

A dynamic system of brain networks revealed by fast transient EEG fluctuations and their fMRI correlates

B. Hunyadi^{1,2}, M. Woolrich^{3*}, A. Quinn³, D. Vidaurre³, and M. De Vos^{4*}

¹ STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Department of Electrical Engineering (ESAT), KU Leuven

² imec, Leuven, Belgium

³ Oxford Centre for Human Brain Activity, Wellcome Centre for Integrative Neuroimaging, Department of Psychiatry, University of Oxford

⁴ Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford, United Kingdom

* equal contribution

ABSTRACT

Resting state brain activity has become a significant area of investigation in human neuroimaging. An important approach for understanding the dynamics of neuronal activity in the resting state is to use complementary imaging modalities.

Electrophysiological recordings can access fast temporal dynamics, while functional magnetic resonance imaging (fMRI) studies reveal detailed spatial patterns. However, the relationship between these two measures is not fully established. In this study, we used simultaneously recorded electroencephalography (EEG) and fMRI, along with Hidden Markov Modelling, to investigate how network dynamics at fast sub-second time-scales, accessible with EEG, link to the slower time-scales and higher spatial detail of fMRI. We found that the fMRI correlates of fast transient EEG dynamic networks show highly reproducible spatial patterns, and that their spatial organization exhibits strong similarity with traditional fMRI resting state networks maps. This further demonstrates the potential of electrophysiology as a tool for understanding the fast network dynamics that underlie fMRI resting state networks.

I. INTRODUCTION

The human brain is a complex system of large neuronal populations, which interact to perform different cognitive, motor and perceptual tasks. Brain regions that tend to work together during task execution also co-activate during rest [Smith2009], giving rise to so-called resting state networks (RSNs) [Damoiseaux2006]. Using fMRI, highly similar resting state network maps have been found consistently across subjects, studies and even different analysis techniques [Beckmann2005, VanDijck2010]. However, the slow dynamics of fMRI signals preclude the investigation of the fast dynamic interactions between brain regions, which are essential to explain the actual neural mechanisms that underlie behaviour [Britz2010, Florin2015].

Electrophysiological recordings, on the other hand, can more directly access brain network dynamics at fast time-scales. Seed-based functional connectivity analysis [dePasquale2010; Brookes2012] or independent component analysis (ICA) [Liu2017], which are commonly used techniques to reveal fMRI RSNs, have been applied to MEG data [Brookes2011; de Pasquale2010] and to source-space EEG data [Liu2017], revealing similar spatial patterns. Microstate analysis, which was specifically developed to study spontaneous EEG fluctuations, builds on the observation that the EEG topographies show quasi-stable periods in the sub-second range [Lehmann2009]. Such EEG microstates were shown to be relatively similar across subjects age groups [Koenig2002]. Recently, techniques based on Hidden Markov Modelling (HMM) have been used to identify consistent patterns of fast, sub-second transient brain states that resemble fMRI-derived RSNs [Baker2014; Vidaurre2017b, Vidaurre2017c].

In order to establish a more complete link between electrophysiologically-derived brain networks and fMRI-derived brain networks, simultaneous EEG and fMRI recordings are needed [Florin2015]. This can provide direct evidence of how the temporal fluctuations are related between the two modalities in ongoing, resting brain activity. In the pioneering work of [Goldman2002], the fMRI signal showed significant decrease in multiple cortical regions, and significant increase in the thalamus and insula in response to fluctuations in EEG alpha band power. Subsequent work found, in response to alpha and beta power increase, fMRI deactivations and activations in networks associated with attention and spontaneous cognitive operations, respectively [Laufs2003]. Not only power increase, but global synchronization in alpha band was proven to be linked to fMRI fluctuations in the corresponding network [Jann2009]. In a symmetric, data-driven approach comparing spontaneous fMRI fluctuations derived by ICA and EEG rhythms at typical frequency bands within 1-50Hz, Mantini et al found a more complex relationship between the two modalities [Mantini2007]. Namely, each ICA-based fMRI RSN were characterized by a specific electrophysiological signature, combining all the different brain rhythms. Interestingly, different parsing of resting-state EEG activity can lead to unique associations across the networks of the two modalities. [Hiltunen2014] found highly specific matching across independent components of infra-slow (0.01-0.1Hz) EEG and a selected subset of ICA-based fMRI RSNs. In another study, four of the well-known set of RSNs were uniquely linked to the four typical microstate maps [Britz2010]. This finding is particularly interesting as it suggests that much faster fluctuations of brain activity underlie RSNs than what is evident from fMRI recordings. The scale-free behaviour of microstates over a wide range of time-scales may provide an explanation to this finding [Vandeville2010, Gschwind2015].

Our approach builds on these results to investigate further the links between the network dynamics at fast sub-second time-scales, accessible with EEG, and brain network activity at slower (>1 sec) time-scales in fMRI, in simultaneously acquired data. To access brain network dynamics at fast (<1 sec) time-scales in electrophysiology we use the Hidden Markov Model (HMM) mathematical framework, which was previously used to identify fast transient dynamic networks in MEG data [Baker2014; Vidaurre2016, Vidaurre2017c]. First, we use the HMM to estimate a dynamic system to model the succession of fast transient networks from EEG so that we can examine its temporal characteristics. Then, using simultaneously recorded fMRI data, we interrogate the relationship of fMRI resting state network maps and the spatial signatures of electrophysiological fast transient networks, in order to bridge the gap between these two data modalities. Following this approach, we found that the fMRI correlates of fast transient networks show highly reproducible spatial patterns, and that their spatial organization shows strong similarity with fMRI resting state network maps. These results suggest that there is a dynamic system of fast transient brain states underlying fMRI resting state network fluctuations. Although previous studies have already shown a link between these two modalities, these links were restricted only to a subset of resting state networks, and established a one-to-one correspondence between them and specific EEG phenomena. In contrast, we show a more complex interrelationship between a large number of fast transient EEG networks and the full set of known resting state networks, leading to a richer characterisation of resting state brain dynamics.

II. MATERIALS AND METHODS

1. Data acquisition and preprocessing

Simultaneous EEG and fMRI data were recorded from 24 healthy subjects (14 females, mean age: 22.7 years, ranging between 18-26 years), all right handed, with no history of neurological or psychiatric disorders. Data were recorded during 6 minutes while the subjects were asked to rest with eyes closed. Due to severe head movement (total displacement or scan-to-scan movement larger than 3mm or 1mm, respectively) or high amount of EEG artefacts, data from 3 subjects were excluded from the analysis. Therefore, the final dataset included recordings from 21 subjects (13 females). The experimental procedure was approved by the ethics committee of the University of Oldenburg. The study was conducted as part of a more extensive recording protocol [Puschmann2016], in accordance with the Declaration of Helsinki [World Medical Association, 2013] after having obtained written informed consent from all subjects.

fMRI images were acquired using a 3T Siemens MAGNETOM Vero MRI scanner (Siemens AG, Erlangen, Germany) with a twelve-channel head array. We obtained 261 T2*-weighted gradient echo planar imaging (EPI) volumes with fMRI-contrast (repetition time (TR) = 1500 ms, echo-time (TE) = 30 ms, flip angle $\alpha=70^\circ$, field of view (FOV) = 200 x 200 mm², voxel size 3.1 x 3.1 x 3.1 mm³). Volumes were recorded in 23 transverse slices with a gap of 0.9 mm in between in ascending order. Afterwards, a high-resolution structural volume was also acquired using a T1-weighted magnetization-prepared rapid acquisition gradient echo (MPRAGE) sequence (TR = 1900ms, TE = 2.25ms, $\alpha = 9^\circ$, FoV = 256 x 256 mm², voxel size = 1 x 1 x 1 mm³).

fMRI time series were preprocessed using FSL [Jenkinson]. More specifically, to correct for head motion the fMRI images were aligned to the middle volume using linear rigid-body transformations. Using a brain mask created from the first fMRI volume, voxels outside the brain were excluded. The images were spatially smoothed with a Gaussian kernel of 5 mm full-width-half-maximum, and voxel time series were finally high-pass filtered at 0.01Hz. The individual fMRI data were subject to spatial independent component analysis (ICA) with 20 components using MELODIC, and artefact-related components were selected based on their lack of spatial correlation with known ICA-based RSNs¹ [Shirer2012]. Clean fMRI data were reconstructed by regressing out the artefact components and motion confounds using the ‘unique variance’ approach implemented in the Fix FSL toolbox [Griffanti2015]. Afterwards, fMRI data were coregistered to each subject’s structural image and were normalized to MNI space. Subsequently, fMRI from all subjects were concatenated and 20 spatially independent components were extracted using group ICA. Out of these 20 group components, 4 were identified as being likely caused by cardiac pulsation, considering that their spatial maps included the ventricles and the main arteries [Salimi-Khorshidi2014]. The remaining 16 components were considered neuronal. The preprocessed fMRI for each subject, which was used for the subsequent analysis, was obtained by reconstructing the fMRI time series using the estimated mixing matrix and the 16 neuronal independent components.

EEG signals were recorded during the fMRI data acquisition using 64 MRI-compatible electrodes and amplifier (BrainAmp MR Plus, Brainproducts, Gilching, Germany), sampled at 5000 Hz and electrode impedances kept below 20k Ω . An equidistant electrode layout was used with AFz and Cz as ground and online recording reference, respectively. Eye movements were monitored with an EOG electrode placed below the left eye, while cardiac activity was recorded using an ECG electrode placed on the left lower back.

EEG data were preprocessed as follows. Gradient artefacts were removed using an artefact template subtraction method (FASTR) constructed from averaging the artefacts from 30 consecutive fMRI volumes, residual artefacts were reduced by the optimal basis set (OBS) method [Niazy2005], and an adaptive noise cancellation procedure [Allen2000] was applied.

¹ https://findlab.stanford.edu/functional_ROIs.html

Subsequently, ballistocardiographic artefacts occurring at each cardiac cycle (detected based on the recorded ECG signal) were removed from each EEG channel using the OBS method [Niazy2005, Vanderperren2010], based on the first 4 principal components. These procedures were implemented in the FMRIB plugin for the EEGLAB toolbox [DelormeMakeig2004]. Next, residual scanner artefacts and eye movement artefacts were removed by manually rejecting artefact sources extracted by independent component analysis using fastICA [Hyvarinnen1999]. Manual identification of the artefact sources was guided by the topography, time course, frequency spectrum and kurtosis of the sources, and their temporal correlation with EOG and ECG channels.

After artefact removal, and following [Baker2014], the EEG data were band-pass filtered at 4-30Hz. The amplitude envelopes were computed using the Hilbert transform, which were then downsampled to 40Hz. The resulting time series will be referred to as ‘broadband EEG envelopes’ in the following analysis.

2. HMM-based fast transient networks in EEG

Following the approach previously used by [Baker2014] on MEG data, fast transient networks were identified using a Hidden Markov Model (HMM) on the group concatenated broadband EEG envelopes. Prior to fitting the HMM, the channel dimension of the data were reduced to 9 principal components (PCs, 99% variance of the original data), resulting in the observation matrix $Y=[y_1 \dots y_T]$, with $y_t \in \mathbb{R}^M$, where $M=9$ and T is the number of samples in the concatenated broadband EEG envelopes. Note that retaining 25 PCs (99.9% variance) has led to an inappropriate model where 2 subjects were dominated by a single HMM state. The HMM assumes that a hidden sequence of states underlies the observed EEG envelope data and that for any given time point, the value of the EEG envelopes is assumed to be drawn from a state-specific observation model. Here, each state’s observation model is characterised by a specific amplitude pattern and a network of functional connectivity, where functional connectivity is defined as amplitude correlation.

