first target of microstate syntax analysis should be an understanding of EEG dynamics in general that should then be used to understand microstate sequences post hoc (Chapter 7). This approach has the benefit of allowing for the utilisation of models such as the Recurrent Neural Network (RNN) used here, which use a continuous dynamic signal, rather than the limitations imposed by a sequence of class labels. Whilst this proposed technique may allow for a robust understanding of the relationship between EEG microstate sequences and fMRI BOLD signal, the model architectures used here require further investigation in order to achieve generalisable models.

Across all studies developed here, an important limitation to consider is the issue of EEG-fMRI alignment. Whilst these approaches do indeed allow for a retention of EEG microstate sequencing, and therefore for a direct association between EEG microstate syntax and fMRI, in each case the findings depend upon the offset of timing between the electrical signal observed at the scalp, and the BOLD signal observed in the brain. This is accounted for by the Haemodynamic Response Function (HRF), but in each study here, a simple single peak, gamma response function was used as a reference, with the time

to peak at 6 seconds. It is known that different regions of the brain have different response functions, and can differ between subjects (Handwerker et al., 2004). In the methods of association proposed between EEG and fMRI here, the choice made for offset time to peak affects all results. Alignment is a critical issue when considering these approaches, and therefore caution should be taken when applying these methods without first considering the HRF.

8.2 Relevance to the Field

8.2.1 Derived EEG Microstates

The microstates identified here in resting state recordings exhibited similar topographies to those identified in the field. The canonical set of four (Koenig et al., 2002; Lehmann et al., 1987) as well as the less common non-canonical meta-microstate F (Koenig et al., 2023; Tarailis et al., 2023) were all identified. All results reported here therefore can be used as contributions to the field and may aid in better understanding the commonly identified microstates.

The non-canonical microstate F has previously been identified in few studies (Tarailis et al., 2023) and has been associated with the anterior Default Mode Network (DMN) (Custo et al., 2017) and the medial prefrontal cortex (Bréchet et al., 2019). Here, it was found using the generally adopted GLM method that microstate F correlated with activity across the group in the right temporal occipital fusiform cortex, as well as the left occipital pole (Figure 13). These findings have not been described previously, and further investigation is needed.

The microstate parameters reported may be used as an addition to the field when investigating the resting state, eyes-open. Additionally, report of EEG microstate parameters in the conventional event-mode (Figure 12), along with the same parameters reported in peak-mode (Figure 28), allow for a comparison between derivation modes.

The proposal of frequency as a new parameter, as well as the reporting of n-gram parameters at lengths greater than 1, are also novel metrics that may be useful if adopted by the field in future investigations of n-grams.

8.2.2 Development of Microstate and EEG-fMRI Association Methodologies

It was highlighted in the literature review that current investigations of microstates do not sufficiently consider the temporal relationship between EEG and fMRI (Chapter 2) or the impact of microstate syntax (Chapter 3) sufficiently.

With regards to microstate syntax, a distinct lack of consideration has been made in the literature when defining microstate sequences, defining parameters for sub-sequences, or indeed updating existing preprocessing methods to account for syntax.

Sequence modes have not been defined previously, meaning that comparison between studies regarding microstate syntax is difficult without careful scrutiny

of each individual studies methodologies. Additionally, the continuous nature of the EEG time series is also not considered when investigating syntax. Using the EEG gradient space as a continuous signal to extract the microstate sequence from post-hoc, is a novel means of syntax investigation which does not reduce the EEG signal down to a discrete sequence of classes immediately.

With regards to the relationship between EEG and fMRI signals, the standard practice is the use of a GLM, with signal from the EEG data used as a regressor for voxel-wise attempts at association with fMRI BOLD signal. Whilst this approach does show a general relationship between the two recordings, there are shortcomings in the approach. The down-sampling of EEG to the same sample frequency as the fMRI destroys the dynamic structure of EEG. It has been shown that the EEG microstate time series has a short term memory, in that information is retained in the time series in a non-Markovian manner (von Wegner et al., 2018). Destroying the syntax of microstate sequencing makes associating the patterns of EEG activity with fMRI signal impossible.

Furthermore, studies which associate the GLM output of voxel-wise statistical maps to fMRI signals do so through association with static Functional Connectivity (FC) networks. These networks are parcellations (e.g., Figure 4), which do not allow a network or region to co-activate with more than one set of regions. This means that more complex connectivity between brain regions cannot be associated to the EEG signal in these models.

In the studies laid out here, multiple attempts were made to alleviate these shortcomings. In the first study, fMRI CAPs were used as global activation patterns over static FC states.

8.2.3 Findings

Study 1 (Chapter 5) firstly demonstrated the standard set of results that are laid out in traditional investigations of the relationship between EEG microstates and simultaneously recorded fMRI BOLD signal. The reports of microstate parameters, as well as GLM outputs contribute to the findings reported in other studies. Most notably is the fMRI activity pattern associated with microstate F, which is underrepresented in microstate literature (Tarailis et al., 2023).

Furthermore, the statistically significant differences identified between microstate n-gram parameters during different simultaneously occurring fMRI global activation patterns such as CAPs, validates the use of dynamic states as proposed in Chapter 2, as more recent studies have also done (Abreu et al., 2021). The use of these global patterns shows that EEG microstates need not be associated with spatial parcellations of fMRI BOLD signal, as is common practice (Britz et al., 2010; Michel & Koenig, 2018; Musso et al., 2010; Xu et al., 2020; Yuan et al., 2012).

Whilst the associations on the subject-level did not necessarily generalise to the group-level, these initial results validate an association between global fMRI activation patterns and EEG microstate syntax. The lack of consistency between participants should inform future methodological designs which wish to investigate microstate n-grams parameters specifically, and microstate syntax

more in generally.

Moving to Study 2 (Chapter 6), the use of an fMRI gradient space is a novel approach (Brown et al., 2022), not previously used to associate fMRI and EEG signals. The method reduces the spatial dimensionality of the fMRI data whilst still retaining its functional significance.

The demonstration of a uniform density of EEG microstates within the fMRI gradient space is a helpful visualisation which demonstrates the uniformity of their occurrence throughout the time series, and highlights the need for more sophisticated relationships to be drawn between EEG microstates and fMRI BOLD signal in general. It may be the case therefore that the underlying structure of EEG sequences may be what drives associations between EEG microstates and fMRI BOLD signal, rather than one-to-one associations assumed by previously used methods.

Furthermore, the results of Study 2 demonstrated peak-mode (and the preprocessing pipeline used to derive them) as a valid means of EEG microstate derivation. The use of such a method (as is discussed in Chapter 3) may be a more appropriate means of investigating microstate n-grams, and microstate syntax in general, due to the isolation of analysis to GFP peaks, as well as the possibility for transition from a microstate class to another occurrence of the same microstate class. The parameter of duration in the context of peak-mode also indicates the amount of time between consecutive GFP peaks, something that is not accounted for in the conventional event-mode (Michel & Koenig, 2018).

The random forest regressor models used to predict fMRI gradient space coordinates using co-occurring microstate n-gram parameters demonstrated further connections between EEG syntax and fMRI activity. Whilst overall relationships at the group-level could not be established, the predictive power of all models with an R^2 value greater than zero indicates an interaction between the two measures. Further investigation into these results is needed to confirm whether associations are due to chance or whether they are participant specific. Future work should measure the R^2 values with shuffled n-gram labels to establish a baseline R^2 for each n-gram to compare with the observed value.

Finally, the findings of Study 3 (Chapter 7) were derived using an EEG gradient space as a novel means of observing EEG activity. Investigation of the density of EEG activity within this space shows the utility of EEG microstates, allowing them to be conceptualised as attractor points within a multidimensional space, as has been suggested in the past (Milz et al., 2017). Investigating GFP peaks and labelling each peak with its corresponding microstate class within the space allows for a continuous representation of EEG activity which can be subject to a wider range of analyses methods, yet still be investigated as a microstate sequence post hoc.

