

Can fMRI functional connectivity index dynamic neural communication?

Sonsoles Alonso Martínez^a, Alberto Llera^b, Gert ter Horst^a and Diego Vidaurre^{c,d}

^a Cognitive Neuroscience Center, Department of Biomedical Sciences of Cells & Systems, University Medical Center Groningen, University of Groningen, 9713 Netherlands

^b Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 Netherlands

^c Center for Functionally Integrative Neuroscience, Department of Clinical Medicine, Aarhus University, 8000 Denmark

^d Department of Psychiatry, University of Oxford, OX3 7JX UK

Abstract: In order to continuously respond to a changing environment and support self-generating cognition and behaviour, neural communication must be highly flexible and dynamic at the same time than hierarchically organized. While whole-brain fMRI measures have revealed robust yet changing patterns of statistical dependencies between regions, it is not clear whether these statistical patterns —referred to as functional connectivity— can reflect dynamic large-scale communication in a way that is relevant to cognition. For functional connectivity to reflect actual communication, we propose three necessary conditions: it must span sufficient temporal complexity to support the needs of cognition while still being highly organized so that the system behaves reliably; it must be able to adapt to the current behavioural context; and it must exhibit fluctuations at sufficiently short timescales. In this paper, we introduce principal components of connectivity analysis (PCCA), an approach based on running principal component analysis on multiple runs of a time-varying functional connectivity model to show that functional connectivity follows low- yet multi-dimensional trajectories that can be reliably measured, and that these trajectories meet the aforementioned criteria to index flexible communication between neural populations and support moment-to-moment cognition.

Introduction

The widespread existence of weak correlations between neurons in the brain is a well-known observed phenomenon (Averbeck et al., 2006; Averbeck and Lee, 2004; Cohen and Kohn, 2011; Nirenberg and Latham, 2003). An important question in theoretical neuroscience is about the mechanistic role of these correlations; particularly, whether or not these correlations convey information by themselves, i.e. as a separate “channel” above and beyond the firing rates and precise temporal ordering of the neuron’s firing events. If that were the case, then these correlations should be able to modulate their magnitude and configuration dynamically according to the external circumstances and the internal state of the animal (Cohen and Kohn, 2011; Nienborg and Cumming, 2009). Here, we ask a similar question at the macroscopic level using functional magnetic resonance imaging (fMRI) from humans, which, given its spatial coverage, might be better suited than the microscopic level to index higher-level aspects of cognition (Dehaene and Naccache, 2001). It is known that

patterns of voxel activations can encode for a variety of task-related variables (Hasson et al., 2009) and cognitive states (Haynes and Rees, 2006; Shine et al., 2019a). Furthermore, it has been shown that session-average patterns of covariance (or functional connectivity, FC) between regions are phylogenetically conserved (Lu et al., 2012), map the functional specialization of the different brain areas (Margulies et al., 2016), and can reflect clinical (Fox and Greicius, 2010) and psychological variations across the population (Finn et al., 2015; Smith et al., 2015). The question here is whether these (second-order) patterns of covariation can reflect actual long-range communication, and therefore relate to the current cognitive state of the subject and its perpetual, self-organized change —i.e. above and beyond the (first-order) voxel activation patterns. We reasoned that at least three necessary conditions must be met: (i) covariances across regions should exhibit temporal modulations of enough complexity to be able to support the flexibility of long-range neural communication and cognition; (ii) the nature of these modulations must be able to adapt to the behavioural context; and (iii) these modulations should span reasonably short timescales (compatible for example with the timescales of a thought process).

To find whether these conditions are met in real data, we introduce principal components of connectivity analysis (PCCA), an approach that combines multiple hidden Markov model (HMM) estimates from randomized initializations, and principal component analysis (PCA) to extract the latent temporal structure across HMM runs. PCCA revealed a multidimensional axis of temporal covariation in macroscopic FC that fits all the above requirements. Along this axis, robust trajectories of cross-region *covariation* unfold reliably across several dimensions (first condition) and at relatively short timescales (third condition). While the spatial connectivity patterns of these dimensions generalise across the population of subjects, the estimated temporal trajectories were subject-specific and were strongly modulated by the behavioural condition of the experiment (second condition). These results suggest that the estimated low-dimensional FC trajectories can not only be measured reliably, but suit the requirements to index dynamic interregional communication.

Results

We analysed fMRI data (repetition time, TR = 0.72s) from the Human Connectome Project (HCP; van Essen et al., 2013) from 100 subjects across three behavioural conditions: resting-state (rest), working memory (WM) task, and motor task. We developed PCCA, a new approach based on PCA and multiple runs of the HMM to identify reliable patterns of time-varying FC (**Fig. 1**). We used the resting-state data projected onto 25 independent components (ICs), and applied the group 25-IC parcellation from this dataset to extract the time series of the WM and motor task datasets (**Fig. 1A**). Then, we used the HMM to represent the data as a series of state activations, where each state represents a pattern of FC (Vidaurre et al., 2017); here, the HMM was set up to have $K = 12$ states (**Fig. 1B**). Note that the HMM is a probabilistic model whose inference estimates both the state activation probabilities (i.e.

the temporal parameters) and the state parameters (i.e. the spatial parameters) simultaneously and in a data-driven way. Critically, such parameter inference depends on an optimization procedure that typically starts from a random initialization. This means that different runs of the algorithm may result in somewhat different results. Here, we used this variability to our advantage by running the algorithm multiple times (on each behavioural condition separately) and performing PCA on the resulting state time series (**Fig. 1C**). By doing so, we reduced the dimensionality of the data into a small set of orthogonal principal components (PCs) that capture the dominant fluctuations in FC, where each PC has an associated covariance matrix. The PC-specific covariance matrices were computed as weighted averages of the HMM state covariance matrices (**Fig. 1D**); each weight, β , (one per HMM state and PC component) was calculated as the variance of the corresponding PC time series during the time that the corresponding HMM state was active (see β calculation in **Fig. 1D** and Methods for more details).