Specifically, we denote the hidden state variables by $s=\{s_1 \dots s_T\}$, which is a vector of state labels indicating which of the K hidden states is active at each point in time. The EEG envelope data observed in a particular state k is assumed to be generated from a multivariate normal distribution, with parameters $\Theta_k=\{\mu_k, \Sigma_k\}$, where $\mu_k \in \mathbb{R}^M$ is the mean and $\Sigma_k \in \mathbb{R}^{M \times M}$ is the covariance matrix. The observation model for state k can then be formally written as

$$P(y_t | s_t = k, \theta) \sim N(\mu_k, \Sigma_k)$$

In practice, the inference process estimates a probability for each state to be active at each time point, such that the data at each time point is effectively modelled as a mixture of Gaussian distributions. After the inference, the Viterbi algorithm [Bishop2006] can be used to compute the “hard-assigned” (non-probabilistic) sequence of states with the highest likelihood. (In this paper, figures are based on the so-called Viterbi path for simplicity). The HMM then assumes that the transitions between the hidden states are *Markovian*, i.e. the probability to transition to another state depends directly only on the current state; more formally, the probability of the next state is conditionally independent of previous states given the current state:

$$P(s_t | s_1 \dots s_{t-1}) = P(s_t | s_{t-1}) = \pi_t$$

where $\pi_t \in \mathbb{R}^{K \times K}$ is the transition probability matrix, which is assumed to be stationary. The (i,j) element of the transition probability matrix is the probability of transitioning from state i to state j , denoted by $P(s_t=j | s_{t-1}=i)$. The initial state probability $P(s_0)$ is parameterized by π_0 . The full true posterior probability of the model is written as

$$P(y, s, \theta) = P(s_0 | \pi_0) \prod_t^T P(s_t | s_{t-1}, \pi_t) P(y_t | s_t, \theta) P(\pi_t) P(\pi_0) P(\theta)$$

where the priors $P(\pi_t)$, $P(\pi_0)$ and $P(\Theta)$ are chosen to be conjugate distributions. As in [Vidaurre2017a], we used a variational Bayesian (VB) inference, which is implemented in the HMM-MAR toolbox² to obtain the full posterior distributions on the model parameters. Finally, the states u_t active at each time point t are assigned by the Viterbi algorithm [Bishop2006].

3. Temporal dynamics

We describe the temporal characteristics of the inferred HMM states using a number of measures. The fractional occupancy expresses the fraction of time spent in each state relative to the total duration T (in samples) of the experiment:

$$\text{fractional occupancy} = \frac{1}{T} \sum_t (u_t == k)$$

The mean *life (or dwell) times* of each state is the time spent in a state before transitioning to a new state on average:

$$\text{mean life time} = \frac{\sum_t u_t == k}{\text{number of occurrences } (k)}$$

We also considered the temporal interaction between the inferred states. This was achieved by inspection of the transition probability matrix, π_t , which describes how likely a transition is from a given state to another state. To inspect interactions on a longer time scale, the fractional occupancies of states were computed over 10s long sliding-windows, and the correlations between these windowed fractional occupancies were quantified for each pair of states.

4. HMM state activation and EEG power fluctuations

To investigate how each HMM state is related to fluctuations of global EEG power, the broadband EEG envelopes were averaged across the channels. The obtained global broadband EEG power fluctuation time course and the HMM state time courses were compared using Pearson's correlation coefficient. A positive or negative correlation means that a given HMM state is active at times when the global broadband EEG power is high or low, respectively. Moreover, to inspect the EEG spatial properties of the states, we regressed the state time courses on the concatenated multichannel broadband EEG envelopes, as in [Baker2014]. The resulting regression coefficients were used to construct EEG topographies for each state.

5. fMRI correlates of fast transient networks

a. Computation and Significance testing

The HMM state time courses were convolved with the canonical hemodynamic response function of the SPM toolbox (delay of response 6s, delay of undershoot 16s, length of kernel 32s) and downsampled to 1/1.5 Hz to match the temporal sampling rate of the fMRI time courses. Using these transformed HMM state time courses as independent variables and the preprocessed fMRI time courses as dependent variables in a voxel-wise multiple regression model, *beta* maps were computed for each of the K states.

To assess the significance of the beta values, a voxel-wise null distribution for the beta coefficients in each of the K global HMM states was computed by using 1000 Monte-Carlo simulated state time courses. These simulated state time courses were obtained with dynamics matched to the real state

² <https://github.com/OHBA-analysis/HMM-MAR>

time courses by sampling random state sequences using the transition probability matrix obtained from the real HMM results (obtained with all 21 subjects' data). In a multiple regression analysis, beta coefficients between the Monte-Carlo simulated state time courses (independent variables) and the fMRI time course (dependent variable) were computed in each voxel. We refer to HMM-EEG-fMRI maps as the result of thresholding the beta maps according to this null-distribution at $p < 0.05$.

b. Comparison with ICA-based fMRI RSNs

To assess the relationship between the fast, transient networks of EEG and resting state networks, the HMM-EEG-fMRI maps were compared to resting-state network maps obtained with group-ICA of fMRI (referred to as ICA-fMRI in the rest of the paper) from the same multi-subject dataset in the following two-step procedure. First, the RSNs were compared with a known functional ROI atlas³ [Shirer2012]. We found a nearly one-to-one correspondence between the networks of this atlas and the extracted ICA-fMRI maps (see Supplementary Figure 2). Therefore, we assigned the interpretation and naming to our ICA-fMRI maps based on the corresponding networks of this atlas. Afterwards, we matched the HMM-EEG-fMRI maps with the data-driven ICA-fMRI maps. As a comparison metric, we computed the Pearson's correlation coefficients between the vectorised HMM-EEG-fMRI activation maps and the vectorised ICA-fMRI maps. Using the same Monte-Carlo simulated state time courses and resulting beta-maps as above, we generated a null distribution of correlation values between 1000 simulated HMM-EEG-fMRI beta maps and the ICA-fMRI maps, which was used to assess the significance of the correlations obtained in the true model. A correlation was considered significant at $p < 0.05$ (after Bonferroni's correction for multiple comparisons).

6. Reproducibility of HMM-EEG-fMRI beta maps

The original set of 21 subjects was randomly split into 2 groups, one with 10 and the other with 11 subjects. Ten different random splits were generated. Subsequently, HMM-based inference of fast transient states was performed separately on each group of each random split, and the HMM-EEG-fMRI beta maps were constructed using the same approach as described above in section 5a. In what follows, the HMM results based on all 21 patients will be referred to as 'global', while the ones based on the 2 subsets of subjects will be referred to as 'group 1' and 'group 2'. Due to the independent HMM inference, the ordering of the corresponding states in each group is ambiguous, and different from the state ordering of the global model. Therefore, each HMM-EEG-fMRI beta map obtained in the split-half analysis was separately matched to the global fMRI maps, based on the correlation of the vectorised maps. Afterwards, the correlations between the matched HMM-EEG-fMRI beta maps of group 1 and group 2 were computed in order to quantify their reproducibility.

For comparison, we also computed the reproducibility of ICA-fMRI maps and HMM-fMRI maps, where the HMM states were inferred directly on the fMRI data as in [Vidaurre2017a, Vidaurre2017b]. Prior to fitting the HMM, the fMRI data were preprocessed as described in section II/1. Then, similarly to the EEG processing pipeline, the spatial dimension of the data was reduced to 14 PCs to retain 99% of variance. In order to measure the reproducibility of the ICA-fMRI maps and HMM-fMRI maps, we used the same 10 random splits as for the analysis of the HMM-EEG-fMRI maps. As above, the maps obtained in each split are ordered differently and need to be matched based on their correlation to the global ICA-fMRI maps

³ https://findlab.stanford.edu/functional_ROIs.html

and the global HMM-fMRI beta maps, respectively. Finally, the reproducibility of the maps was expressed based on the correlation of the matched maps across group 1 and group 2.

In summary, by comparing the reproducibility of ICA-fMRI maps and HMM-EEG-fMRI maps we aim to investigate whether the HMM-EEG-fMRI analysis pipeline produces fMRI spatial maps that are as robust and reproducible as the ones produced by the widely used ICA-fMRI approach. Moreover, given that our proposed HMM-EEG-fMRI approach differs from the classic ICA-fMRI analysis in terms of both data modality (EEG versus fMRI) and methodology (HMM versus ICA), we use the comparison with the HMM-fMRI beta maps as an intermediate step between the two approaches in order to clarify the potential differences between their reproducibility.

7. Summary of network analysis approaches

Table 1 below shows an overview of all the different EEG- and fMRI-based network analysis approaches we have performed.

Table 1: A summary of network analysis approaches

	HMM-EEG	HMM-EEG-fMRI	HMM-fMRI	ICA-fMRI
Input	PCA-reduced EEG sensor envelopes	HMM-EEG state activation time courses, fMRI voxel time courses	PCA-reduced fMRI data	fMRI data
Analysis step 1	HMM with K=16	Multiple regression	HMM with K=16	Spatial groupICA with M=20 components
Output step 1	16 hidden states and their activation time courses	HMM-EEG-fMRI “activation” (fMRI+) and “deactivation” (fMRI-) maps	16 hidden states and their activation time courses	20 spatial maps and their time courses
Input step 2	EEG sensor envelopes, HMM-EEG state activation time courses	-	fMRI data, HMM-fMRI state activation time courses	-
Analysis step 2	Multiple regression	-	Multiple regression	-
Output step 2	HMM-EEG topographies	-	HMM-fMRI “activation” and “deactivation” maps	-
Interpretation	Dynamic sequence of fast transient EEG networks and their spatial characterization in EEG sensor space	Activation and deactivation maps consist of voxels where the HRF-convolved HMM-EEG activation time course significantly positively or negatively correlates with the fMRI voxel time course	Activation and deactivation maps consist of voxels where the HMM-fMRI activation time course significantly positively or negatively correlated with the fMRI voxel time course	Resting state network (RSN) maps which are mutually spatially independent and their corresponding temporal fluctuations which are representative of the voxel time courses within each map

III. RESULTS

We inferred a Hidden Markov Model (HMM) with $K=16$ hidden states (see SI for a discussion on the choice of the number of states) from the temporally concatenated EEG of the multi-subject dataset after dimensionality reduction (see Methods). According to our hypothesis, each hidden state corresponds to a distinct brain network. As such, the HMM provides a detailed spatiotemporal characterization of the underlying system of brain networks. The HMM state time course reveals when each brain network is active and how they are sequenced. Moreover, regressing the state time course on the multichannel broadband EEG envelopes or the fMRI voxel time courses reveals the anatomical properties of the networks, in terms of state EEG sensor-space topographies or HMM-EEG-fMRI activation maps, respectively. For illustration, Figure 1 shows a representative 1s long segment of multichannel broadband EEG envelope from subject 1, along with the corresponding inferred state time course. During this 1s period, 3 out of the 16 states activate in turn (a 4th state being active very briefly). The computed state topographies and HMM-EEG-fMRI maps are shown below.