Although preliminary results failed to identify a classifier which could correctly identify EEG gradient space trajectories as a given fMRI CAP between both training and test partitions, further investigation of the method using other RNN architectures may lead to identification of a better fitting model. Additionally, a regressor may be utilised rather than a classifier in the future,

with the target being an fMRI gradient space coordinate rather than a CAP label. These methods developments are valuable to the field of microstate syntax investigation, and to attempts of simultaneous EEG-fMRI analysis in general.

8.2.4 Usefulness of EEG Microstates and n-Grams

Previous studies have demonstrated that microstates and microstate n-grams are valuable biomarkers for identifying neurological disorders and diseases (Lehmann et al., 2005; Nishida et al., 2013; Serrano et al., 2018; Vellante et al., 2020). Investigations of EEG microstate sequences have demonstrated a complexity that is represented as a short-term memory in the sequence (von Wegner et al., 2017) that is yet to be understood.

The methods used in Chapters 5 and 6 use microstate n-gram parameters as tools of their analysis. It is worth highlighting that such parameters do not consider the individual microstate parts. The measurements for ABC for example, do not consider the parameters of A, B, C, AB and BC. Whilst parameters were collected for each of these 1- and 2-grams, a summarising measurement that takes all into account at once would be valuable for investigations of n-grams, and would not limit comparison of n-grams to their own length. There are clear dependencies between lengths that are yet to be addressed.

The EEG gradient method does appear to demonstrate a shortcoming of the microstate method in general however. In the gradient space, the moment-to-moment EEG signal is conceptualised simply as a trajectory moving through the multidimensional gradient space. Since microstates are derived via clustering, the microstates are a segmentation of this space where the label associated with each EEG data point indicates the region of the space that it is occupying. Studies that have investigated EEG microstate parameters such as duration, occurrence, and coverage have been attempting to understand the average behaviour of subsections of the overall trajectory. Mean duration for example is the average length of each subsection of the trajectory where the microstate label remains the same.

Whilst the microstate approach does uncover some general characteristics of the EEG signal with regards to its topography, the specifics are lost. There is possibility of investigating the dynamics of the signal beyond the labels associated. Two sub-trajectories that are labelled with microstate A for 100ms could both be moving in different directions in the space, and could be occupying different regions of "A space", but would be seen as identical when considered solely as microstates with parameters. This problem extends to the investigation of microstate n-grams, they are longer sub-trajectories that pass through the pre-defined microstate boundaries. The percentage occurrence of microstate n-grams may show how often these transitions across boundaries occur, but there is a lack of consideration for the differences in these sub-trajectories.

I believe that future investigations of microstates should keep these points in mind. Microstate labels simplify the trajectory through discretisation into a sequence of labels, and allow for methods that shed light on the quasi-stable topographies that appear around GFP peaks. But, the information about the

EEG signal that is lost when applying microstate analysis should not be forgotten, and methods should be developed that retain it.

8.3 Future Directions

8.3.1 Application of Methodologies to Cognitive Manipulation

The dataset used in this study was collected on 18 individuals (3 removed during cleaning). The dataset used was a simultaneous recording of EEG-fMRI in the resting state. This dataset itself includes some limitations that should be highlighted. Firstly, the resting state may not be the most appropriate cognitive state to investigate when attempting to identify significant connections between EEG and fMRI. Previous studies that have attempted to draw connections between the two imaging modalities with microstates in the resting state have done so with weak correlation values (Britz et al., 2010; Yuan et al., 2012). It perhaps may have been more suitable to investigate methods using tasks with well established activity patterns in the fMRI domain, and to work backwards from there, attempting to draw associations with each of the EEG microstates this way. An event-based task design may have worked in synergy with the temporally-based approaches developed here. Future studies should begin with applying the methodologies outlined in this thesis at both subject- and group-levels to the task data collected from the same individuals during a sustained attention (Choice Reaction Time (CRT)) task, Fagerholm et al. (2015)).

The simultaneous EEG-fMRI data were collected during two versions of the CRT task: one with blocks of rest and another with no rest. Participants were shown an arrow pointing either left or right. The task was to react with a button push with the left or right index finger as quickly as possible for the same side as the arrow was pointing. As can be seen from Figure 43, the EEG microstates derived from the resting state data used in Chapters 5 to 7 are very similar to the EEG microstates derived from the EEG data during performance of two CRT tasks.

The microstates are very similar between studies, though it appears that meta-microstate F may change to G during the task (see Figure 2, or Koenig et al. (2023)). The high similarity between the two cognitive states may allow for more robust modelling, as comparison could be used to control for noise.

It is also worth highlighting that the methods developed here; especially the EEG gradient space, may be valuable for investigation of EEG without association to fMRI as was the case with this dataset. The notion of conceptualising the EEG signal as a trajectory in a multi-dimensional space is a powerful one, and may be a valuable method in instances where dynamics of EEG are at the centre of investigation. Studies that are attempting to understand transition boundaries between sleep states for example (Benington et al., 1994), may benefit from such an approach. Future studies will focus on developing the EEG gradient space approach and applying it to such investigations.

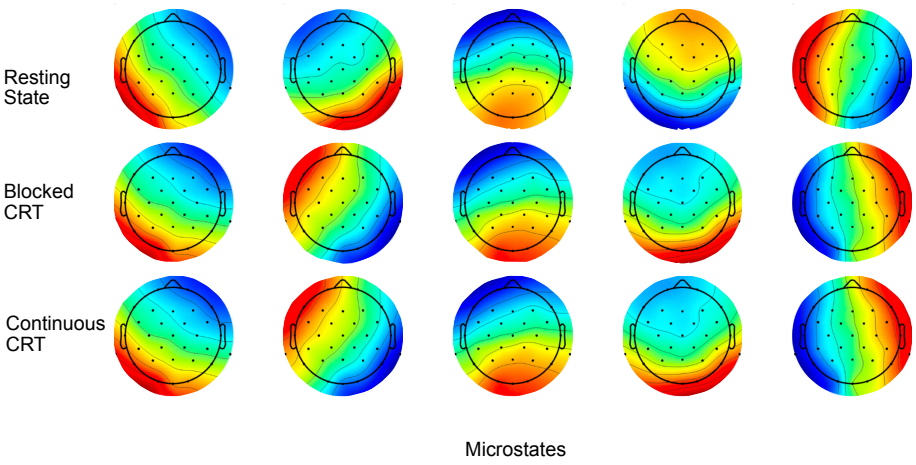

Figure 43: EEG microstates generated using 5 cluster centres for rest, block CRT and continuous CRT tasks, top to bottom. Microstates have not been labelled, but have been aligned based on similarity. Note that polarity colour is unimportant, only pole locations are considered in microstate analysis.

8.3.2 Future Methodological Developments

Future studies are likely to benefit from developing and refining the methodologies proposed here. For example, the Study 1 (Chapter 5) method which measured distances between microstate n-gram distributions may be developed upon by using other metrics than the number of edges over a given threshold.

The implementation of a gradient space methodology to observed EEG signal in future investigations may further the understanding of EEG microstate syntax by investigating EEG sequences post-hoc, following gradient space analysis. This method allows for free investigation of the EEG time series which may then be reduced to a discretised sequence of microstates post-hoc. The models used then do not need to be constrained to those that investigate transitions between microstates.

Future studies that wish to investigate the relationship between EEG microstates and the fMRI BOLD signal should in general consider the principles laid out in the literature reviews of Chapters 2 and 3. New methods aimed at investigating EEG microstate syntax must ensure that syntactic structure is retained at all phases. The investigations relating EEG and fMRI BOLD signal dynamics should aim to retain the resolution of recorded EEG, finding alternative approaches to overcoming the issue of temporal resolution of the EEG and fMRI data.

8.3.3 Epsilon-Machines

One of the methods proposed for the study of complex microstate syntax is using epsilon machines (Nehaniv & Antonova, 2017). Epsilon machines are dis-

crete dynamical system automata with state-dependent probabilities on different future observations (Crutchfield, 1994; Rhodes & Nehaniv, 2009). Syntactic structure is well studied in formal language theory, where different classes of automata models act as recognisers of a "language" of possible sequences (Hopcroft et al., 2001).