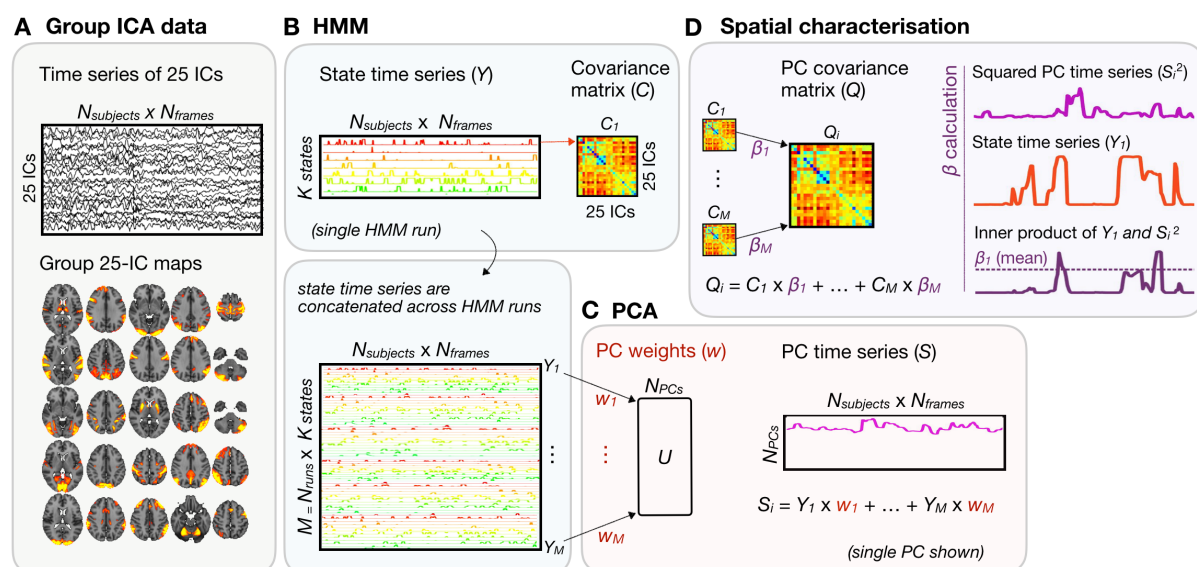


Fig. 1 Graphical summary of PCCA. **A** processed BOLD time series of 25 ICs obtained from 100 subjects of the HCP. **B** HMM was applied to the concatenated BOLD time series across subjects (25 ICs by $[N_{\text{subjects}} \times N_{\text{frames}}]$) resulting in $K = 12$ states, each with associated state time series (i.e., the probability of a given state to be active or not at every time point) and covariance matrix C_m across 25 ICs. The HMM was run several times (N_{runs}), each starting from a different random initialisation. **C** PCA was applied to the concatenated state time series across runs ($[N_{\text{runs}} \times K]$ by $[N_{\text{subjects}} \times N_{\text{frames}}]$), projecting the $N_{\text{runs}} \times K$ rows into a reduced set of PCs explaining maximum variance. **D** The covariance matrix of each PC was estimated as a linear combination of the covariance matrices of the original states, using the betas of each state (β_m) as their weights. The betas were computed as the weighted average of each original state time series with the square time series of a given PC. HMM: hidden Markov model(ling); ICs: independent components; PCA: principal component analysis. PCCA: principal components of connectivity analysis.

In what follows, we show that time-varying FC follows reliable trajectories that meet the aforementioned criteria. First, we show that such trajectories are low- but multi-dimensional. Second, we demonstrate that the nature of these trajectories differs across behavioural

conditions. Third, we show that FC exhibits changes within reasonably short timescales that are arguably compatible with the timescale of cognition and behaviour. In addition, we confirm that FC trajectories are not determined just by amplitude modulations. Finally, we verify the reproducibility of our results on new model configurations and a separate set of subjects.

Functional connectivity reliably embeds into multidimensional temporal trajectories

We used our PCCA approach to capture modulations in brain network activity, in the sense of changing patterns of covariance across regions. By running the HMM multiple times, we leveraged the variability of the HMM inference by refactoring the different estimations into reliable, latent trajectories of time-varying FC. By doing so, we explored the complexity of these trajectories, finding that they robustly span multiple dimensions.

First, to quantify the stability of the separate instances of the HMM inference, we computed a Pearson correlation between the state time series generated across multiple runs. Because the ordering of the states within a run is arbitrary, the states were aligned between runs using the Hungarian algorithm (Munkres, 1957), so that the first state in one run corresponds to the first state in another run, and so forth. We obtained correlation coefficients between every pair of aligned states for each pair of runs. **Fig. 2A** confirmed that the states inferred by the HMM were (to some extent) variable between the different runs of the inference (Vidaurre et al., 2019). Note that, because of the nature of the Hungarian algorithm, the states were ordered from best to worst aligned, therefore creating a natural ordering of the states from most to least replicable.