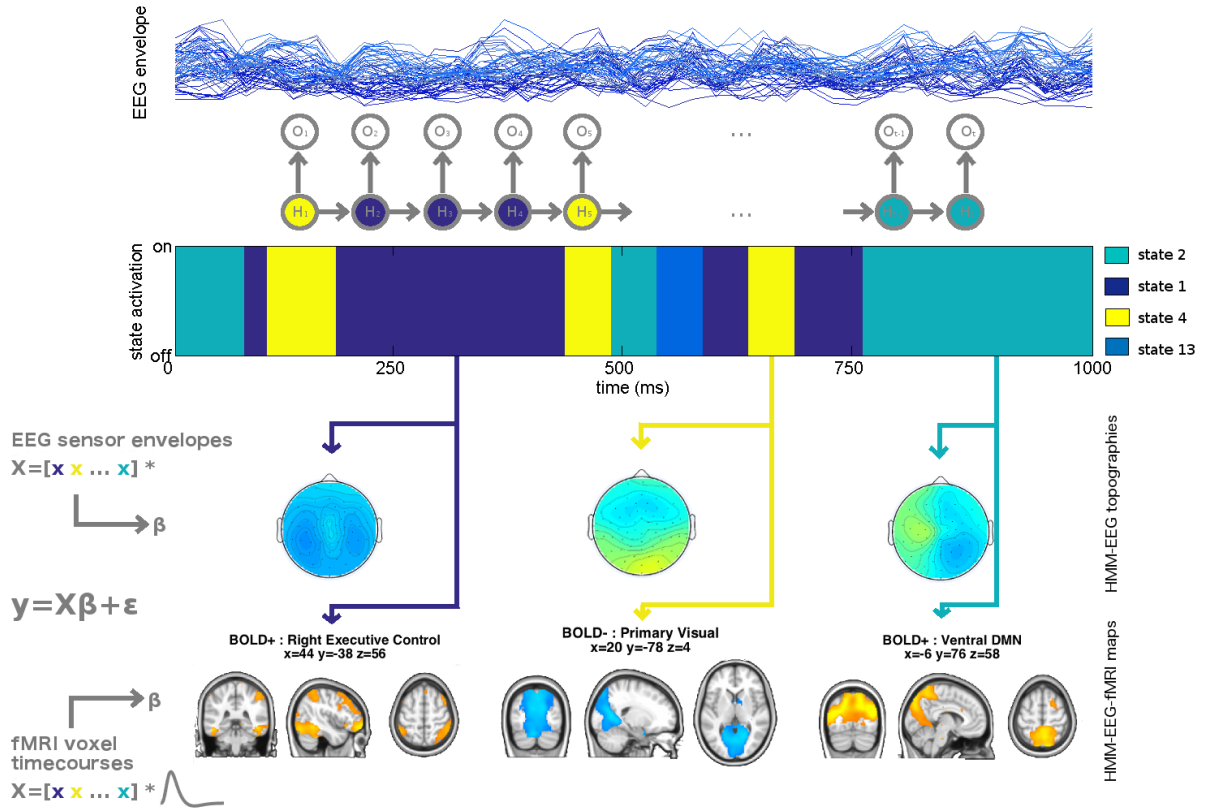


Figure 1: Fast transient EEG networks analysis pipeline: We fitted a $K=16$ HMM model on the Hilbert envelopes of the multichannel EEG, i.e. the observations (O_t). Each of the 16 hidden states (H_t) corresponds to a distinct brain network. One state stays active during a certain short time interval (indicated in green, dark blue, yellow or light blue colour), and then the brain transitions to the next state. Multiple regression analysis is performed separately at every EEG sensor or fMRI voxel, using the broadband EEG envelopes or fMRI time course data as dependent variables (y), and the state activation time courses (convolved with the hemodynamic response function in the fMRI case) as independent variables (X). Finally, assembling the computed regression coefficients (β) per state, we create HMM-EEG topographies or HMM-EEG-fMRI activations maps, respectively, to characterise the spatial organization of fast transient networks.

1. Spatial organization of fast transient networks

The spatial organization of resting state brain activity has been widely studied using fMRI, and reliable resting state network maps have been reproduced in numerous studies. Here, we investigate whether the fMRI correlates of fast transient EEG dynamic networks, i.e. the HMM-EEG-fMRI maps, resemble canonical fMRI RSN maps. As described in section II/4b, this similarity is assessed in terms of the spatial correlation between the (unthresholded) HMM-EEG-fMRI maps and the (unthresholded) RSN maps extracted from fMRI using ICA, i.e. the ICA-fMRI maps. Note that there is an important difference between the ICA-fMRI maps and HMM-EEG-fMRI maps: most voxels in the ICA-fMRI maps are positive, and the range of positive voxel values is much larger than that of the negative values. As such, only the positive side of the maps is typically interpreted. On the contrary, many HMM-EEG-fMRI maps have a considerable amount of positive and negative voxels as well, both of which can be interpreted. In what follows, we will refer to the positive and the negative side of the maps as HMM-EEG-fMRI “activation” (fMRI+) and “deactivation” (fMRI-). This reflects whether the HRF-convolved HMM-EEG state time-courses are positively or negatively temporally correlated with the fMRI data in a particular voxel.

In the top panel of **Figure 2** we visualize the results of the similarity analysis. Six out of 16 HMM states have HMM-EEG-fMRI maps with both “activations” and “deactivations” that significantly spatially correlate (or anti-correlate) with one of the ICA-fMRI maps. Five HMM states have HMM-EEG-fMRI maps with only “activations”, and five others have only “deactivations” that match (i.e. spatially significantly correlate with) an ICA-fMRI map. Six HMM-EEG-fMRI activation maps and two deactivation maps correlate significantly with more than one ICA-fMRI.

Overall, each HMM-EEG-fMRI map matches with at least one ICA-fMRI map, and, similarly, each known ICA-fMRI map (as described by the atlas of [Shirer2012]) matches with at least one HMM-EEG-fMRI map. Although there is no one-to-one correspondence, each fast transient network can be characterized by the combination of only a small number of ICA-fMRI RSNs, up to 3. In light of this, we will refer to each state according to its best matching ICA-fMRI RSN in both positive and negative direction, if it exists. For example, we refer to state 6 as *Auditory+/dorsal DMN-*, to state 11 as *Precuneus-*, to state 9 as *Sensorimotor+*, and state 4 as *Primary Visual-*. These examples are also visualized in **Figure 2B** along with other representative examples. In Supplementary Figure 3 we show all HMM-EEG-fMRI maps together with their best matching ICA-fMRI RSNs. Note that the ordering/numbering of the states is arbitrary. In summary, these significant spatial correlations (and anticorrelations) between HMM-EEG-fMRI and ICA-fMRI maps suggest that there are fast transient networks captured by EEG underlying fMRI resting state network fluctuations.