Epsilon machines model the syntactic structure of the observed sequence of states. Unlike fitting parameters to the architecture of a pre-given model, the structure of an epsilon machine is derived from the data, which is first built naively and is then minimised into a more fundamental model. The epsilon machine iterates through each observation in the sequence and collects the transitions from each state to the next. After iterating through the whole sequence, the epsilon machine has a count of transitions from each state to every other state that occurred in the dataset, generating a probabilistic distribution. This distribution predicts the next observation given the current observation. Figure 44 shows a visualisation of the process.

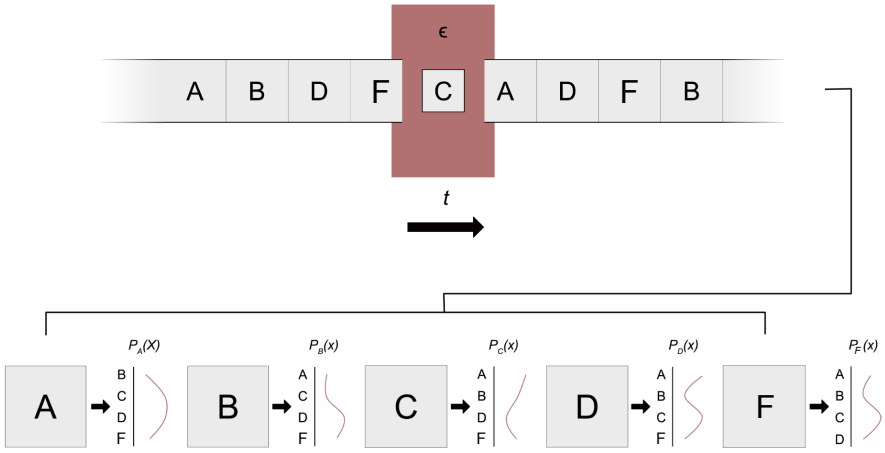

Figure 44: Visualisation of the epsilon machine. An epsilon machine iterates through each state in the input sequence (top) and builds a probability distribution of transition from each state to every other possible state (bottom). This results in a matrix of transition probabilities which can be used to generate a probabilistic automaton. This process can be done with n-grams of longer lengths (i.e., bottom microstates shown here would be each n-gram with its own distribution of transitions to the next possible state).

A benefit of the epsilon machine approach is the ability to use short length sequences as input, rather than a single event as observation. In application to microstates, n-grams can be used as states. For example, a 4-gram can be used as an input observation, not just a single observation of microstate A.

Using an epsilon machine which considers n-grams as input events instead of

single letters, causes a calculation of transition probabilities from each length n sequence to the set of next possible length n sequences. For example, $ABCD$ has the possibility to transition to $BCDA$, $BCDB$ and $BCDC$ if using the canonical EEG microstate set in event mode (see Chapter 3 for explanation). Using length 4-grams as an example, the generated automaton would be a much more complex web of connections than the simple transitions between microstates, where each state would be a 4-gram, rather than just a single letter. Following the building of the 4-gram epsilon machine, each n-gram is an event that has an associated transition probability distribution to all other possible states.

Once this basic structure has been derived, the epsilon machine can be "minimised". The minimisation process involves identifying the *causal* structure of the sequence. If two 4-grams have very similar probabilistic distributions to the set of possible subsequent observations, they are binned into the same "causal state". In this context, binning two 4-grams into the same causal state identifies them as having the same predictive value when considering the following observation.

Figure 45 shows the 1-gram epsilon machine generated from the observed event-mode microstate sequences of the participants.

Figure 46 then shows a schematic of a 4-gram event-mode epsilon machine after minimisation. Note how the minimisation process bins multiple n-grams into a single causal state due to their similar probabilities of transition to the next microstate.

Using a model such as this allows for a description of the EEG syntax that is non-Markovian, something that has already been observed in past studies (von Wegner et al., 2018). Preliminary results have observed some minimisation of n-grams into causal states, suggesting that these n-grams have a similar predictive power in predicting the next state.

Future work should utilise epsilon machines for the investigation of the relationship between EEG microstates and fMRI CAPs (or fMRI classes in general) in simultaneously recorded EEG/fMRI data. The microstate sub-sequences that are occurring during each fMRI state class can be isolated and used to build epsilon machines. Figure 47 shows a representation of such a model.

The epsilon machines generated at the microstate level are for individual fMRI CAP states. A machine can also be generated at the CAP level using the sequence of CAPs as input. Hence a "multi-level" epsilon machine which simultaneously describes the transitions between microstate n-grams and fMRI CAPs can be generated, and subsequently interrogated.

This is a purely discrete model of the time series which only considers transitions between states. Utilising such a model in-tandem with a continuous model such as the gradient space regression models proposed above would allow for a double-pronged approach to the investigation of microstate syntax which avoids the pitfalls of a purely continuous, or purely discrete investigation.

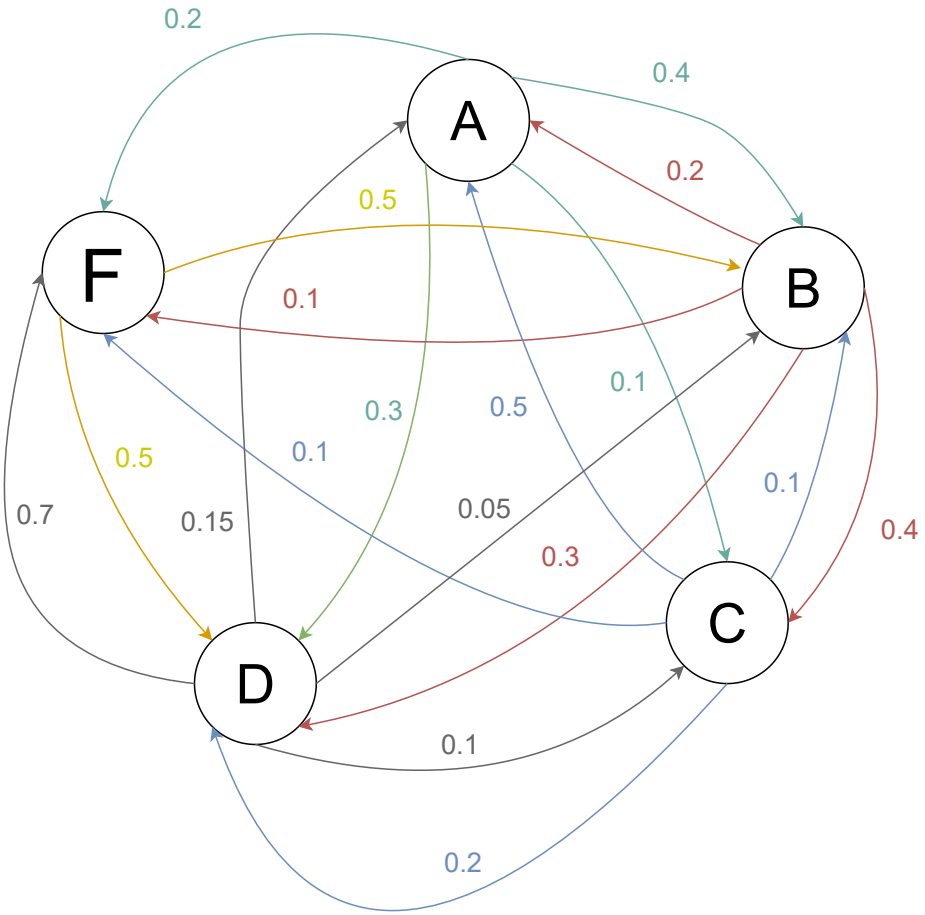

Figure 45: Group 1-gram event-mode epsilon machine. Nodes in the graph are microstates, directed edges are transition ratios. Process of minimisation is not possible at the 1-gram level, so causal states are equivalent to microstate classes.