As illustrated in **Fig 1**, we exploited this variability by applying PCA across the state time series from 200 HMM runs. This produced a set of orthogonal PCs that optimally explained maximum time-varying FC variability across states and runs (i.e., 2400 states obtained from 200 runs of the model, with 12 states per run). As shown in **Fig. 2B**, the top five PCs alone retained 44% of the total temporal variance present across all of the states, suggesting that the temporal network variability in the data has more than five dimensions. Focusing on these five dimensions, we assessed the reliability of these PCs across 30 different PCA runs, where each PCA run was performed on 200 HMM runs from the original pool of 1000 HMM runs. We then estimated the reliability of each PC by computing the Pearson correlation between the PC time series of each pair of PCA repetitions. The average coefficient across all pairs was very close to $r = 1.0$ for PC1 and PC2, and $r = 0.99$ for PC3, PC4 and PC5, confirming that our approach is extremely stable against estimation noise (of which other approaches suffer systematically: Hindriks et al., 2016; Vidaurre et al., 2019). To quantify the relation between the number of HMM runs (N_{runs}) and the stability of the results, we repeated this process for various choices of N_{runs} . **Fig. 2C** shows that the set of states generated over 30 PCA runs was already able to make the five PCs highly stable, ($r > 0.92$). While this stability decreased for lower-order PCs, adding more HMM runs did increase the number of stable PCs. Additionally, since the aim of applying PCA was to integrate information across HMM runs, we verified that

the PCs were not representations of specific runs. As observed in **Fig. 2D** the average state contribution of each HMM run to the PCs are similar across runs (particularly concerning their contributions to higher-order PCs), demonstrating that the PCs effectively represent a low-dimensional embedding across all runs.

Similar conclusions were drawn from the analysis of task fMRI data, as depicted in **Supplementary Fig. 1**: during the WM and the motor task conditions, FC followed reliable trajectories across multiple dimensions. Compared to the resting-state condition, there is a higher variability between the states of different HMM runs during the task conditions, particularly during the WM task. Due to this increased variability, more HMM runs were thus needed to ensure stable PCA results for the task conditions. However, similarly to the resting-state condition, the top five PCs were highly stable during both task conditions ($r > 0.9$); these retained around 40% of the total variance present across the 2400 states, suggesting that FC organises into even higher dimensions.

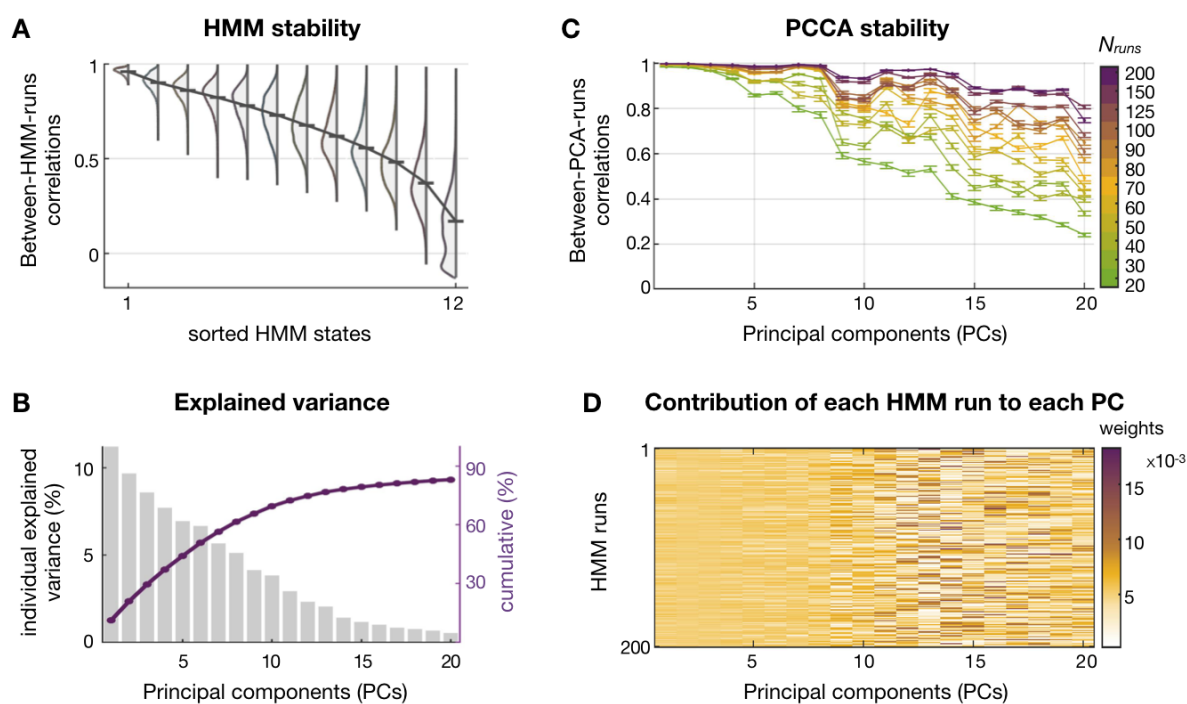


Fig. 2. Resting-state time-varying FC spans multiple temporal dimensions that can be reliably measured.

A Pearson correlation coefficient between states for each pair of HMM runs (aligned to coincide in their ordering), for $K = 12$ states and 200 HMM runs (i.e. 19900 pairs). **B** Plot representing the percentage of individual (bars) and cumulative (line) explained variance by the top 20 PCs that resulted from applying PCA to the $K = 12$ states from 200 HMM runs (i.e., $12 \times 200 = 2400$ states). **C** Pearson correlation coefficient (y-axis) of the 20 PCs (x-axis) between each pair of 30 PCA runs, for several numbers of HMM runs (from 20 to 200); error bars represent the standard error of the mean correlation across all pairs of PCA runs. **D** Average state contribution (PCA weights) of each HMM run to the final PCCA model for the top 20 PCs.

These results indicate that, despite the stochastic nature of the HMM, momentary changes in FC can be reliably captured; and that these changes occur in a multidimensional space, where at least five dimensions could be reliably measured.

FC fluctuations depend on the behavioural context

We examined whether the spatial patterns of the estimated FC trajectories were specific to the different behavioural conditions (i.e., rest, WM and motor). For this purpose, we transformed the covariance matrices of the top 20 PCs of each behavioural condition into correlation matrices and then applied the Fisher z-transformation on the off-diagonal elements of the matrices; these were then vectorised and compared to one another using Pearson correlation (**Fig. 3A**). To determine whether the overall correlation within conditions was significantly higher than between conditions we used permutation testing (significance level of 0.05; 10000 permutations); see **Fig 3B**. See brain connectivity graphs and connectivity maps of the top five PCs for each behavioural condition in **Supplementary Fig. 2**.