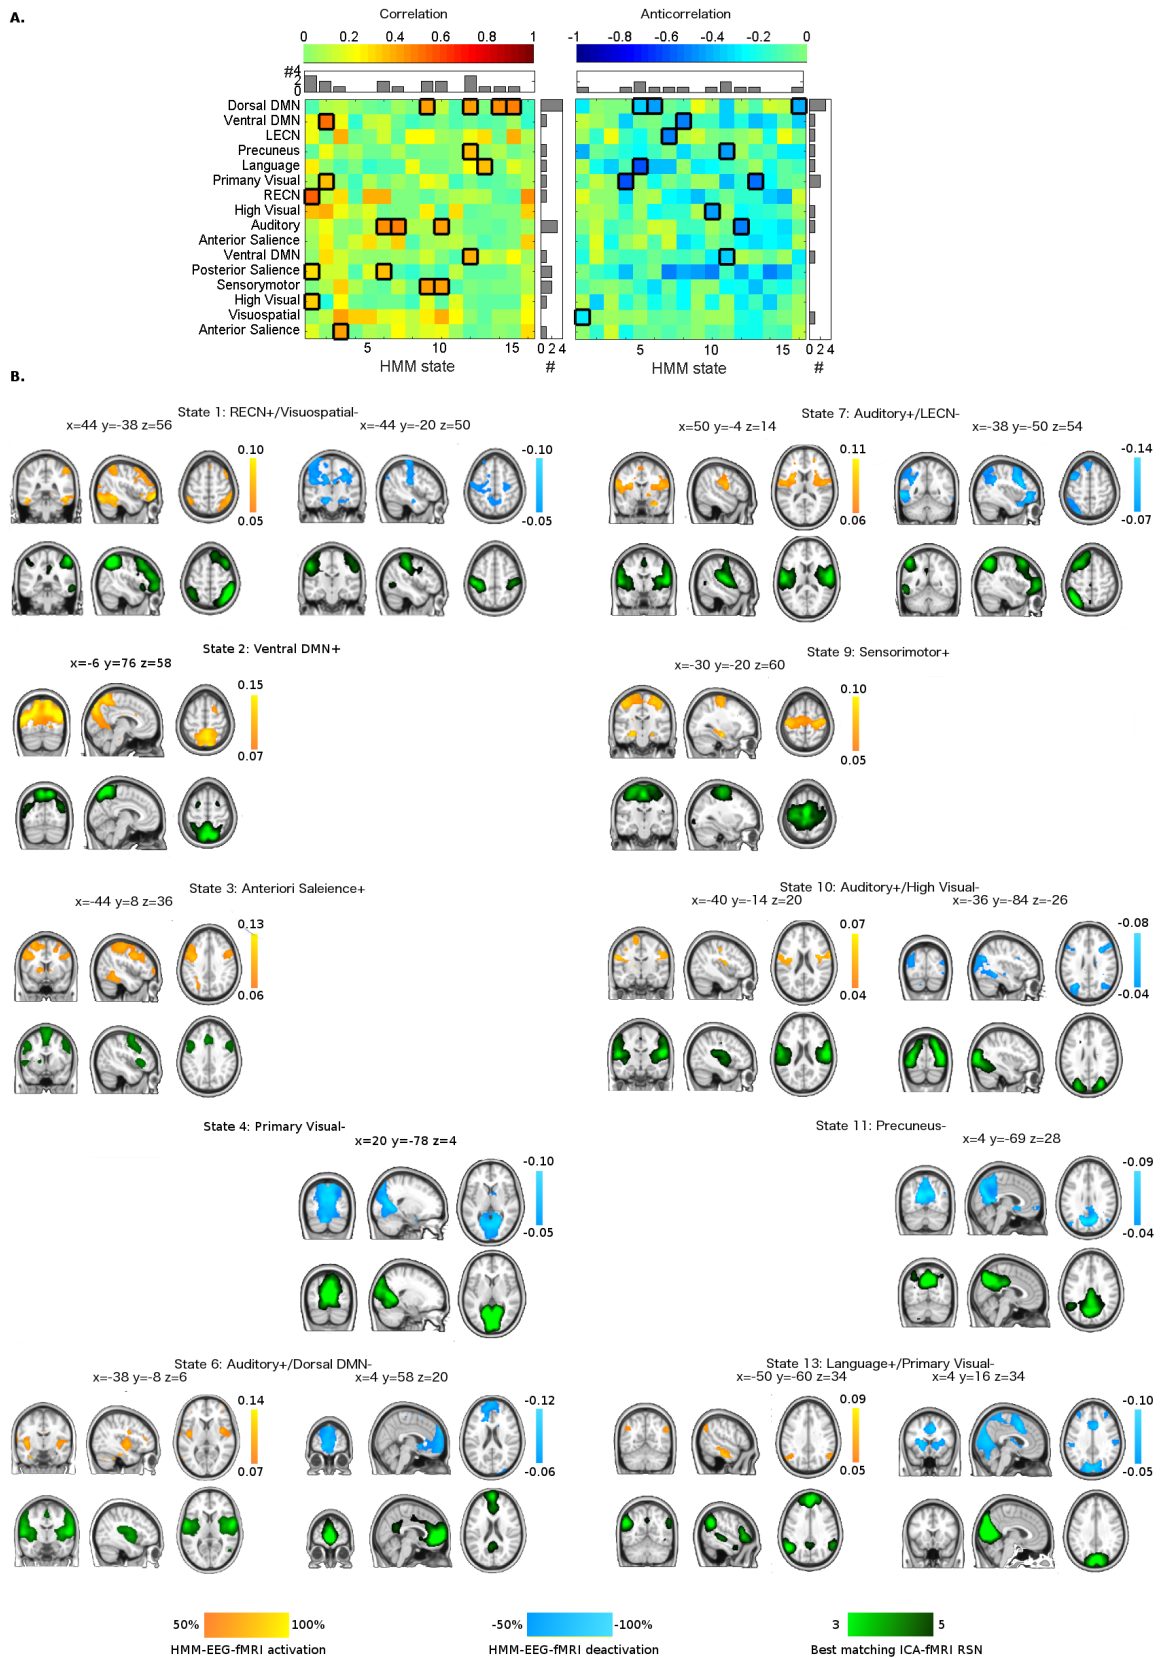


Figure 2: Network maps obtained by correlating fMRI data with EEG-derived, HRF-convolved brain state time-courses concatenated across subjects (HMM-EEG-fMRI) show similar spatial patterns as fMRI resting-state networks (ICA-fMRI). **A.** Spatial correlations (left) and anticorrelations (right) between unthresholded ICA-fMRI maps and unthresholded HMM-EEG-fMRI maps. Significant correlations are indicated with a black frame. Since ICA-fMRI maps are predominantly positive, positive spatial correlations always correspond to HMM-EEG-fMRI “activation” maps, and negative spatial correlations correspond to HMM-EEG-fMRI “deactivation” maps; where HMM-EEG-fMRI “activation” or “deactivation” reflects where

the HRF-convolved HMM-EEG state time-courses are positively or negatively temporally correlated with the fMRI data respectively. Bar plots above and next to the (anti)correlation matrices show the number of significant matches (i.e. spatial correlations) per HMM-EEG-fMRI map and ICA-fMRI map. B. Examples of the HMM-EEG-fMRI beta maps (red/yellow colour map for HMM-EEG-fMRI “activation”, and blue colour map for “deactivation”) are shown for a subset of the HMM brain states, together with their best matching ICA-fMRI RSNs (green colour map). Maps are not shown when no match was found between the HMM-EEG-fMRI and ICA-fMRI maps.

2. Reproducibility of HMM-EEG-fMRI maps

One remarkable property of RSN maps, obtained by ICA-fMRI, is their spatial stability across different runs, subjects or studies. We argue that the interpretation of the spatial organization of fast transient networks is only meaningful *if* the HMM-EEG-fMRI maps are robust as well. Therefore, we assessed the reproducibility of the HMM-EEG-fMRI maps across different datasets. Figure 3 shows the results of the split-half reproducibility analysis, computed based on 10 random splits of the original multi-subject data into 2 groups.

Reproducibility is assessed in terms of spatial correlation. We ordered the maps according to their reproducibility, and compared them with the reproducibility of ICA-fMRI and HMM-fMRI spatial maps. Although there is a tendency for ICA-fMRI maps to be more reproducible than HMM-EEG-fMRI maps, the difference does not reach significance (Wilcoxon rank sum test, corrected for multiple comparisons). Half of the ICA-fMRI maps are significantly more reproducible than the corresponding HMM-fMRI maps, however, there is no difference in reproducibility between HMM-EEG-fMRI and HMM-fMRI maps, adding to our confidence that the EEG-derived HMM state time courses are relevant to fMRI. The tendency for ICA-fMRI maps to be more reproducible is possibly due (at least partially) to the higher complexity of the HMM models, which, unlike ICA-fMRI, includes information of functional connectivity. Another factor contributing to the slightly reduced reproducibility of HMM-EEG-fMRI networks is the inherently weak correlation between EEG and fMRI.




Figure 3: Reproducibility of HMM-EEG-fMRI maps in comparison with ICA-fMRI RSN maps and HMM-fMRI maps computed on the same dataset. The maps are ordered based on their average reproducibility (indicated as solid lines) across the 10

random splits, for each method. Note, therefore, that the best reproducible ICA-fMRI map might correspond to a different network than the best reproducible HMM-EEG-fMRI map. Significant differences (Wilcoxon rank sum test, corrected for multiple comparisons with a false discover rate (FDR)<0.05) are marked above the plot with a red asterisk for HMM-fMRI vs ICA-fMRI. The difference between HMM-EEG-fMRI vs ICA-fMRI and HMM-EEG-fMRI or HMM-fMRI did not reach significance for any of the maps. The names of the networks, ordered according to their reproducibility, is listed below the figure.

3. A dynamic system of fast transient networks

Figure 34 shows the HMM-EEG state fractional occupancies (total time spent in a state relative to the total duration of the data) on the left and the distribution of HMM state life-times (time spent in a state before transitioning to a new state) on the right. The states are stable for 50-70ms on average, typically up to 125-200ms (95th percentile of life times), and the states are occupied 5-10% of the time, on average every 0.6-2s. One exception is state 6 (Auditory+, dorsalDMN-), which tends to be stable for longer periods of time (130ms on average, typically up to 400ms), but visited infrequently, with a fractional occupancy of 1%, on average every 13 seconds. We considered the possibility that this state represents an auditory response to scanner noise, or perhaps a residual artefact. However, the activation time course of the state did not correspond to the repetitive scanning sequence; neither did the state EEG topography (see Supplementary Figure 3) hint towards an artefact-related nature of this state.

The dynamic interaction between states can be interpreted based on the state transition probability matrix, shown on the left, and the correlations between the windowed fractional occupancy time courses, shown on the right. Notably, we observe strong transition probabilities between networks involved in different stages of visual processing, such as the visuospatial (state 1; fMRI-), the primary visual (state 4; fMRI-), and high visual (state 10; fMRI-) networks. Also, one can observe a tendency for transitions between various default mode and higher cognitive networks, such as from right executive control/visuospatial network (state 1 fMRI+/fMRI-) to ventral DMN (state 2; fMRI+), from left executive control network (state 7; fMRI-) to Precuneus (state 11; fMRI-), from dorsal DMN (state 15; fMRI+) to the anterior salience network (state 3; fMRI+), or from the dorsal DMN (state 16; fMRI-) to the ventral DMN (state 2; fMRI+).

Based on the most probable transitions from a certain network to the other, one can construct preferred processing pathways within the brain. Once such pathway could be, for instance, dorsal DMN (state 15, fMRI+) → anterior salience network (state 3, fMRI+) → from right executive control/visuospatial network (state 1 fMRI+/fMRI-) → ventral DMN (state 2; fMRI+).

Notably, the dorsal default mode network (DMN) is associated with a number of different HMM states, sometimes in anti-correlation with the auditory network (state 6, 12), but also as an individual activation or deactivation (state 14, 15 and 16, respectively). Similarly, the ventral DMN is associated with multiple states (state 2, and 8). These findings can be interpreted together by considering the fractional occupancy correlations of the networks, shown in the bottom right of Figure 3. For instance, states 15 and 16, which capture different anatomical regions within the default mode network (see Supplementary Figure 3), show highly correlated fractional occupancy time courses. This suggests a close functional integration of otherwise two distinct underlying networks. On the other hand, states 2 and 8, both within the ventral DMN, appear to be active in time periods when the other is silent. Notice that the right and left executive control network (ECN) seem to be functionally incompatible as well, but each are strongly related to either of the ventral DMN states: the

ECN (state 1; fMRI+) is active at times when state 2 is active, while the left ECN (state 7, fMRI-) is active when state 8 is active. This can be interpreted as a functional incompatibility or segregation of the networks.

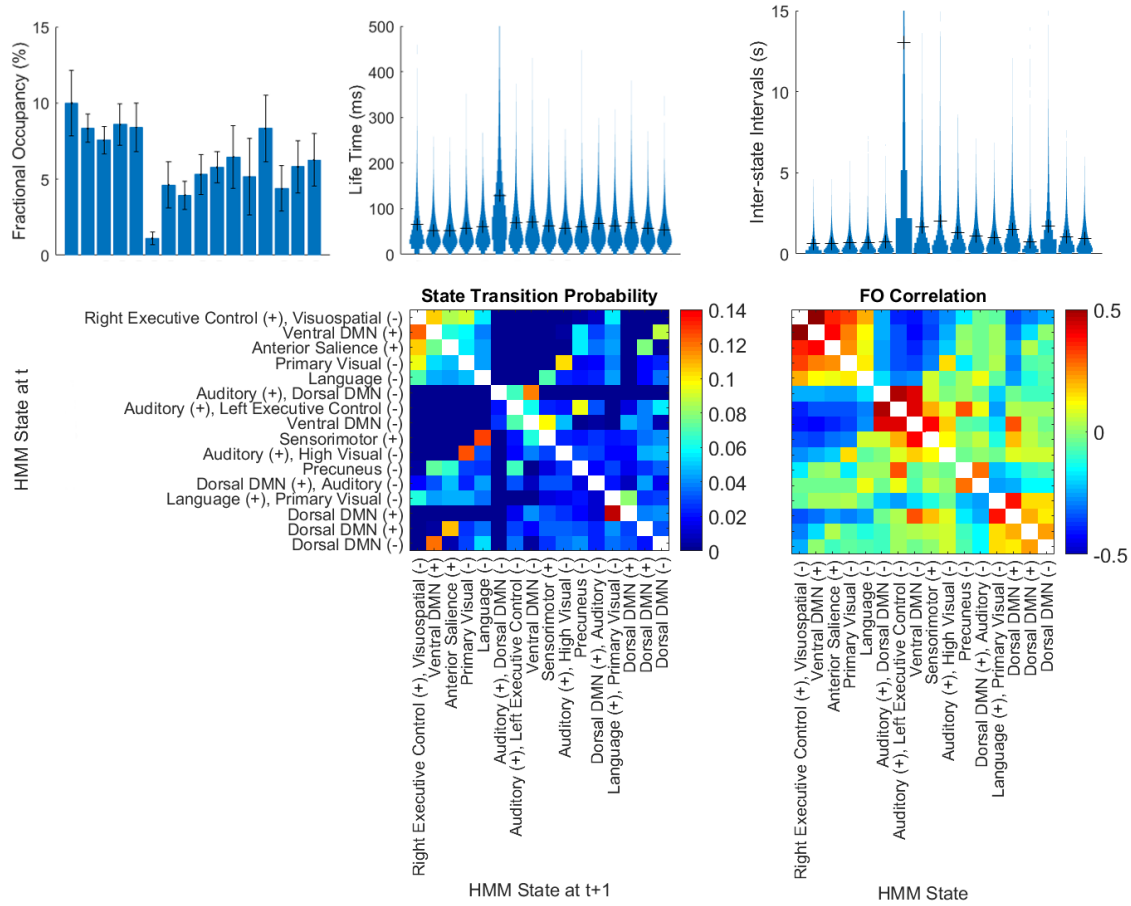


Figure 4 Temporal characteristics and dynamics of the HMM-based EEG networks. Top Left: Fractional occupancy of HMM states, with error bars indicating the standard deviation from the mean. Top Middle: Distribution of HMM state life times. Black crosses indicate the mean life time. Top right: Distribution of Inter-state Intervals. HMM-based EEG networks are ordered as indicated by the figure labels of the bottom panels. Bottom Left: state transition probability matrix. Self-transition probabilities are the highest for all states, but were omitted in this figure to emphasize the inter-relationships of states. Bottom Right: Correlation of the HMM state fractional occupancy time courses within 10s long sliding-window intervals.