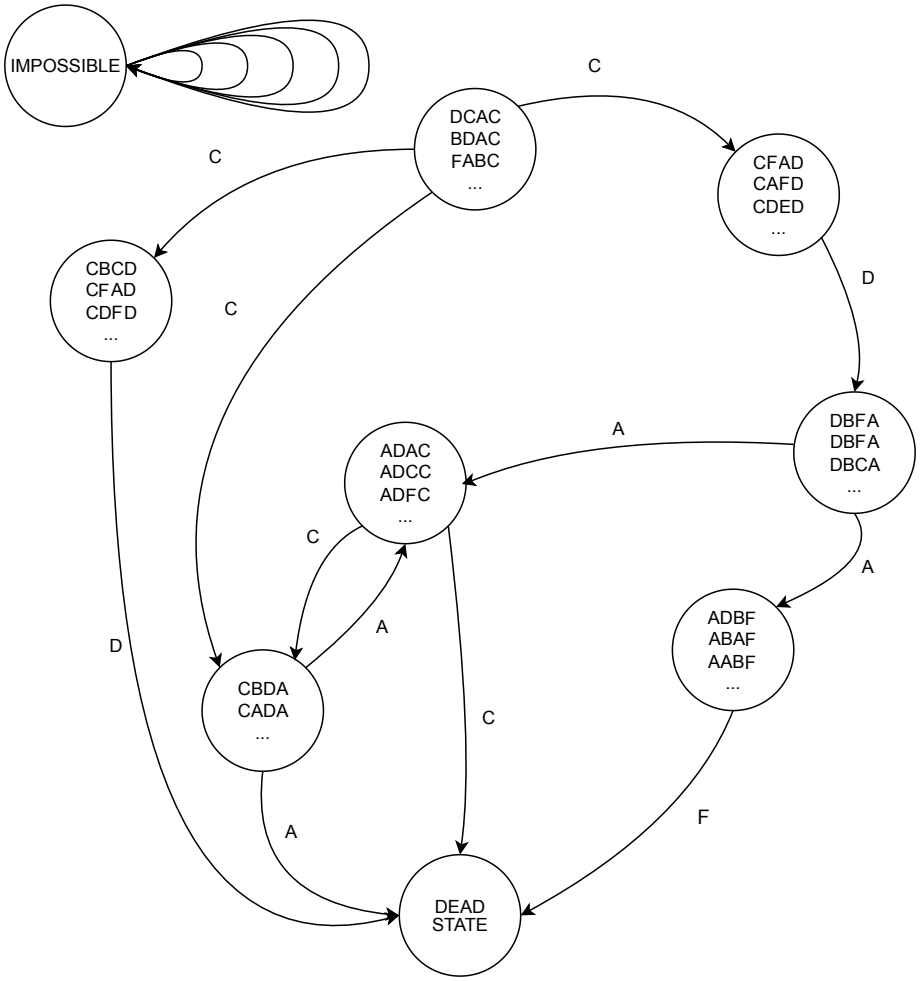

Figure 46: Visualisation of a minimised event-mode 4-gram epsilon machine, not generated from data. The 4-grams with similar probability distributions are binned into the same causal state, meaning it is assumed that the underlying causal state is the driver of retained information through the syntax, rather than the n-grams themselves. The impossible state includes the list of all possible 4-grams that did not occur in the input sequence, and the dead state is the state transitioned into at the end of a sequence (i.e., at the end of a participants observed sequence).

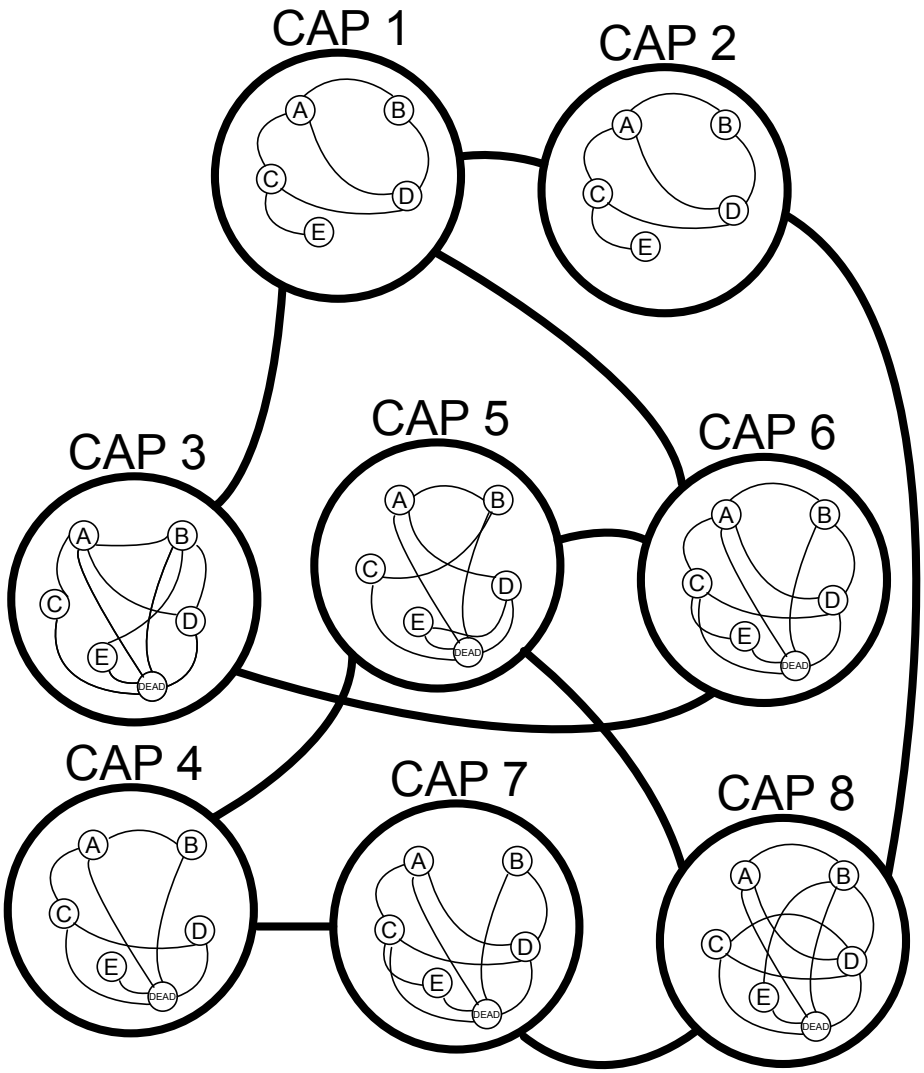

Figure 47: Concept for a "multi-level" epsilon machine. The EEG microstates which are occurring simultaneously during a given fMRI CAP are isolated and an epsilon machine is built using only those n-grams, for each CAP. An epsilon machine is then generated with the fMRI CAP sequence as input to get an fMRI-level epsilon machine. The result would be an overarching CAP-level epsilon machine which transitions between nested microstate-level epsilon machines. In this example, the microstates generated are A to E, and there are eight CAPs.

9 General Conclusion

It was demonstrated here, using a number of novel methodologies, firstly that there is a marked difference between Electroencephalography (EEG) microstate n-gram parameters during simultaneously occurring global Functional Magnetic Resonance Imaging (fMRI) Co-Activation Pattern (CAP) patterns, and that these associations could not be established across participants. This was extended to the investigation of fMRI gradient space, where preliminary results suggested that the relationship between EEG microstate n-gram parameters and fMRI Blood-Oxygenation Level Dependency (BOLD) signal could be established as a dynamic signal rather than as differences between classes of activity patterns. Finally, the EEG gradient space was derived as an alternative to microstate sequences and n-grams, as a potential means of liberating microstate analysis from the investigation of parameters and transition probabilities. Results were generated using the preprocessing pipeline proposed here that was developed specifically with EEG microstate syntax in mind, and also develop upon existing methods that investigate the relationship between simultaneously recorded EEG-fMRI data, prioritising the retaining of EEG temporal resolution. It is suggested that future investigations into the relationship between EEG microstates and fMRI BOLD signal utilise these methods and develop them further to elucidate the functional significance of EEG microstates.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., … Zheng, X. (n.d.). TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems.