First, we found that changes in FC were condition-specific, that is, they depend on the behavioural context. As indicated by the red block in **Fig. 3B**, the overall similarity across PCs (i.e., the averaged correlation across the top 20 PCs) was significantly higher within than between conditions ($p < 0.0001$). Second, the largest differences were observed between the rest and task conditions, i.e. the similarity of the spatial patterns between the two tasks was always higher than between a task and the rest condition ($p < 0.0001$; lower green triangle in **Fig. 3B**). Although the WM and motor tasks were designed to target different cognitive processes (i.e., working memory function and motor execution, respectively), there are however common aspects to the two tasks such as motor coordination and visual processing, and potentially other factors related to attention (e.g., following instructions), that might be causing the resting-state to be more markedly different. Third, we observed that the differences between rest and task PCs were even more pronounced for the WM than for the motor task ($p < 0.0001$), which could partially be explained by the higher experimental demands of the WM task compared to the motor task. Fourth, we found that the spatial connectivity patterns of each condition were more diverse (i.e., lower correlation across PCs) during rest than during task conditions ($p < 0.01$; upper green triangle in **Fig. 3B**), probably due to the higher capacity of the participants to engage in a more unconstrained type of cognition during the resting-state (Wang et al., 2018).

In summary, these results indicate that the dominant modulations in FC meaningfully relate to the behavioural context and that these modulations are more diverse during rest than during task conditions —reflecting the unconstrained nature of the resting-state.




Fig. 3. The spatial patterns of the FC trajectories are condition-specific. **A** Correlation (lower triangular) and partial correlation (upper triangular) coefficients across the top 20 PCs within and between conditions: rest, WM and motor (each row and column is a different PC). **B** Significant differences (lower in green and higher in red; significance level of 0.05) of the averaged similarity of the spatial patterns (of rows with respect to columns) for all pairs of within (rest, WM, motor) and between (rest-WM, rest-motor, WM-motor) conditions.

FC fluctuates at timescales that are compatible with cognition

Having shown that the nature of the estimated FC trajectories is context-dependent, we sought to investigate the temporal scales of these modulations. We reasoned that dynamic neural communication and ongoing cognition should be reflected at relatively short timescales. In particular, focusing on the memory task, we asked how differences in cognitive abilities (i.e., WM performance) and cognitive demands (i.e., WM load) would relate to each of the different timescales underlying the FC trajectories.

To dissect the various FC-related timescales, we used the empirical mode decomposition (EMD, Huang et al., 1998), a data-driven approach that allowed us to split each FC trajectory (i.e. each PC time series) into different so-called intrinsic mode functions (IMFs). Each IMF covers a certain range of frequency, with one instantaneous frequency value per time point. Focusing on the top five PCs, the EMD yielded ten IMFs for each PC time series (see **Fig. 4A** and **Supplementary Fig. 3**), where the first IMF (IMF1) corresponded to the fastest scale (with the range 0.03Hz to 0.31Hz covering 90% of the instantaneous frequency values) and the last IMF (IMF10) to the slowest scale (90% range of 0.02Hz to 0.09Hz). To assess the relationship between IMF-specific PC time series and both WM variables, we used canonical correlation analysis (CCA) and permutation testing —with 10000 permutations per test accounting for the family structure of the HCP data (Smith et al., 2015; Winkler et al., 2015). Specifically, we ran a separate test for each IMF (i.e. data-driven frequency band) and each behavioural variable (WM performance and WM load) to assess the canonical correlation between the IMF's variance across the top 5 PCs (which can be conceptualised as the “energy” of the various FC trajectory dimensions within an EMD-estimated frequency band) and the

behavioural aspects of the task. In terms of the behaviour, WM performance was encoded as a two columns matrix, containing accuracy and reaction time; here, permutations were performed across subjects. WM load was encoded as a single vector, which elements are minus one for 0-back and plus one for 2-back blocks; here, permutations were performed within each subject across blocks and runs. False discovery rate (FDR; Benjamini and Hochberg, 1995) was used to correct across multiple comparisons.

For illustration, **Fig. 4B** shows how the IMF3's energy of the first three PCs varies across subjects. IMF3 was chosen since it exhibited a highly significant relation to behaviour (see below) and spans relatively fast timescales (mean frequency is 0.01 Hz), therefore being a reasonable candidate to reflect aspects of dynamic, whole-brain neural communication. To further illustrate how the FC trajectories vary across subjects, **Fig. 4C** projects the IMF3 time series for the subject with the lowest (in blue) and the highest (in red) IMF3 energy. For WM performance, the (first) canonical correlation was statistically significant ($r = 0.45$, $p < 0.01$, $p\text{-FDR} = 0.045$). **Fig. 4D** shows a scatter plot of the canonical covariate related to WM performance vs the canonical covariate related to the IMF3-related energies. For WM load, the fastest IMFs (IMF1–IMF3) showed a significant ($p\text{-FDR} < 0.05$) canonical correlation with WM load of, respectively, $r = 0.31$, $r = 0.36$ and $r = 0.29$. **Fig. 4E** shows the distribution of the canonical covariates for IMF1 to IMF3, obtained as an optimal linear projection of their energy values for low (in yellow) and high load (in purple). See **Supplementary Table 1** for a display of all canonical correlations, permutation p-values and FDR p-values.

Altogether, these analyses revealed that the three modes (IMF1–IMF3) that capture relatively fast changes in FC (i.e., from seconds to tens of seconds) relate to behavioural information such as cognitive abilities and cognitive demands.