DISCUSSION

Spontaneous brain activity at rest is organized into structured, recurring patterns. The exact structure and source of these patterns are currently subject to intense research [Florin2015, Vidaurre2017a, Vidaurre2017b, Vidaurre 2017c, Feige2017]. Electrophysiological recordings (e.g. EEG and MEG) allow the exploration of the fine temporal dynamics of brain activity, whereas fMRI captures slow fluctuations at high spatial detail. In this paper, we use

simultaneously recorded EEG and fMRI data to investigate resting state brain activity. Despite their fundamentally different time scale and nature of the measured signal and the assumptions behind the analysis methods, we find remarkable similarities between the fMRI correlates of fast transient EEG networks and fMRI resting state networks. Previous studies have reported correspondence between fMRI and EEG power fluctuation [Mantini2007], temporally independent (infra-)slow EEG fluctuations [Yuan2016, Hiltunen2014, Liu2015] and microstates [Britz2010, Yuan2012], suggesting a close link between these modalities. The scale-free behaviour of EEG microstates over a wide range of time-scales, including slow fluctuations characteristic to the BOLD signal, may provide an explanation to this link [VanDeVille2010]. In an attempt to theoretically explain the observable EEG microstate sequences, a stochastic model of the microstate process was proposed [Gartner2015]. In a more recent study [von Wegner2017], an information-theoretical analysis of microstate sequences suggested that microstates behave in a non-Markovian way, the transition rule between the states are non-stationary, yet, they are short-range correlated for large time lags. Hence, our framework is different: we propose a novel way of modelling EEG activity which conveys additional information on the properties of EEG data. First, unlike microstates the HMM formalises the features that make the states unique using a generative model. Second, the HMM state observation models make use not just activity levels, but of second order statistics (variance and correlations) in the HMM state observation model. As such, our HMM states contain functional connectivity information. Moreover, the inclusion of variances and covariances potentially allows us to separate states with similar mean activations but different covariances which microstates would consider to be the same. Third, the HMM explicitly models Markovian dependency, i.e. temporal interactions between states through the transition probability matrix, and is therefore tuned to finding states that repeat in a predictable way. The Markovian model helps constrain and regularise the detection of states, but importantly does not preclude the HMM states also showing long-range dependencies, which relate to scale-free behaviour. Future research efforts should aim to confirm long-range dependency emerges in HMM framework [Gschwind2015]. Notably, the predictable sequences revealed by HMM seem to correlate with behavioural traits (Vidaurre et al., 2017a). Finally, these predictable sequences are characterised by fast transitions (states are active for periods of ~50-100ms), as opposed to (infra-)slow power fluctuations (range of 2s and above) [Hiltunen2014] explored by other studies. In summary, HMM states are essentially different from microstates or temporally independent fluctuations. As a consequence, finding spatially meaningful fMRI correlates to these HMM states is also a novel finding.

We conducted extensive analyses to verify the robustness of our results for different numbers of inferred networks (states). Importantly, when analysing EEG and fMRI data individually, our results were in agreement with previous literature [Baker2014, Shirer2012, Damosieux2006, Beckmann2005]. More specifically: we extracted fMRI resting state networks using independent component analysis, showing an almost one-to-one correspondence with a previously known functional atlas [Shirer2012]; and we modelled EEG data using a HMM to reveal so-called fast transient networks, showing very similar (though slightly faster) dynamics as in [Baker2014].

Strikingly, we found large similarities with the results of [Baker2014] despite the important methodological differences. While Baker et al. modelled source-space MEG data, we here used sensor-space EEG envelopes, as our 64-channel recordings did not represent sufficient spatial variance for reliably mapping these to source space. By regressing the HMM state time courses on the MEG source-space envelopes, Baker et al. obtained spatial maps resembling

known resting state networks; similarly, we regressed the state time courses back on the EEG envelopes to produce EEG topographies. Although the interpretation of these topographies is not straightforward (as they are based on power time courses instead of dipolar time courses), in this paper we show that, by regressing on the simultaneously recorded fMRI data, we are able to obtain a high-resolution spatial interpretation of these fast transient networks.

The HMM-EEG-fMRI spatial maps obtained using this approach have an intimate anatomical correspondence to well-known fMRI resting state networks, showing significant correlation or anticorrelation with at least one RSN. Importantly, we do not observe a one-to-one correspondence. Although previous studies did establish one-to-one correspondence between networks revealed by different analysis methods [Hiltunen2014, Yuan et al., 2016 and Liu et al., 2017], this is not necessarily expected or desired. Instead, we argue that no one description is definitive, but the various approaches provide different and potentially all useful, complementary perspectives on the data. The main contribution of our HMM approach analysing EEG network is the temporal detail it provides, particularly in terms of the non-random transitioning between states and captured by the transition probability matrix in the HMM. We found, for example, that EEG networks which were associated with various “visual RSNs” (primary visual, high visual and visuospatial network) had high transition probabilities across each other. Note that similar interrelationship was also revealed with HMM-fMRI analysis, however, at a slower time scale (see Supplementary Figure 4.) Moreover, we could also construct preferred pathways within the brain by following the most probable transitions from one state to another. One such pathway starts from the dorsal DMN (state 15, fMRI+), then transitions to the anterior salience network (state 3, fMRI+), followed by the right executive control/visuospatial network (state 1 fMRI+/fMRI-). Note that state 15 and state 1 show antagonistic behaviour (see 10s fractional occupancy correlation matrix). This is in line with the well-known phenomenon that during times of cognitive load activity levels in the executive network increase while in the default mode network decrease and vice versa. It has been hypothesized that the anterior salience network drives this switch [Sridharan2008, Goulden2014]. Our result confirms this theory, providing electrophysiological evidence for the role of the ACN in the switching mechanism, at an appropriate sub-second time scale.

We observe that certain RSNs are associated with multiple fast transient EEG networks. We noted two different phenomena underlying this observation. First, transient HMM-EEG networks might separately capture regions that together belong in a single traditional RSN. For example, we found that the dorsal default mode network may split up to a posterior and an anterior cluster, consistent with the findings of recent studies [Vidaurre2017c]. Second, we observe that the fMRI correlates of two distinct EEG networks (exhibiting antagonistic temporal behaviour) capture most of the spatial information of a well-known RSN. One possible interpretation of this phenomenon is that the same RSN, while executing different tasks, may exhibit different patterns of electrophysiological activity. The presence of agonistic behaviour within certain groups of networks and antagonistic behaviour between these groups of networks altogether suggests the presence of distinct subsystems within the ensemble of brain networks. This is consistent with previous work, where, focusing on the temporal organization of fMRI resting state dynamics, [Vidaurre2017a] found a reproducible hierarchical structure within brain networks across subjects into two distinct high-level metastates.

The fMRI-ICA RSN maps, inferred using *spatial* ICA, predominantly show positive activations. By contrast, the fMRI temporal correlates of the fast transient HMM-EEG

networks (i.e. the HMM-EEG-fMRI spatial map), had considerably large regions of activations as well as deactivations. Similar observations were made in [Smith2012] using *temporal* ICA of fMRI. [Smith2012] interpreted this phenomenon as the presence of “considerable deactivation synchronized with activation” with respect to a certain brain function. Putting aside the modality and temporal resolution differences, this similarity between the HMM and temporal ICA, and their difference to spatial ICA, is possibly due to the relationship between the modelling assumptions made in the different computational approaches. More specifically, in temporal ICA the assumption of temporal independence will discourage temporal overlap, whereas spatial ICA will discourage spatial overlap. In the HMM, the state mutual exclusion assumption will prohibit temporal overlap, rendering it more similar to temporal ICA than spatial ICA. As a consequence, we might get a better correspondence between temporal ICA and HMM, despite the existing differences between them (i.e. HMM using second order statistics, Markovian temporal regularisation and accessing faster time scales). Unfortunately, the small amount and low sampling rate of fMRI data prohibits the use of temporal ICA. Therefore, comparing these two analysis approaches is an area for future work.

Due to the fact that we found no one-to-one correspondence between fMRI-ICA RSN maps and HMM-EEG-fMRI maps, these network assignments do not provide a precise anatomical description of the fast transient networks. Further research should be carried out to improve the description of these networks. A possible approach would be to apply ICA on fMRI using a high number of components to parcellate the brain into small segregated areas [Allen2014]. Then, the HMM-EEG state maps could be reconstructed onto such a localized dictionary. Such analysis will be an important step in follow-up studies, especially when describing HMM states specific to certain tasks (in a cognitive neuroscience experiment) or specific to certain neurological diseases.

One inherent limitation in many EEG-fMRI studies (including ours), is that one must make assumptions about the hemodynamic response linking the fast and transient electrophysiological events to the slow fMRI responses. Here, we use a single linear canonical HRF. However, some evidence suggests that the HRF can vary across individuals and brain regions [Aguirre1998, Miezin2000]. Also, it has been shown that a change in fMRI signal may be observed several seconds prior to an electrophysiological event. [Feige2017, de Munck2007]. While short time lags may be explained by HRF variability, delays of several seconds are in contradiction with the traditional belief that EEG-related neuronal activity causes the vascular response and concomitant increase in blood oxygenation. Instead, the observed phenomenon might be explained by the presence of preferred sequences of network activations, i.e. brain states [Feige2017]. Future work, using the proposed methodology and simultaneous EEG-fMRI methodology, will hopefully provide further evidence in these lines.

In summary, our method opens new opportunities in investigating the dynamical properties and interrelationships of functional brain networks. We believe that it can find applications in cognitive neuroscience studies to investigate processing pathways of complex tasks in the brain; as well as in clinical neuroscience, to explore the mechanisms behind certain neurological diseases, e.g. how the brain evolves from healthy to epileptic behaviour.