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., & Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, *8*.

Abreu, R., Jorge, J., Leal, A., Koenig, T., & Figueiredo, P. (2021). EEG Microstates Predict Concurrent fMRI Dynamic Functional Connectivity States. *Brain Topography*, *34*(1), 41–55.

Allen, P. J., Josephs, O., & Turner, R. (2000). A Method for Removing Imaging Artifact from Continuous EEG Recorded during Functional MRI. *NeuroImage*, *12*(2), 230–239.

Amico, E., Gomez, F., Perri, C. D., Vanhaudenhuyse, A., Lesenfants, D., Boveroux, P., Bonhomme, V., Brichant, J.-F., Marinazzo, D., & Laureys, S. (2014). Posterior Cingulate Cortex-Related Co-Activation Patterns: A Resting State fMRI Study in Propofol-Induced Loss of Consciousness. *PLOS ONE*, *9*(6), e100012.

Antonova, E., Holding, M., Suen, H. C., Sumich, A., Maex, R., & Nehaniv, C. (2022). EEG microstates: Functional significance and short-term test-retest reliability. *Neuroimage: Reports*, *2*(2), 100089.

Artoni, F., Maillard, J., Britz, J., Seeber, M., Lysakowski, C., Bréchet, L., Tramèr, M. R., & Michel, C. M. (2022). EEG microstate dynamics indicate a U-shaped path to propofol-induced loss of consciousness. *NeuroImage*, *256*, 119156.

Bakhshayesh, H., Fitzgibbon, S. P., Janani, A. S., Grummett, T. S., & Pope, K. J. (2019). Detecting connectivity in EEG: A comparative study of data-driven effective connectivity measures. *Computers in Biology and Medicine*, *111*, 103329.

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1457), 1001–1013.

Benington, J. H., Kodali, S. K., & Heller, H. C. (1994). Scoring Transitions to REM Sleep in Rats Based on the EEG Phenomena of Pre-REM Sleep: An Improved Analysis of Sleep Structure. *Sleep*, *17*(1), 28–36.

Benjamini, Y., & Yekutieli, D. (2005). False Discovery Rate–Adjusted Multiple Confidence Intervals for Selected Parameters. *Journal of the American Statistical Association*, *100*(469), 71–81.

Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.-M., & Robbins, K. A. (2015). The PREP pipeline: Standardized preprocessing for large-scale EEG analysis. *Frontiers in Neuroinformatics*, *9*.

Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*(4), 537–541.

Braga, R. M., & Leech, R. (2015). Echoes of the Brain: Local-Scale Representation of Whole-Brain Functional Networks within Transmodal Cortex. *The Neuroscientist*, *21*(5), 540–551.

Bréchet, L., Brunet, D., Birot, G., Gruetter, R., Michel, C. M., & Jorge, J. (2019). Capturing the spatiotemporal dynamics of self-generated, task-initiated thoughts with EEG and fMRI. *NeuroImage*, *194*, 82–92.

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32.

Britz, J., Van De Ville, D., & Michel, C. M. (2010). BOLD correlates of EEG topography reveal rapid resting-state network dynamics. *NeuroImage*, *52*(4), 1162–1170.

Brodbeck, V., Kuhn, A., von Wegner, F., Morzelewski, A., Tagliazucchi, E., Borisov, S., Michel, C. M., & Laufs, H. (2012). EEG microstates of wakefulness and NREM sleep. *NeuroImage*, *62*(3), 2129–2139.

Brown, J. A., Lee, A. J., Pasquini, L., & Seeley, W. W. (2022). A dynamic gradient architecture generates brain activity states. *NeuroImage*, *261*, 119526.

Buckner, R. L. (1998). Event-related fMRI and the hemodynamic response. *Human Brain Mapping*, *6*(5-6), 373–377.

Burt, J. B., Demirtaş, M., Eckner, W. J., Navejar, N. M., Ji, J. L., Martin, W. J., Bernacchia, A., Anticevic, A., & Murray, J. D. (2018). Hierarchy of transcriptomic specialization across human cortex captured by structural neuroimaging topography. *Nature Neuroscience*, *21*(9), 1251–1259.

Calhoun, V. D., & Adali, T. (2016). Time-Varying Brain Connectivity in fMRI Data: Whole-brain data-driven approaches for capturing and characterizing dynamic states. *IEEE Signal Processing Magazine*, *33*(3), 52–66.

Case, M., Zhang, H., Mundahl, J., Datta, Y., Nelson, S., Gupta, K., & He, B. (2017). Characterization of functional brain activity and connectivity using EEG and fMRI in patients with sickle cell disease. *NeuroImage: Clinical*, *14*, 1–17.

Chen, J. E., Chang, C., Greicius, M. D., & Glover, G. H. (2015). Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. *NeuroImage*, *111*, 476–488.

Chu, C., Wang, X., Cai, L., Zhang, L., Wang, J., Liu, C., & Zhu, X. (2020). Spatiotemporal EEG microstate analysis in drug-free patients with Parkinson's disease. *NeuroImage: Clinical*, *25*, 102132.

Coifman, R. R., Lafon, S., Lee, A. B., Maggioni, M., Nadler, B., Warner, F., & Zucker, S. W. (2005). Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences*, *102*(21), 7426–7431.

Croce, P., Quercia, A., Costa, S., & Zappasodi, F. (2020). EEG microstates associated with intra- and inter-subject alpha variability. *Scientific Reports*, *10*(1), 2469.

Croft, R., & Barry, R. (2000). Removal of ocular artifact from the EEG: A review. *Neurophysiologie Clinique/Clinical Neurophysiology*, *30*(1), 5–19.

Crutchfield, J. P. (1994). Observing Complexity and the Complexity of Observation. In H. Haken, H. Atmanspacher, & G. J. Dalenoort (Eds.), *Inside Versus Outside* (pp. 235–272, Vol. 63). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-48647-0_14

Custo, A., Van De Ville, D., Wells, W. M., Tomescu, M. I., Brunet, D., & Michel, C. M. (2017). Electroencephalographic Resting-State Networks: Source Localization of Microstates. *Brain Connectivity*, *7*(10), 671–682.

Damoiseaux, J. S., Rombouts, S. A. R. B., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M., & Beckermann, C. F. (2006). Consistent resting-state networks across healthy subjects. *PNAS*, *103*(37), 6.

D'Croz-Baron, D. F., Bréchet, L., Baker, M., & Karp, T. (2021). Auditory and Visual Tasks Influence the Temporal Dynamics of EEG Microstates During Post-encoding Rest. *Brain Topography*, *34*(1), 19–28.

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.

Dennis, E. L., Jahanshad, N., Toga, A. W., McMahon, K. L., de Zubicaray, G. I., Martin, N. G., Wright, M. J., & Thompson, P. M. (2012). Test-Retest Reliability of Graph Theory Measures of Structural Brain Connectivity. In N. Ayache, H. Delingette, P. Golland, & K. Mori (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012* (pp. 305–312). Springer. https://doi.org/10.1007/978-3-642-33454-2_38

Diezig, S., Denzer, S., Achermann, P., Mast, F. W., & Koenig, T. (2022). EEG Microstate Dynamics Associated with Dream-Like Experiences During the Transition to Sleep. *Brain Topography*.

Faber, P. L., Lehmann, D., Barendregt, H., Kaelin, M., & Gianotti, L. R. R. (2005). Increased duration of EEG microstates during meditation. The KEY-Institute for Brain-Mind Research.

Fagerholm, E. D., Lorenz, R., Scott, G., Dinov, M., Hellyer, P. J., Mirzaei, N., Leeson, C., Carmichael, D. W., Sharp, D. J., Shew, W. L., & Leech, R. (2015). Cascades and Cognitive State: Focused Attention Incurs Subcritical Dynamics. *Journal of Neuroscience*, *35*(11), 4626–4634.