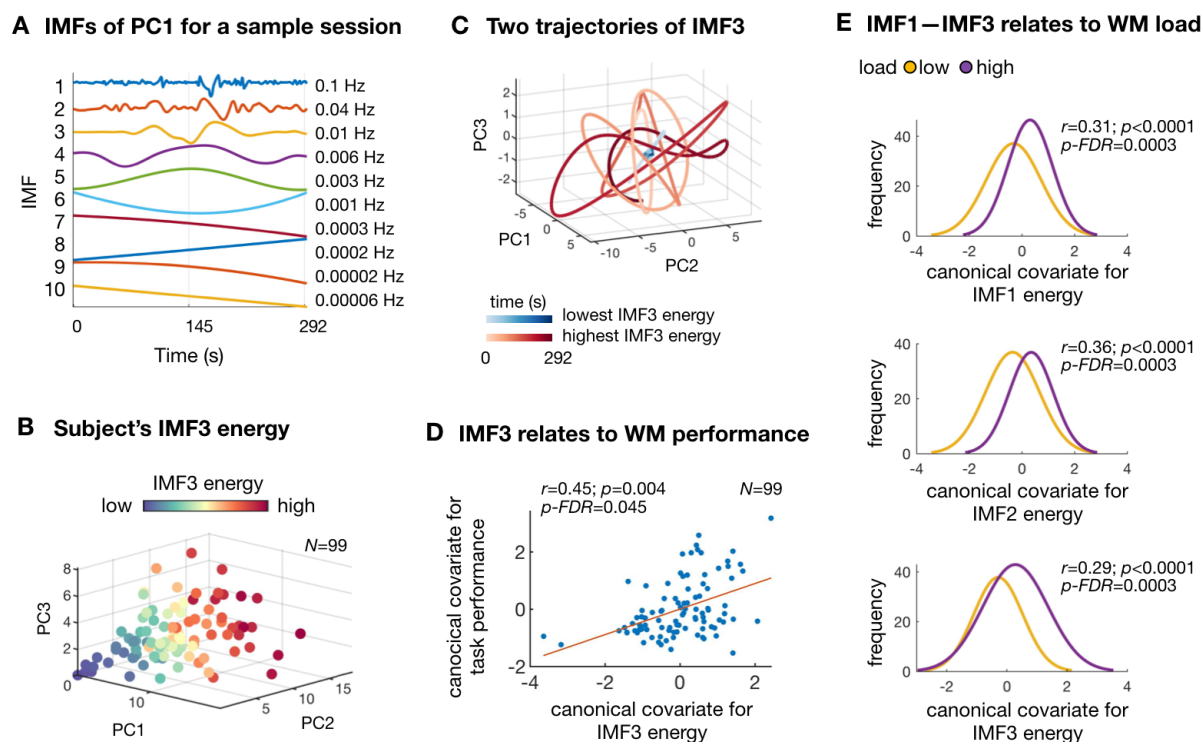


Fig. 4. Fast changes in FC are strongly associated with cognitive abilities (WM performance) and cognitive demands (WM load). **A** IMFs from an illustrative session obtained from PC1's time series; on the right, median instantaneous frequency of each IMF. **B** Individual subject values of the IMF3's energy for the top three PCs, coloured from low to high energy (blue to red colour bar). **C** Two subject's trajectories of IMF3 with the lowest (blue) and the highest (red) energy; colour bars indicate the time of the session from light to dark colours. **D** Scatterplot of the relationship between the canonical covariate for IMF3's energy values and the canonical covariate for the WM performance scores ($N = 99$ subjects); each dot represents the values for one subject. **E** Distribution of the canonical covariates for the energies of IMF1–IMF3, for low (in yellow) and high WM load (in purple) trials.

Functional connectivity trajectories are not driven by simple activation patterns

Since we are interested in whether dynamic cross-regional communication can be related to FC trajectories, we sought to verify that such trajectories are not determined just by amplitude modulations. To disambiguate whether the PC modulations contain information that is unique to time-varying FC (i.e. above and beyond just regional activations and amplitude modulations), we examined how much of the temporal covariation in FC can be predicted by changes in the raw BOLD signal (represented as 25 IC time series). Using cross-validated regularized ridge regression (Hoerl and Kennard, 1970) on each PC separately, we show that the main patterns of temporal fluctuations in FC during rest cannot be well explained by changes in order-1 activations (**Fig. 5**). Specifically, the IC time series can only explain 1% of the variance of the top four PCs and 11% of the fifth PC. In contrast, the order-2 interactions of the signal (i.e. the pairwise product of the IC time series, which is directly related to Pearson-based FC), explained considerably more variance of the PCs (values ranging between 18% and 43%). Note that because the predictions are cross-validated, this comparison is not biased by the fact that the number of order-2 interactions (300 time series,

one per pair of ICs) is larger than the number of order-1 activations (25 IC time series). Moreover, since the order-2 interactions also contain the information of the amplitude itself (within the diagonal of the covariance matrices), we repeated the prediction after regressing out the raw signal of the PC time series in a cross-validated fashion (Snoek et al., 2019). On average, less than 2% of the explained variance was lost after accounting for the raw signal, indicating that the PCs do contain information that is unique to time-varying FC.

During the task conditions, however, the changes in order-1 time series contributed more to explain the fluctuations in FC. This was expected since the evoked responses are synchronised across brain regions by the task experimental design, creating task-induced correlations and therefore contributing to the FC estimations. This effect was more pronounced for the motor task than for the WM task condition. Specifically, the IC time series could explain up to 48% of the variance of the first PC during the motor task, and 32% of the second PC during the WM task. Order-2 interactions of the signal had similar contributions to the amplitude for the first two PCs of the WM task and the first three PCs of the motor task. For the remaining PCs, order-2 interactions explained more variance than the IC time series. As before, we repeated the prediction after correcting for the order-1 activations, finding that the variance explained by order-2 interactions decreased around 12% for the WM task, and 21% for the motor task, after deconfounding for the IC time series. This loss of variance indicates that a considerable proportion of the fluctuations in FC are driven by changes in signal amplitude that are likely evoked by task responses.