CONCLUSION

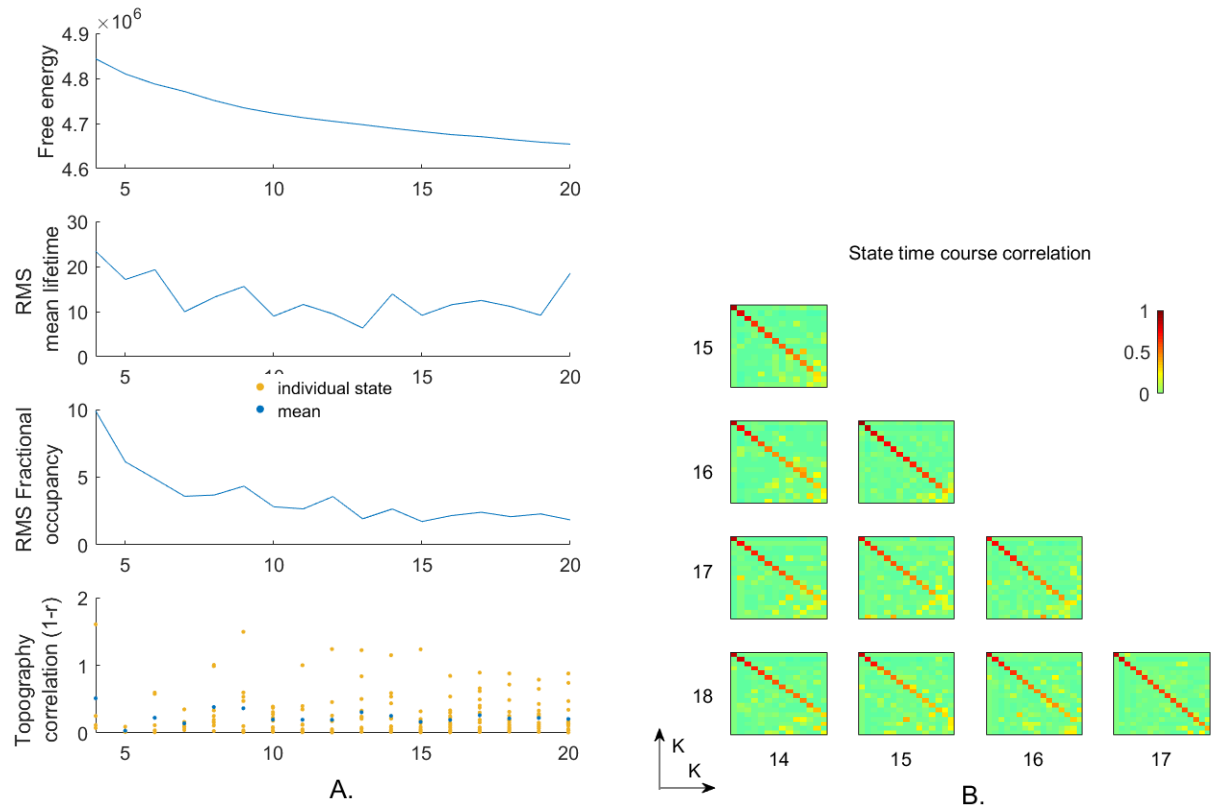
Using simultaneous EEG-fMRI recordings, we provide new evidence that fast, sub-second electrophysiological network activity underlies slow fMRI resting state network fluctuations.

In particular, we inferred a dynamic system of fast transient brain states from EEG data, which recurs in regulated sequences. Moreover, we gave spatial and functional interpretation to these states by showing that their HRF convolved correlates in the simultaneously acquired fMRI data are closely related to known fMRI resting state networks. This provides a more comprehensive spatiotemporal characterization of resting state brain function, spanning across different scales of spatial temporal detail.


Supplementary material

1. Model selection for HMM

The chosen number of states $K=16$ was determined as follows. First, we used the free energy of the model, which represents a lower bound of the so-called model evidence and represents a trade-off between model complexity and data likelihood (i.e. how well the model represents the data). A lower free energy value indicates a better (accurate but not too complex) model. Second, we considered different measures of robustness across different input datasets. This is done by splitting the original multisubject dataset randomly in two subgroups and fitting the HMM on both models separately. We would then choose a value of K such that the states have similar spatial and temporal properties in both groups. We measured spatial consistency as the average spatial correlation between the corresponding state EEG topographies. In order to inspect the consistency of temporal properties, we computed the root mean squared difference between the fractional occupancies and mean lifetimes across the split-half models. Supplementary Figure 1a depicts the different measures taken into account for model selection. Free energy shows a steady decrease, without reaching a plateau, suggesting the use of models with many states. Similarly, the consistency of fractional occupancy across the models increases with the number of states. However, there is no real trend for the consistency of mean life times or EEG topographies. Although combining these criteria does not give a clear preference, it does suggest that a higher number of states is beneficial. Finally, we chose a HMM with $K=16$ for further investigation although different choices of K could be reasonable as well. In order to make sure that this fairly arbitrary choice does not have a large influence on the results, we also inspected the stability of the model against the perturbation of K . We fitted an HMM model on the original multisubject dataset with different K , ranging from 12 to 20. Then, we compared the state time courses obtained by the different models. Supplementary Figure 1b shows the results of this analysis, in terms of the correlation of the best matching pairs of states across the models. Note that we reordered the states such that the corresponding state appear in the same order in both models and the ones with the highest correlation appear first. In general, large correlation values appear almost along the full length of the diagonals, and very small values appear elsewhere, except for a few entries. This indicates a good one-to-one correspondence of states across models with different K . For example, comparing models with $K=15$ and $K=16$, 14 states appear in both models, as indicated by the 14 large values in the diagonal. The last few states show small to moderate correlations across each other, which might indicate that these states are combined together when reducing the number of states from 16 to 15. In summary, these cross-state correlations indicate that the results are reasonably stable with respect to the choice of the number of states K .

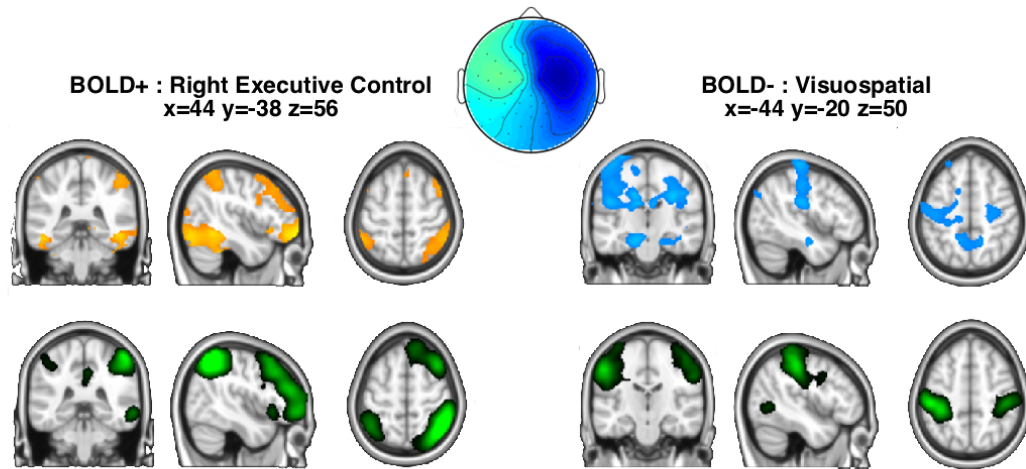


Supplementary Figure 1: A. Model selection based on free energy, and consistency of spatial and temporal properties of states across HMM models fitted on 2 randomly split subgroups of subjects. Consistency measures include root mean square difference of mean lifetime and fractional occupancy across models and correlation of state topographies. B. Correlation between matched state time courses across models with different number of states K .

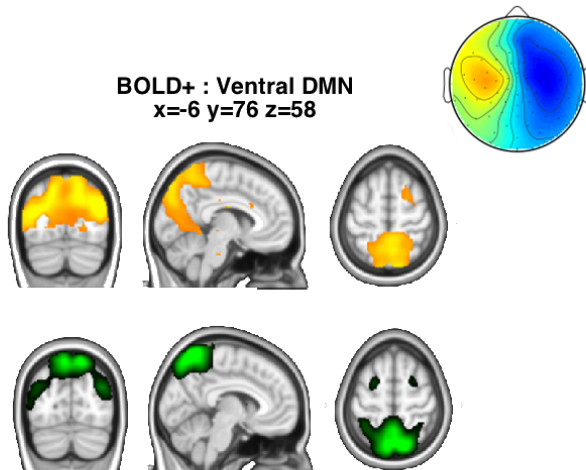


Supplementary Figure 2: Matching resting state networks maps derived from our multisubject fMRI dataset using groupICA to the networks from the functional ROI atlas in [Shirer2012] based on spatial correlation. There is a near one-to-one correspondence. We assign the name and interpretation to each ICA-based RSN map according to its best matching network from the atlas. The correlation values with the best matching network is indicated in brackets after the assigned names. A few exceptions include 5 RSNs, which, after visual inspection, were interpreted as follows: RSN 18 – posterior salience network; RSN 15, 17, 19 and 20 – artifact.