Férat, V., Arns, M., Deiber, M.-P., Hasler, R., Perroud, N., Michel, C. M., & Ros, T. (2021). Electroencephalographic Microstates as Novel Functional Biomarkers for Adult Attention-Deficit/Hyperactivity Disorder. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*.

Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences*, *103*(26), 10046–10051.

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., Hampton, J. M., Coalson, R. S., Nguyen, A. L., McDermott, K. B., Shimony, J. S., Snyder, A. Z., Schlaggar, B. L., Petersen, S. E., Nelson, S. M., & Dosenbach, N. U. F. (2017). Precision Functional Mapping of Individual Human Brains. *Neuron*, *95*(4), 791–807.e7.

Grech, R., Cassar, T., Muscat, J., Camilleri, K. P., Fabri, S. G., Zervakis, M., Xanthopoulos, P., Sakkalis, V., & Vanrumste, B. (2008). Review on solving the inverse problem in EEG source analysis. *Journal of Neuro-Engineering and Rehabilitation*, *5*(1), 25.

Gutierrez-Barragan, D., Basson, M. A., Panzeri, S., & Gozzi, A. (2019). Infraslow State Fluctuations Govern Spontaneous fMRI Network Dynamics. *Current Biology*, *29*(14), 2295–2306.e5.

Haak, K. V., & Beckmann, C. F. (2020). Understanding brain organisation in the face of functional heterogeneity and functional multiplicity. *NeuroImage*, *220*, 117061.

Halgren, M., Ulbert, I., Bastuji, H., Fabó, D., Erőss, L., Rey, M., Devinsky, O., Doyle, W. K., Mak-McCully, R., Halgren, E., Wittner, L., Chauvel, P., Heit, G., Eskandar, E., Mandell, A., & Cash, S. S. (2019). The generation and propagation of the human alpha rhythm. *Proceedings of the National Academy of Sciences*, *116*(47), 23772–23782.

Handwerker, D. A., Ollinger, J. M., & D'Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *NeuroImage*, *21*(4), 1639–1651.

Haufe, S., Nikulin, V. V., Müller, K.-R., & Nolte, G. (2013). A critical assessment of connectivity measures for EEG data: A simulation study. *NeuroImage*, *64*, 120–133.

Heeger, D. J., & Ress, D. (2002). What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience*, *3*(2), 142–151.

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780.

Hopcroft, J. E., Motwani, R., & Ullman, J. D. (2001). Introduction to automata theory, languages, and computation, 2nd edition. *ACM SIGACT News*, *32*(1), 60–65.

Huntenburg, J. M., Bazin, P.-L., & Margulies, D. S. (2018). Large-Scale Gradients in Human Cortical Organization. *Trends in Cognitive Sciences*, *22*(1), 21–31.

Huster, R. J., Debener, S., Eichele, T., & Herrmann, C. S. (2012). Methods for Simultaneous EEG-fMRI: An Introductory Review. *Journal of Neuroscience*, *32*(18), 6053–6060.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Penna, S. D., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., Handwerker, D. A., Keilholz, S., Kiviniemi, V., Leopold, D. A., de Pasquale, F., Sporns, O., Walter, M., & Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, *80*, 10.1016/j.neuroimage.2013.05.079.

Iannetti, G. D., Niazy, R. K., Wise, R. G., Jezzard, P., Brooks, J. C. W., Zambreanu, L., Vennart, W., Matthews, P. M., & Tracey, I. (2005). Simultaneous recording of laser-evoked brain potentials and continuous, high-field functional magnetic resonance imaging in humans. *NeuroImage*, *28*(3), 708–719.

Jbabdi, S., Sotiropoulos, S. N., & Behrens, T. E. (2013). The topographic connectome. *Current Opinion in Neurobiology*, *23*(2), 207–215.

Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, *17*(2), 825–841.

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, *62*(2), 782–790.

Joel, S. E., Caffo, B. S., van Zijl, P. C., & Pekar, J. J. (2011). On the relationship between seed-based and ICA-based measures of functional connectivity. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, *66*(3), 644–657.

Jorge, J., van der Zwaag, W., & Figueiredo, P. (2014). EEG–fMRI integration for the study of human brain function. *NeuroImage*, *102*, 24–34.

Karapanagiotidis, T., Vidaurre, D., Quinn, A. J., Vatansever, D., Poerio, G. L., Turnbull, A., Ho, N. S. P., Leech, R., Bernhardt, B. C., Jefferies, E., Margulies, D. S., Nichols, T. E., Woolrich, M. W., & Smallwood, J. (2020). The psychological correlates of distinct neural states occurring during wakeful rest. *Scientific Reports*, *10*(1), 21121.

Katayama, H., Gianotti, L. R. R., Isotani, T., Faber, P. L., Sasada, K., Kinoshita, T., & Lehmann, D. (2007). Classes of Multichannel EEG Microstates in Light and Deep Hypnotic Conditions. *Brain Topography*, *20*(1), 7–14.

Khanna, A., Pascual-Leone, A., & Farzan, F. (2014). Reliability of Resting-State Microstate Features in Electroencephalography (T. Koenig, Ed.). *PLoS ONE*, *9*(12), e114163.

Kikuchi, M., Koenig, T., Munesue, T., Hanaoka, A., Strik, W., Dierks, T., Koshino, Y., & Minabe, Y. (2011). EEG Microstate Analysis in Drug-Naive Patients with Panic Disorder. *PLOS ONE*, *6*(7), e22912.

Kindler, J., Hubl, D., Strik, W., Dierks, T., & Koenig, T. (2011). Resting-state EEG in schizophrenia: Auditory verbal hallucinations are related to shortening of specific microstates. *Clinical Neurophysiology*, *122*(6), 1179–1182.

Kleinert, T., Koenig, T., Nash, K., & Wascher, E. (2023). On the Reliability of the EEG Microstate Approach. *Brain Topography*.

Klem, G. H., Lüders, H. O., Jasper, H. H., & Elger, C. (1999). The ten-twenty electrode system of the International Federation. The International Federation of Clinical Neurophysiology. *Electroencephalography and Clinical Neurophysiology. Supplement*, *52*, 3–6.

Koenig, T., Diezig, S., Kalburgi, S. N., Antonova, E., Artoni, F., Brechet, L., Britz, J., Croce, P., Custo, A., Damborská, A., Deolindo, C., Hein-

richs, M., Kleinert, T., Liang, Z., Murphy, M. M., Nash, K., Nehaniv, C., Schiller, B., Smailovic, U., ... Michel, C. M. (2023). EEG-Meta-Microstates: Towards a More Objective Use of Resting-State EEG Microstate Findings Across Studies. *Brain Topography*.

Koenig, T., Prichep, L., Lehmann, D., Sosa, P. V., Braeker, E., Kleinlogel, H., Isenhart, R., & John, E. (2002). Millisecond by Millisecond, Year by Year: Normative EEG Microstates and Developmental Stages. *NeuroImage*, *16*(1), 41–48.

Lanfranco, R. C., Rivera-Rei, Á., Huepe, D., Ibáñez, A., & Canales-Johnson, A. (2021). Beyond imagination: Hypnotic visual hallucination induces greater lateralised brain activity than visual mental imagery. *NeuroImage*, *239*, 118282.

Leech, R., Scott, G., Carhart-Harris, R., Turkheimer, F., Taylor-Robinson, S. D., & Sharp, D. J. (2014). Spatial Dependencies between Large-Scale Brain Networks. *PLOS ONE*, *9*(6), e98500.

Lehmann, D., Ozaki, H., & Pal, I. (1987). EEG alpha map series: Brain microstates by space-oriented adaptive segmentation. *Electroencephalography and Clinical Neurophysiology*, *67*(3), 271–288.