Fig. 5. PCCA time series cannot be well explained by amplitude. Percentage of explained variance (R^2 , y-axis) of the top 5 PCs (x-axis) by the (order-1) IC time series, and order-2 interactions with and without deconfounding for amplitude. Components are coloured according to 3 conditions: rest, WM and motor.

Reproducibility of the results

We sought to evaluate the reproducibility of our results with respect to five different variations of the experimental setup (Fig. 6, left to right): using a separate set of subjects, using the original dataset with different model configurations (changing the fixed number of HMM states from 12 to 6 and to 16), and using different data configurations (reducing the length of the time series from 4800 to 810 frames; and using a subsample with half of the subjects). We focused on the resting-state, free of task-related activations that could trivially explain the reproducibility scores, and used the top 20 PCs (instead of five) for completeness.

For each of these five experimental variations, we first extracted the top 20 PCs following the same procedure as described in **Fig. 1**. As discussed, we transformed the covariance matrices of each of these 20 PCs into correlation matrices, applied the Fisher z-transformation on the off-diagonal elements of the matrices and vectorised them. These vectors were then compared to those of the original dataset using Pearson correlation. Reproducibility was assessed by comparing whether the correlation between aligned PCs was significantly higher than between nonaligned PCs. The statistical significance of the difference was determined using permutation testing (significance level of 0.05; 10000 permutations).

We found that the overall changes in FC (as captured by the top 20 PCs) were highly reproducible for all five experimental variations. That is, the correlation between aligned PCs was significantly higher than between non-aligned PCs ($p < 10^{-5}$). In particular, when using a separate set of subjects (first column of **Fig. 6**) only three out of 20 aligned PCs (orange) fall below the 75th percentile of the correlations between non-aligned PCs (green). With respect to changes in model configurations, reducing the number of fixed states from 12 to 6 (second column of **Fig. 6**) returned 14 very similar PCs; while increasing it from 12 to 16 (third column of **Fig. 6**) returned 19 very similar PCs, suggesting that our approach is robust to the number of HMM states within a reasonable range. When reducing the length of the time series to 810 frames (the length of the WM task), 15 of the 20 PCs could still be recovered even though only a fifth of the original number of volumes were available (fourth column of **Fig. 6**). Lastly, only one of the 20 PCs could not be recovered when estimating the PCCA model on just a half of the subjects from the original dataset (fifth column of **Fig. 6**).




Fig. 6. Reproducibility of the time-varying FC patterns for the top 20 PCs in the resting-state. From left to right, reproducibility of our results on a separate set of subjects (i.e., 100 subjects, 4600 frames, 12 states), and with respect to changes in the number of states (from 12 to 6 and 16 per HMM run), to different scan lengths (from 4600 frames to 810), and different sample sizes (from 100 to 50 subjects). Boxplots represent the inter-quartile range (IQR; between the

75th and 25th percentiles) and median (horizontal line) of the correlation values (one per dot) between aligned PCs (orange) and nonaligned PCs (green).

Discussion

Can the temporal variability of patterns of fMRI FC across multiple regions reflect information transfer between the underlying neuronal populations, or, rather, they just reflect stable anatomical pathways? We propose that, for the former case to be true, these correlations should fulfil the following criteria: (i) its latent dimensionality must be relatively low to be reliable, but still high enough to accommodate the flexibility that brain communication requires; (ii) it should adapt to the behavioural context; and (iii) it should exhibit relevant changes at fast timescales. Using Principal components of connectivity analysis, a new approach based on multiple runs of the HMM in combination with PCA, we showed that reliable estimations of FC patterns can be represented as low-dimensional trajectories of order-2 interactions. Importantly, these could not be trivially explained from order-1 signal activations. These FC-related trajectories were found to meet the proposed criteria, suggesting that they may faithfully reflect at least some aspects of large-scale, cross-regional communication.

The fact that not only one but several FC dimensions could be reliably estimated might seem at odds with previous work where network activity was shown to be hierarchically organized into two groups of states or metastates (Vidaurre et al., 2017). Indeed, having the states strongly clustered into two metastates essentially defines a one-dimensional axis of variation for FC. Two main reasons explain this discrepancy. First, in Vidaurre et al. (2017) the HMM was run only once, whereas here we have used multiple runs; this allowed us to explore the space of solutions exhaustively. Second, and most importantly, while in the previous study PCA was applied on the fractional occupancies (an aggregated measure where each session is defined as how much time the subject spent on each state), here we have applied PCA on the state time series. That is, in the previous study PCA was run on a (*no. of subjects* by *no. of states*) matrix, and here PCA is applied on a much larger ($[no. of subjects \times no. of time points \times no. of sessions]$ by $[no. of states \times no. of HMM runs]$) matrix. Therefore, whereas the metastates in Vidaurre et al. (2017) fundamentally reflect between-subject variability of *both* time-varying and time-averaged FC, PCCA characterises the actual temporal dimension of the FC patterns, where multiple dimensions are shown to be at play.