HMM state 1
Global EEG power -




HMM state 2
Global EEG power -




HMM state 3
Global EEG power -




HMM state 4
Global EEG power -




BOLD- : Primary Visual
 $x=20$ $y=-78$ $z=4$




HMM state 5
Global EEG power -




BOLD- : Language
 $x=-6$ $y=54$ $z=28$




HMM state 6
Global EEG power +



BOLD+ : Auditory
 $x=-38$ $y=-8$ $z=6$



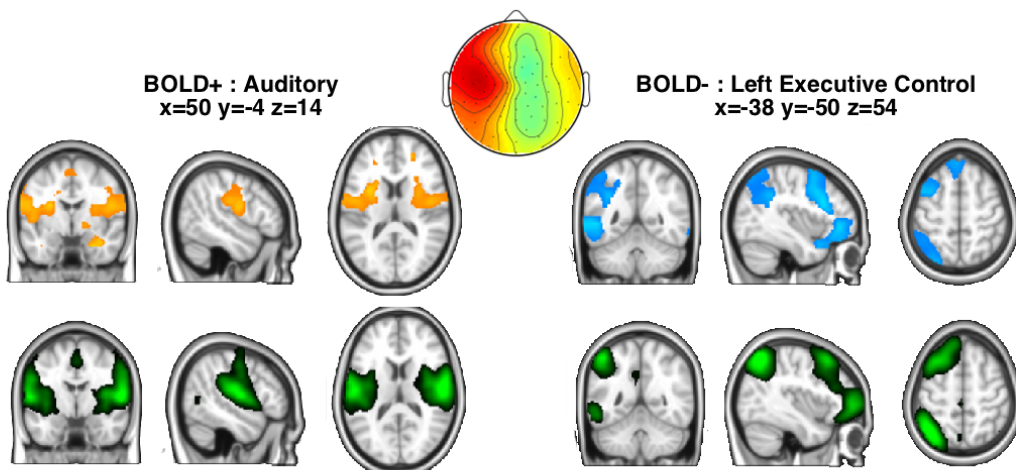
BOLD- : Dorsal DMN
 $x=4$ $y=58$ $z=20$



HMM state 7
Global EEG power +


BOLD+ : Auditory
x=50 y=-4 z=14

BOLD- : Left Executive Control
x=-38 y=-50 z=54



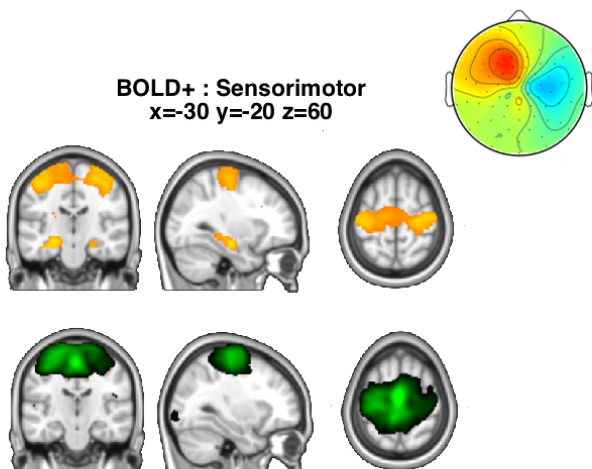
HMM state 8
Global EEG power +

BOLD- : Ventral DMN
x=-14 y=-74 z=58

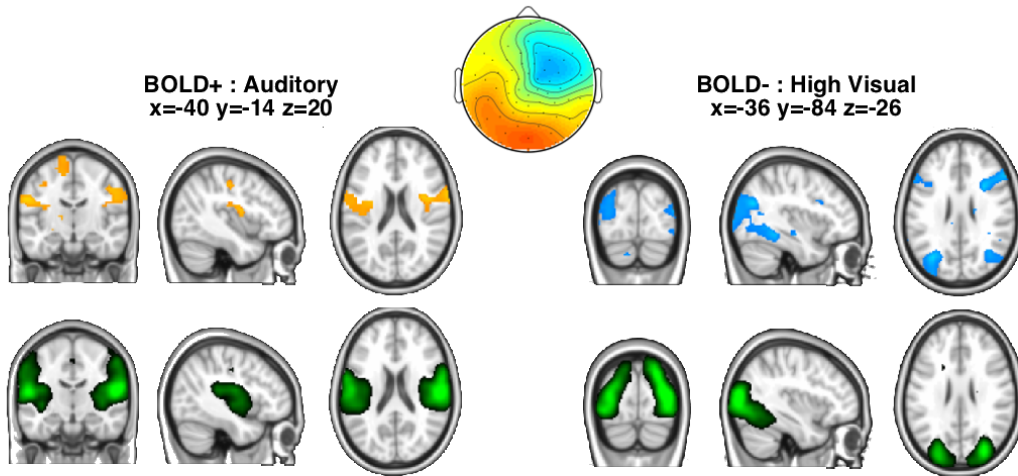


HMM state 9
Global EEG power +

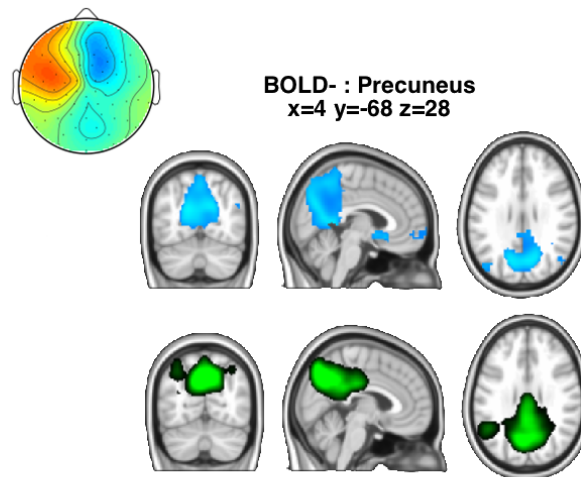
BOLD+ : Sensorimotor
x=30 y=-20 z=60



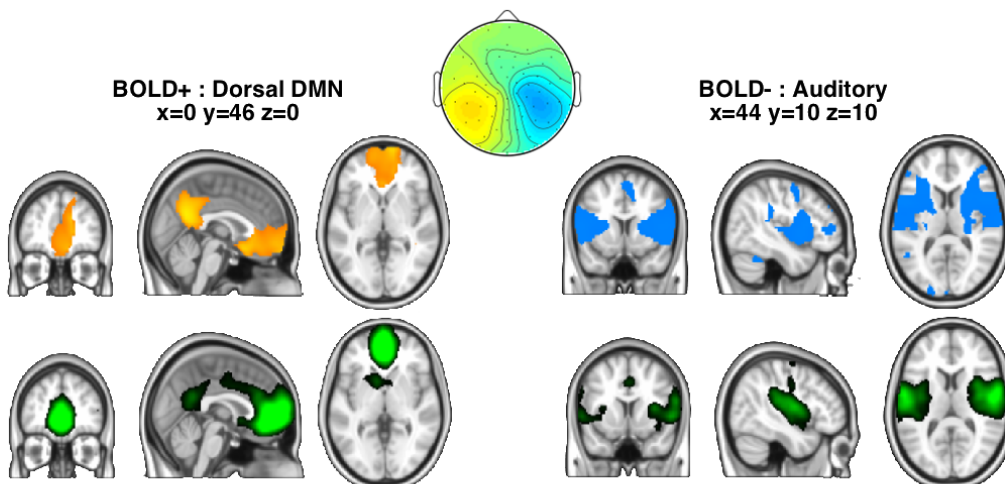
HMM state 10
Global EEG power +



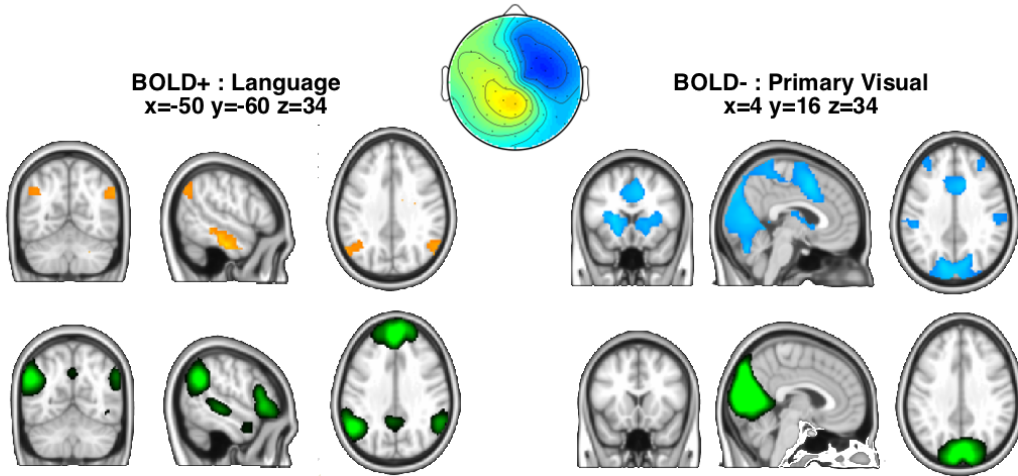
HMM state 11
Global EEG power +




HMM state 12
Global EEG power -




HMM state 13
Global EEG power -




HMM state 14
Global EEG power +

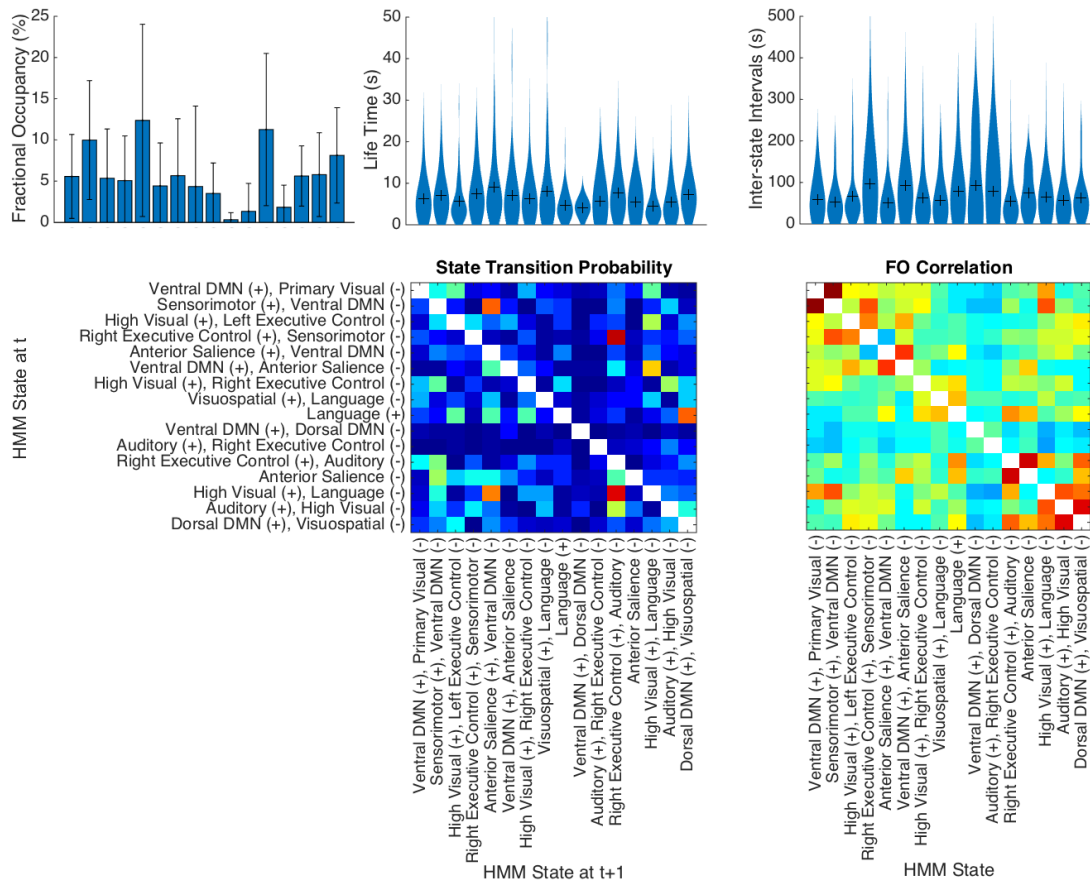


HMM state 15
Global EEG power +






Supplementary Figure 3: Spatial characterization of the HMM-EEG states. EEG topographies (colorbar is the same for all, shown for state 16), HMM-EEG-fMRI maps (orange: activations, blue: deactivations) and the corresponding best matching group-ICA maps (green) are shown for each state. We also indicate whether each HMM state is active at times the global EEG power increases (+) or decreases (-). Significance thresholded HMM-EEG-fMRI beta maps are shown in all cases except for the deactivations of state 10 and 13, where thresholding did not result in clear maps. Activations, indicated as ‘BOLD+’ are on the left, while deactivations, indicated as ‘BOLD-’ are on the right. The colours range between 50-100% of the largest voxel value for positive voxels and 50-100% of the smallest voxel value for negative voxels.




Supplementary Figure 4: Temporal characteristics and dynamics of the HMM-based fMRI networks. The same notation and visualization principles are used as for HMM-based EEG networks on Figure 2. In comparison with HMM-EEG networks, the

life times and interstate-intervals of HMM-fMRI states are longer (5-10s long and 50-100s long, on average, respectively). State fractional occupancies are in similar range (up to 12%) as in HMM-EEG (5-10%). The state transition probability matrix shows a similar structure as for HMM-EEG, namely, transition probabilities are in general low, except for a few selected transitions. This could mean that there are certain preferred processing pathways in the brain. Our previous finding, that one such preferred pathway follows the different visual networks is confirmed with HMM-fMRI as well. Consistently with HMM-EEG, a modular structure of the fractional occupancy matrix can be observed (although less clear), which suggests the presence of different subsystems in the brain. Regarding spatial organization, both HMM-EEG-fMRI and HMM-fMRI have both activations and deactivations, and multiple matches with RSNs (not shown). There is no one-to-one correspondence between HMM-EEG and HMM-EEG-fMRI maps either. In summary, the main observations about the brain's dynamic structure are similar using EEG and fMRI; but the same phenomena can be observed at different time-scales. These findings are in line with the notion of scale-free behavior.