Lehmann, D., Faber, P. L., Galderisi, S., Herrmann, W. M., Kinoshita, T., Koukkou, M., Mucci, A., Pascual-Marqui, R. D., Saito, N., Wackermann, J., Winterer, G., & Koenig, T. (2005). EEG microstate duration and syntax in acute, medication-naïve, first-episode schizophrenia: A multi-center study. *Psychiatry Research: Neuroimaging*, *138*(2), 141–156.

Lei, X., Ostwald, D., Hu, J., Qiu, C., Porcaro, C., Bagshaw, A. P., & Yao, D. (2011). Multimodal Functional Network Connectivity: An EEG-fMRI Fusion in Network Space. *PLOS ONE*, *6*(9), e24642.

Li, M., Dahmani, L., Wang, D., Ren, J., Stocklein, S., Lin, Y., Luan, G., Zhang, Z., Lu, G., Galiè, F., Han, Y., Pascual-Leone, A., Wang, M., Fox, M. D., & Liu, H. (2021). Co-activation patterns across multiple tasks reveal robust anti-correlated functional networks. *NeuroImage*, *227*, 117680.

Liu, X., & Duyn, J. H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, *110*(11), 4392–4397.

Liu, X., Zhang, N., Chang, C., & Duyn, J. H. (2018). Co-activation patterns in resting-state fMRI signals. *NeuroImage*, *180*, 485–494.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869–878.

Manganas, S., & Bourbakis, N. (2017). A Comparative Survey on Simultaneous EEG-fMRI Methodologies. *2017 IEEE 17th International Conference on Bioinformatics and Bioengineering (BIBE)*, 1–8.

Mantini, D., Perrucci, M. G., Del Gratta, C., Romani, G. L., & Corbetta, M. (2007). Electrophysiological signatures of resting state networks in the human brain. *Proceedings of the National Academy of Sciences*, *104*(32), 13170–13175.

Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., Jefferies, E., & Smallwood, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, *113*(44), 12574–12579.

Mesulam, M. M. (1998). From sensation to cognition. *Brain*, *121*(6), 1013–1052.

Michel, C. M., & Koenig, T. (2018). EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: A review. *NeuroImage*, *180*, 577–593.

Milz, P., Faber, P., Lehmann, D., Koenig, T., Kochi, K., & Pascual-Marqui, R. (2016). The functional significance of EEG microstates—Associations with modalities of thinking. *NeuroImage*, *125*, 643–656.

Milz, P., Pascual-Marqui, R., Achermann, P., Kochi, K., & Faber, P. (2017). The EEG microstate topography is predominantly determined by intracortical sources in the alpha band. *NeuroImage*, *162*, 353–361.

Mitra, P., & Bokil, H. (2007, December). *Observed Brain Dynamics*. Oxford University Press.

Miyakoshi, M., Jurgiel, J., Dillon, A., Chang, S., Piacentini, J., Makeig, S., & Loo, S. K. (2020). Modulation of Frontal Oscillatory Power during Blink Suppression in Children: Effects of Premonitory Urge and Reward. *Cerebral Cortex Communications*, *1*(1), tgaa046.

Mulert, C., & Lemieux, L. (2023, January). *EEG - fMRI: Physiological Basis, Technique, and Applications*. Springer Nature.

Mullen, T., Kothe, C., Chi, Y. M., Ojeda, A., Kerth, T., Makeig, S., Cauwenberghs, G., & Tzyy-Ping Jung. (2013). Real-time modeling and 3D visualization of source dynamics and connectivity using wearable EEG. *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2184–2187.

Munkres, J. (1957). Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics*, *5*(1), 32–38.

Murray, M. M., Brunet, D., & Michel, C. M. (2008). Topographic ERP Analyses: A Step-by-Step Tutorial Review. *Brain Topography*, *20*(4), 249–264.

Musaeus, C. S., Nielsen, M. S., & Høgh, P. (2019). Microstates as Disease and Progression Markers in Patients With Mild Cognitive Impairment. *Frontiers in Neuroscience*, *13*.

Musso, F., Brinkmeyer, J., Mobascher, A., Warbrick, T., & Winterer, G. (2010). Spontaneous brain activity and EEG microstates. A novel EEG/fMRI analysis approach to explore resting-state networks. *NeuroImage*, *52*(4), 1149–1161.

Nehaniv, C. L., & Antonova, E. (2017). Simulating and reconstructing neurodynamics with Epsilon-automata applied to electroencephalography (EEG) Microstate Sequences. *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–9.

Niazy, R. K., Beckmann, C. F., Iannetti, G. D., Brady, J. M., & Smith, S. M. (2005). Removal of FMRI environment artifacts from EEG data using optimal basis sets. *NeuroImage*, *28*(3), 720–737.

Nishida, K., Morishima, Y., Yoshimura, M., Isotani, T., Irisawa, S., Jann, K., Dierks, T., Strik, W., Kinoshita, T., & Koenig, T. (2013). EEG microstates associated with salience and frontoparietal networks in frontotemporal dementia, schizophrenia and Alzheimer's disease. *Clinical Neurophysiology*, *124*(6), 1106–1114.

O'Gorman, R. L., Poil, S.-S., Brandeis, D., Klaver, P., Bollmann, S., Ghisleni, C., Lüchinger, R., Martin, E., Shankaranarayanan, A., Alsop, D. C., & Michels, L. (2013). Coupling Between Resting Cerebral Perfusion and EEG. *Brain Topography*, *26*(3), 442–457.

Pal, A., Behari, M., Goyal, V., & Sharma, R. (2021). Study of EEG microstates in Parkinson's disease: A potential biomarker? *Cognitive Neurodynamics*, *15*(3), 463–471.

Paquola, C., Wael, R. V. D., Wagstyl, K., Bethlehem, R. A. I., Hong, S.-J., Seidlitz, J., Bullmore, E. T., Evans, A. C., Misic, B., Margulies, D. S., Smallwood, J., & Bernhardt, B. C. (2019). Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLOS Biology*, *17*(5), e3000284.

Pascual-Marqui, R. D. (1999). Review of Methods for Solving the EEG Inverse Problem. *1*(1), 13.

Pascual-Marqui, R. D., Michel, C. M., & Lehmann, D. (1995). Segmentation of brain electrical activity into microstates: Model estimation and validation. *IEEE Transactions on Biomedical Engineering*, *42*(7), 658–665.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, *12*(85), 2825–2830.

Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, *198*, 181–197.

Poulsen, A. T., Pedroni, A., Langer, N., & Hansen, L. K. (2018). Microstate EEGlab toolbox: An introductory guide. *bioRxiv*.

Preti, M. G., Bolton, T. A., & Van De Ville, D. (2017). The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage*, *160*, 41–54.

Pruim, R. H. R., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K., & Beckmann, C. F. (2015). ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage*, *112*, 267–277.

Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W., LaMantia, A.-S., & White, L. (2019, June). *Neurosciences*. De Boeck Superieur.

Rhodes, J., & Nehaniv, C. L. (2009). *Applications Of Automata Theory And Algebra: Via The Mathematical Theory Of Complexity To Biology, Physics,*

Psychology, Philosophy, And Games. World Scientific Publishing Company.

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex (New York, NY)*, *28*(9), 3095–3114.

Schlegel, F., Lehmann, D., Faber, P. L., Milz, P., & Gianotti, L. R. R. (2012). EEG Microstates During Resting Represent Personality Differences. *Brain Topography*, *25*(1), 20–26.

Schumacher, J., Peraza, L. R., Firbank, M., Thomas, A. J., Kaiser, M., Gallagher, P., O'Brien, J. T., Blamire, A. M., & Taylor, J.-P. (2019). Dysfunctional brain dynamics and their origin in Lewy body dementia. *Brain*, *142*(6), 1767–1782.

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., & Greicius, M. D. (2007). Dissociable Intrinsic Connectivity Networks for Salience Processing and Executive Control. *Journal of Neuroscience*, *27*(9), 2349–2356.