We have investigated whether the estimated FC-related trajectories reflect the behavioural context and cognitive demands by comparing the brain representations of FC fluctuations across three experimental conditions: at rest, during a WM task, and during a motor task. On the one hand, our analysis revealed distinct spatial patterns for each of the three conditions, indicating the behavioural specificity of the estimated trajectories. On the other hand, we found a higher similarity of brain network modulations during the execution of tasks, highlighting the constrained element of both WM and motor tasks as opposed to the unconstrained nature of the resting-state. A plausible explanation is that the task designs elicit patterns of synchronised activity that are to some extent generalisable across tasks (for instance related to attentional processes), and these become reflected in the network

trajectories. This observation aligns with previous evidence that task-induced changes in FC are not simply driven by task-evoked activation (Greene et al., 2020). Recent studies analysing how brain activity unfolds within a low-dimensional space have also found common patterns across several tasks (Saggar et al., 2018; Shine et al., 2019a), and observed that modulations in these patterns have been associated with differences in task demands and task performance (Cornblath et al., 2020; Saggar et al., 2018; Shine et al., 2019b). As opposed to our PCCA approach, however, these latent trajectories were computed based on order-1 statistics and are therefore unspecific to connectivity. Our work extends these findings by demonstrating that cognitive demands and cognitive performance shape the temporal evolution of brain FC (i.e. not just activity) at relatively fast timescales.

These results align with previous reports on the variability in BOLD signals as an important feature of brain function (Garrett et al., 2020, 2010; McIntosh et al., 2008). Using EEG and fMRI, these studies showed that higher brain signal variability facilitates the formation of brain networks and transitions between them, which in turn enhances the brain's dynamical repertoire to enable cognitive function. Complementing these studies, we quantified the degree of variability in the space of brain network interactions, therefore extending the notion of signal variability towards connectivity variability. In doing so, we found that connectivity variability at relatively fast frequencies tracks cognitive load and task performance. Furthermore, we found brain network modulations to be more diverse during rest than during task, reflecting the unconstrained nature of the resting-state and supporting the idea that resting-state FC may provide a richer characterization of brain activity than task-states (Ponce-Alvarez et al., 2015).

Previous research on time-varying FC has often identified a set of brain states occurring within the full scanning session (Lurie et al., 2020), for example using an HMM (Vidaurre et al., 2018) or by performing (e.g. k-means) clustering on sliding window estimates (Allen et al., 2014). Despite the specifics of each approach, the choice of the number of states is a cross-methodological concern. Here, we found that irrespective of the initial number of fixed states ($K = 6, 12$ and 16), PCA was able to reliably capture the same prevalent underlying configurations. The highly robust estimation across choices of K indicates that the relative instability of the HMM inference can be turned into an advantage when harnessed by PCA; see Vidaurre et al., 2019 for a similar argument in the context of statistical testing. Methodologically, PCA might offer a promising alternative to conventional model selection, toward model integration.

In summary, our results indicate that brain network modulations fluctuate at short timescales following reliable patterns that are context-dependent, influenced by cognitive complexity and associated with cognitive abilities. The estimated trajectories are relatively low- but still multi-dimensional, meeting the minimum requisites to reflect the dynamics of information transfer across the brain.

Methods

Data Description

We used fMRI data from 100 subjects from the Human Connectome Project (HCP, van Essen et al., 2013) publicly available. For each subject, we considered data from the resting-state (rest), the working memory (WM) task, and the motor task. One subject's WM task data was missing. Brief descriptions of each condition are provided in **Table 1**; please refer to Van Essen et al., 2012 for full details about the acquisition and preprocessing of the data. In brief, whole-brain fMRI data was acquired at 3T, with a spatial resolution of 2×2×2 mm and a temporal resolution of 0.72 s. All fMRI processing was performed using FSL (Jenkinson et al., 2012) including minimal high-pass temporal filtering (< 0.01 Hz) to remove the linear trends of the data, no low-pass filter, no global signal regression, and artefact removal using independent component analysis (ICA)+FIX (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014). In particular, group spatial-ICA was performed using MELODIC (Beckmann and Smith, 2004) on the resting-state data to obtaining a parcellation of 25 ICs. For ease of comparison across conditions, the time series for each IC was computed —separately for each subject and condition (rest, WM and motor)— using the group 25-ICA parcellation from the resting-state data through dual regression (Filippini et al., 2009). Specifically, the IC time series corresponding to the task sessions were computed from the fully preprocessed data in MNI152 space as delivered in the HCP1200 release; the IC time series of the resting-state sessions were directly obtained from the HCP 'PTN' Parcellation+Timeseries+Netmats (specifically, the first 100 subjects from the 'recon2' version).

Table 1. Data description

Condition	Subjects	Runs	Frames per run	Description
Rest	100	4	1200	Participants were instructed to try to lie still, to keep their eyes open (relaxed fixation), not to think of anything specific, and not to fall asleep (although some likely fall sleep at some point).
Working Memory (WM)	99	2	405	Participants were presented with pictures (places, faces, body parts and tools) in alternating blocks of 0-back and 2-back WM. During 0-back blocks, participants had to respond to any presentation a target cue (presented at the start of the block). During 2-back blocks, participants had to respond whether the current stimulus was the same as two stimuli earlier.
Motor	100	2	284	Participants were presented with visual cues that required them to tap their left or right fingers, squeeze their left or right toes, or move their tongue.

Principal Component of Connectivity Analysis

Principal Component of Connectivity Analysis, or PCCA, is composed of two methodological elements: the HMM and PCA. We next describe these, including how to compute connectivity maps from the resulting principal connectivity components.

Hidden Markov modelling. The HMM was used to extract whole-brain patterns of time-varying FC during the rest and task conditions. In a data driven way, the HMM characterises neural time series of concatenated data using a finite number of K (group-level) states that reoccur across time (Rabiner, 1989), where the k -th state time series represents the probability for that state to be active at each time point; that is, the spatial parameters of the HMM are at the group level, while the temporal parameters are at the subject level. The states themselves are probability distributions within a certain family, and each is characterised by a certain set of parameters. In order to focus on FC, each state was here characterized by a Gaussian distribution with no mean parameter (i.e. with the mean parameter pinned to zero) and a full covariance matrix (representing pairwise covariations across regions); or, equivalently, as a Wishart distribution (Vidaurre et al., 2021). Finally, a K -by- K transition probability matrix is estimated as part of the model, containing the probability of transitioning from one state to another, or to remain in the current state. Here, the number of K states was set to 12 and the HMM was applied to the concatenated time series for all subjects ($[N_{runs} \times K]$ by $[N_{subjects} \times N_{frames}]$) for each condition separately (rest, WM task and motor task; see **Table 1**). See **Fig. 1A-B** for a schematic description of the HMM.