State1: Ventral DMN (+), Primary Visual (-)




State 2: 'Sensorimotor (+), Ventral DMN (-)'




State 3: 'High Visual (+), Left Executive Control (-)'




State 4: 'Right Executive Control (+), Sensorimotor (-)'




State 5: 'Anterior Salience (+), Ventral DMN (-)'




State 6: 'Ventral DMN (+), Anterior Salience (-)'




State 7: 'High Visual (+), Right Executive Control (-)'




State 8: 'Visuospatial (+), Language (-)'




State 9: 'Language (+)'




State 12: 'Right Executive Control (+)'




State 13: 'Anterior Salience (-)'




State 14: 'High Visual (+), Language (-)'



State 15: 'Auditory (+), High Visual (-)'



State 16: 'Dorsal DMN (+), Visuospatial (-)'S



Supplementary Figure 5: Spatial characterization of the HMM-fMRI states. Significance thresholded HMM-fMRI beta maps are shown in all cases except for the activations of state 2 and deactivations of 8, where thresholding did not result in clear maps. States 10 and 11 did not have any significant activations or deactivations (perhaps due to the fact there they were not extracted in all subjects), therefore, are not shown. Deactivation of state 12 did not look similar visually to the Auditory network (despite of the matching procedure indicating it as a match), therefore, it is not shown. Activations, indicated as

'BOLD+' are on the left, while deactivations, indicated as 'BOLD-' are on the right. The colours range between 50-100% of the largest voxel value for positive voxels and 50-100% of the smallest voxel value for negative voxels.

Acknowledgements

We thank Sebastian Puschmann for kindly providing the EEG-fMRI dataset.

BH was supported by imec funds 2017, Fonds voor Wetenschappelijk Onderzoek-Vlaanderen (FWO) and by the European Research Council: The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Advanced Grant: BIOTENSORS (n° 339804). This paper reflects only the authors' views and the Union is not liable for any use that may be made of the contained information. The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust (203139/Z/16/Z). MWW's research is supported by the NIHR Oxford Health Biomedical Research Centre, by the Wellcome Trust (106183/Z/14/Z), and the MRC UK MEG Partnership Grant (MR/K005464/1). MDV's research is supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC).

References

Aguirre, G.K., Zarahn, E., D'Esposito, M. The variability of human, BOLD hemodynamic responses. *Neuroimage* 1998;8(4):360–369.

Allen, P., Josephs, O., Turner, R. A method for removing imaging artifact from continuous EEG recording during functional MRI. *Neuroimage* 2000;12:230–239.

Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T. and Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. *Cerebral cortex*, 24(3):.663-676.

Baker, A., Brookes, M., Rezek, I., Smith, S., Behrens, T., Smith, P., Woolrich, M.W. Fast transient networks in spontaneous human brain activity. *Elife* 2014;3:p.e01867.

Beckmann, C., Smith, S. Tensorial extensions of independent component analysis for multisubject fMRI analysis. *Neuroimage* 2005;25(1):294 – 311.

Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. Springer New York (New York, USA).

Britz, J., Van De Ville, D., Michel, C. Bold correlates of EEG topography reveal rapid resting-state network dynamics. *Neuroimage* 2010; 52(4):1162–1170.

Brookes, M. J., Woolrich, M., Luckhoo, H., Price, D., Hale, J. R., Stephenson, M. C., et al. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences*, 2011; 108(40), 16783–16788

Brookes, M. J., Woolrich, M. W., & Barnes, G. R. Measuring functional connectivity in MEG: a multivariate approach insensitive to linear source leakage. *Neuroimage* 2012; 63(2), 910–920.

Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S., Beckmann, C. Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences* 2006; 103(37):13848–13853.

de Munck, J.C. , Gonçalves, S.I., Huijboom, L., Kuijer, J.P., Pouwels, P.J., Heethaar, R.M., Lopes da Silva, F.H. The hemodynamic response of the alpha rhythm: an EEG/fMRI study. *NeuroImage* 2007, 35 pp. 1142-1151

De Pasquale, F., Della Penna, S., Snyder, A., Lewis, C., Mantini, D., Marzetti, L., Belardinelli, P., Ciancetta, L., Pizzella, V., Romani, G., Corbetta, M. Temporal dynamics of spontaneous MEG activity in brain networks. *Proceedings of the National Academy of Sciences* 2010;107(13):6040–6045.

Delorme, A., Makeig, S. EEGLab: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 2004;134:921.

Feige, B., Spiegelhalder, K., Kiemen, A., Bosch, O., van Elst, L., Hennig, J., Seifritz, E., Riemann, D.. Distinctive time-lagged resting-state networks revealed by simultaneous EEG-fMRI. *NeuroImage* 2017;145:1–10.

Florin, E., Watanabe, M. and Logothetis, N.K., The role of sub-second neural events in spontaneous brain activity. *Current opinion in neurobiology* 2015, 32, pp.24-30.

Gärtner, M., Brodbeck, V., Laufs, H. and Schneider, G.. "A stochastic model for EEG microstate sequence analysis." *NeuroImage* 104 (2015): 199-208.

Goldman, R.I., Stern, J.M., Engel Jr., J., Cohen, M.S., Simultaneous EEG and fMRI of the alpha rhythm. *Neuroreport* 2002; 13:2487-2492.

Goulden, N., Khusnulina, A., Davis, N.J., Bracewell, R.M., Bokde, A.L., McNulty, J.P. and Mullins, P.G. The salience network is responsible for switching between the default mode network and the central executive network: replication from DCM. *Neuroimage* 2014; 99:180-190.

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., et al. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* 2014;95:232–247.

Gschwind, M., Michel, C.M. and Van De Ville, D., 2015. Long-range dependencies make the difference—Comment on “A stochastic model for EEG microstate sequence analysis”. *NeuroImage*, 2015;117:449-455.

Hiltunen, T., Kantola, J., Elseoud, A., Lepola, P., Suominen, K., Starck, T., Nikkinen, J., Remes, J., Tervonen, O., Palva, S., Kiviniemi, V. Infra-slow EEG fluctuations are correlated with resting-state network dynamics in fMRI. *Journal of Neuroscience* 2014;34(2):356–362.

Hyvarinen, A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks* 1999;10(3):626–634.

Jann, K., Dierks, T., Boesch, C., Kottlow, M., Strik, W., Koenig, T. BOLD correlates of EEG alpha phase-locking and the fMRI default mode network. *Neuroimage* 2009; 45:903-916,

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

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

Koenig, T., Prichep, L., Lehmann, D., Sosa, P.V., Braeker, E., Kleinlogel, H., Isenhardt, R. and John, E.R. Millisecond by millisecond, year by year: normative EEG microstates and developmental stages. *Neuroimage*, (2002) 16(1), pp.41-48.

Lehmann, Dietrich, Roberto D. Pascual-Marqui, and Christoph Michel. "EEG microstates." *Scholarpedia* 4, no. 3 (2009): 7632.

Liu, Q., Farahibozorg, S., Porcaro, C., Wenderoth, N., Mantini, D. Detecting large-scale networks in the human brain using high-density electroencephalography. *Human brain mapping* 2017;38(9):4631–4643.

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

Miezin, F.M., Maccotta, L., Ollinger, J., Petersen, S., Buckner, R. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *Neuroimage* 2000;11(6):735–759.

Niazy, R., Beckmann, C., Iannetti, G., Brady, J., Smith, S. Removal of fMRI environment artifacts from EEG data using optimal basis sets. *NeuroImage* 2005;28(3):720–737.

Rukat, T., Baker, A., Quinn, A. and Woolrich, M. Resting state brain networks from EEG: Hidden Markov states vs. classical microstates. *arXiv preprint* 2016; arXiv:1606.02344.

Puschmann, S., Huster, R.J., Thiel, C.M. Mapping the spatiotemporal dynamics of processing task-relevant and task-irrelevant sound feature changes using concurrent EEG-fMRI. *Human Brain Mapping* 2016;37(10):3400-3416.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C., Glasser, M., L., G., Smith, S. Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage* 2014;90:449–68.

Sridharan, D., Levitin, D.J. and Menon, V. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proceedings of the National Academy of Sciences* 2008; 105(34): 12569-12574.

Shirer, W., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral cortex* 2012;22(1):158–165.

Smith, S. M., Fox, P. T., Miller, Karla L., Glahn, David C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A.R. and Beckmann, C.F. Correspondence of

the brain's functional architecture during activation and rest. . Proceedings of the National Academy of Sciences 2009 ;106 (31): 13040-13045

Smith, S., Miller, K., Moeller, S., Xu, J., Auerbach, E., Woolrich, M.W., Beckmann, C., Jenkinson, M., Andersson, J., Glasser, M., Van Essen, D. Temporally-independent functional modes of spontaneous brain activity. proceedings of the national academy of sciences. Proceedings of the National Academy of Sciences 2012;109(8):3131–3136.

Yuan H, Zotev V, Phillips R, Drevets WC, Bodurka J. Spatiotemporal dynamics of the brain at rest—exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. Neuroimage 2012; 60:2062-2072.

Yuan H, Ding L, Zhu M, Zotev V, Phillips R, Bodurka J. Reconstructing Large-Scale Brain Resting-State Networks from High-Resolution EEG: Spatial and Temporal Comparisons with fMRI. Brain Connect. 2016;6(2):122-35.

Van Dijk, K.R., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L. Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. Journal of neurophysiology 2010;103(1):297–321.

Vanderperren, K., De Vos, M., Ramautar, J., Novitskiy, N., Mennes, M., Asseconci, S., Vanrumste, B., Stiers, P., Van den Bergh, B., Wagemans, J., Lagae, L., Sunaert, S., Van Huffel, S.. Removal of BCG artifacts from EEG recordings inside the MR scanner: A comparison of methodological and validation-related aspects. NeuroImage 2010;50(3):920 – 934.

Van de Ville, D., Britz, J. and Michel, C.M. EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. Proceedings of the National Academy of Sciences 2010;107(42), pp.18179-18184.

Vidaurre, D., Smith, S.M., Woolrich, M.W. Brain network dynamics are hierarchically organised in time. Proceedings of the National Academy of Sciences 2017a, 114(48), 12827–12832.

Vidaurre, D., Quinn, A. J., Baker, A. P., Dupret, D., Tejero-Cantero, A., & Woolrich, M. W. Spectrally resolved fast transient brain states in electrophysiological data., Neuroimage 2016; 126, 81–95.

Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A. J., Alfaro-Almagro, F., Smith, S. M., & Woolrich, M. W. Discovering dynamic brain networks from big data in rest and task. NeuroImage 2017b, In press.

Vidaurre, D., Hunt, L.T., Quinn, A.J., Hunt, B.A.E., Brookes, M.J., Nobre A.C. & Woolrich, M.W. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. BiorXiv 2017c.