Seitzman, B. A., Abell, M., Bartley, S. C., Erickson, M. A., Bolbecker, A. R., & Hetrick, W. P. (2017). Cognitive manipulation of brain electric microstates. *NeuroImage*, *146*, 533–543.

Serrano, J. I., del Castillo, M. D., Cortés, V., Mendes, N., Arroyo, A., Andreo, J., Rocon, E., del Valle, M., Herreros, J., & Romero, J. P. (2018). EEG Microstates Change in Response to Increase in Dopaminergic Stimulation in Typical Parkinson's Disease Patients. *Frontiers in Neuroscience*, *12*.

Sikka, A., Jamalabadi, H., Krylova, M., Alizadeh, S., van der Meer, J. N., Danyeli, L., Deliano, M., Vicheva, P., Hahn, T., Koenig, T., Bathula, D. R., & Walter, M. (2020). Investigating the temporal dynamics of electroencephalogram (EEG) microstates using recurrent neural networks. *Human Brain Mapping*, *41*(9), 2334–2346.

Silva, L. R., Amitai, Y., & Connors, B. W. (1991). Intrinsic Oscillations of Neocortex Generated by Layer 5 Pyramidal Neurons. *Science*, *251*(4992), 432–435.

Skrandies, W. (1990). Global field power and topographic similarity. *Brain Topography*, *3*(1), 137–141.

Smallwood, J., Bernhardt, B. C., Leech, R., Bzdok, D., Jefferies, E., & Margulies, D. S. (2021). The default mode network in cognition: A topographical perspective. *Nature Reviews Neuroscience*, *22*(8), 503–513.

Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., Van Essen, D. C., Feinberg, D. A., Yacoub, E. S., & Ugurbil, K. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, *109*(8), 3131–3136.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, *17*(3), 143–155.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., & Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, *106*(31), 13040–13045.

Soni, S., Muthukrishnan, S. P., Sood, M., Kaur, S., & Sharma, R. (2018). Hyperactivation of left inferior parietal lobule and left temporal gyri shortens resting EEG microstate in schizophrenia. *Schizophrenia Research*, *201*, 204–207.

Sten, S., Lundengård, K., Witt, S. T., Cedersund, G., Elinder, F., & Engström, M. (2017). Neural inhibition can explain negative BOLD responses: A mechanistic modelling and fMRI study. *NeuroImage*, *158*, 219–231.

Strelets, V., Faber, P., Golikova, J., Novototsky-Vlasov, V., Koenig, T., Gianotti, L., Gruzelier, J., & Lehmann, D. (2003). Chronic schizophrenics with positive symptomatology have shortened EEG microstate durations. *Clinical Neurophysiology*, *114*(11), 2043–2051.

Strik, W. K., Chiaramonti, R., Muscas, G. C., Paganini, M., Mueller, T. J., Fallgatter, A. J., Versari, A., & Zappoli, R. (1997). Decreased EEG microstate duration and anteriorisation of the brain electrical fields in mild and moderate dementia of the Alzheimer type. *Psychiatry Research: Neuroimaging*, *75*(3), 183–191.

Tait, L., Tamagnini, F., Stothart, G., Barvas, E., Monaldini, C., Frusciante, R., Volpini, M., Guttmann, S., Coulthard, E., Brown, J. T., Kazanina, N., & Goodfellow, M. (2020). EEG microstate complexity for aiding early diagnosis of Alzheimer's disease. *Scientific Reports*, *10*(1), 17627.

Tait, L., & Zhang, J. (2022). +microstate: A MATLAB toolbox for brain microstate analysis in sensor and cortical EEG/MEG. *NeuroImage*, *258*, 119346.

Tarailis, P., Koenig, T., Michel, C. M., & Griškova-Bulanova, I. (2023). The Functional Aspects of Resting EEG Microstates: A Systematic Review. *Brain Topography*.

Teplan, M. (2002). Fundamentals of EEG Measurement. *Measurement Science Review*, *2,2*.

Tomescu, M. I., Rihs, T. A., Rochas, V., Hardmeier, M., Britz, J., Allali, G., Fuhr, P., Eliez, S., & Michel, C. M. (2018). From swing to cane: Sex differences of EEG resting-state temporal patterns during maturation and aging. *Developmental Cognitive Neuroscience*, *31*, 58–66.

Tomescu, M. I., Papasteri, C. C., Sofonea, A., Boldasu, R., Kebets, V., Pistol, C. A. D., Poalelungi, C., Benescu, V., Podina, I. R., Nedelcea, C. I., Berceanu, A. I., & Carcea, I. (2022). Spontaneous thought and microstate activity modulation by social imitation. *NeuroImage*, *249*, 118878.

Tops, M., & Boksem, M. (2011). A Potential Role of the Inferior Frontal Gyrus and Anterior Insula in Cognitive Control, Brain Rhythms, and Event-Related Potentials. *Frontiers in Psychology*, *2*.

Uddin, L. Q., Yeo, B. T. T., & Spreng, R. N. (2019). Towards a Universal Taxonomy of Macro-scale Functional Human Brain Networks. *Brain Topography*, *32*(6), 926–942.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, *80*, 62–79.

Vellante, F., Ferri, F., Baroni, G., Croce, P., Migliorati, D., Pettoruso, M., De Berardis, D., Martinotti, G., Zappasodi, F., & Giannantonio, M. D. (2020). Euthymic bipolar disorder patients and EEG microstates: A neural signature of their abnormal self experience? *Journal of Affective Disorders*, *272*, 326–334.

von Wegner, F., Bauer, S., Rosenow, F., Triesch, J., & Laufs, H. (2021). EEG microstate periodicity explained by rotating phase patterns of resting-state alpha oscillations. *NeuroImage*, *224*, 117372.

von Wegner, F., Tagliazucchi, E., & Laufs, H. (2017). Information-theoretical analysis of resting state EEG microstate sequences - non-Markovianity, non-stationarity and periodicities. *NeuroImage*, *158*, 99–111.

von Wegner, F., Knaut, P., & Laufs, H. (2018). EEG Microstate Sequences From Different Clustering Algorithms Are Information-Theoretically Invariant. *Frontiers in Computational Neuroscience*, *12*.

Vos de Wael, R., Royer, J., Tavakol, S., Wang, Y., Paquola, C., Benkarim, O., Eichert, N., Larivière, S., Xu, T., Misic, B., Smallwood, J., Valk, S. L., & Bernhardt, B. C. (2021). Structural Connectivity Gradients of the Temporal Lobe Serve as Multiscale Axes of Brain Organization and Cortical Evolution. *Cerebral Cortex*, *31*(11), 5151–5164.

Wagstyl, K., Ronan, L., Goodyer, I. M., & Fletcher, P. C. (2015). Cortical thickness gradients in structural hierarchies. *NeuroImage*, *111*, 241–250.

Winkler, I., Haufe, S., & Tangermann, M. (2011). Automatic Classification of Artifactual ICA-Components for Artifact Removal in EEG Signals. *Behavioral and Brain Functions*, *7*(1), 30.

Xu, J., Pan, Y., Zhou, S., Zou, G., Liu, J., Su, Z., Zou, Q., & Gao, J.-H. (2020). EEG microstates are correlated with brain functional networks during slow-wave sleep. *NeuroImage*, *215*, 116786.

Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). NeuroSynth: A new platform for large-scale automated synthesis of human functional neuroimaging data. *Frontiers in Neuroinformatics*, *5*.

Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., & Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165.

Yu, Y., Si, X., Hu, C., & Zhang, J. (2019). A Review of Recurrent Neural Networks: LSTM Cells and Network Architectures. *Neural Computation*, *31*(7), 1235–1270.

Yuan, H., Zotev, V., Phillips, R., Drevets, W. C., & Bodurka, J. (2012). Spatiotemporal dynamics of the brain at rest — Exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. *NeuroImage*, *60*(4), 2062–2072.

Zanesco, A. P., Denkova, E., & Jha, A. P. (2021). Associations between self-reported spontaneous thought and temporal sequences of EEG microstates. *Brain and Cognition*, *150*, 105696.