Principal component analysis and component characterization. PCA was applied on the HMM state time series to integrate the time-varying FC estimations across many runs. In short, the state time series generated over multiple runs (N_{runs}) were attached together into a matrix (Y) with as many rows as frames in the data set ($N_{subjects} \times N_{frames}$) and M columns ($N_{runs} \times K$ states); see **Fig. 1B-C**. PCA was then performed by eigendecomposing $Y Y^T$, which returns a set of M eigenvectors and eigenvalues according to the equation:

$$Y Y^T = U \Lambda U^T \quad (1)$$

where U contained the (orthonormal) eigenvectors on its columns and Λ contained the corresponding eigenvalues in its diagonal (**Fig. 1C**). These eigenvectors, or PCs, represented patterns of variation across all time-varying FC states, while the eigenvalues described the contribution of each PC in terms of explained variance. The PCs (columns of U) were arranged according to their relative contributions to Y , from higher to lower eigenvalues.

To estimate the PC-specific spatial FC patterns for the i -th PC, we computed one covariance matrix for each PC (denoted as Q_i) by weighting the covariance matrices of the original states (C_m) with weights β_m as follows:

$$Q_i = \sum_m \beta_m C_m \quad (2)$$

The weights β_m were calculated as the variance of the (1 by [$N_{subjects} \times N_{frames}$]) PC time series S_i , during the time when state m is active:

$$\beta_m = Y_m^T S_i^2 / |Y_m|_1 \quad (3)$$

where S_i represents the time series of the i -th temporal PC, Y_m is the probability of that state to be active or not at every time point (**Fig. 1D**), and $|Y_m|_1$ represents the sum of all elements of Y_m (i.e. the L^1 norm of the vector). Note that it is not possible to estimate the PCs' covariance matrices using the PCA weights directly, because these are not necessarily positive and therefore the resulting matrix would not necessarily be positive definite (a mathematical requirement for any covariance matrix). Therefore, Eq. (3) is defined by convention, and does not follow from first principles.

To facilitate visualization, the spatial connectivity map of each PC was then derived using the first eigenvector of the PC covariance matrix (Q_i), which represented the main pattern of FC for that PC. Note that this is a different eigendecomposition than Eq. (1), which is done on the PC time series. The brain connectivity maps shown in **Supplementary Fig. 2** were created by overlaying these eigenvectors on brain surface space.

Empirical mode decomposition

Using an iterative algorithm, the EMD aims at decomposing multi-frequency time series into a number of oscillatory modes referred to as Intrinsic Mode Functions (IMFs), where each IMF relates to one fundamental frequency (Huang et al., 1998; Quinn et al., 2021). Here, we used the EMD to decompose each FC trajectory (i.e. PC time series) into a set of IMF-specific time series of the same length as the PC time series. The number of IMFs were determined in a data-driven way by the intrinsic temporal-spectral characteristics of the PC time series. The decomposition of the PC time series yielded ten IMFs (IMF1–IMF10), where each IMF covered a certain frequency range; see **Fig. 4A** and **Supplementary Fig. 3**. The instantaneous frequencies of each IMF were obtained using the Hilbert-Huang Transform (Huang et al., 2009).

Code availability

The code used to analyse the data in this study comprises our open-access code repository (<https://github.com/OHBA-analysis/HMM-MAR>) and custom Matlab scripts publicly available at <https://github.com/sonsolesalonsomartinez/PCCA>.

References

- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cereb. Cortex* 24, 663–676. <https://doi.org/10.1093/cercor/bhs352>
- Averbeck, B.B., Latham, P.E., Pouget, A., 2006. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–66. <https://doi.org/10.1038/nrn1888>
- Averbeck, B.B., Lee, D., 2004. Coding and transmission of information by neural ensembles. *Trends Neurosci.* <https://doi.org/10.1016/j.tins.2004.02.006>
- Beckmann, C.F., Smith, S.M., 2004. Probabilistic Independent Component Analysis for Functional Magnetic Resonance Imaging. *IEEE Trans. Med. Imaging* 23, 137–152. <https://doi.org/10.1109/TMI.2003.822821>
- Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *J. R. Stat. Soc. Ser. B* 57, 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Cohen, M.R., Kohn, A., 2011. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–9. <https://doi.org/10.1038/nn.2842>
- Cornblath, E.J., Ashourvan, A., Kim, J.Z., Betzel, R.F., Ciric, R., Adebimpe, A., Baum, G.L., He, X., Ruparel, K., Moore, T.M., Gur, R.C., Gur, R.E., Shinohara, R.T., Roalf, D.R., Satterthwaite, T.D., Bassett, D.S., 2020. Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. *Commun. Biol.* 3, 261. <https://doi.org/10.1038/s42003-020-0961-x>
- Dehaene, S., Naccache, L., 2001. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79, 1–37. [https://doi.org/10.1016/s0010-0277\(00\)00123-2](https://doi.org/10.1016/s0010-0277(00)00123-2)
- Essen, D.C. Van, Smith, S.M., Barch, D.M., Behrens, T.E.J.J., Yacoub, E., Ugurbil, K., Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J.J., Yacoub, E., Ugurbil, K., 2013. The WU-Minn Human Connectome Project: An overview. *Neuroimage* 80, 62–79. <https://doi.org/10.1016/j.neuroimage.2013.05.041>
- Filippini, N., MacIntosh, B.J., Hough, M.G., Goodwin, G.M., Frisoni, G.B., Smith, S.M., Matthews, P.M., Beckmann, C.F., Mackay, C.E., 2009. Distinct patterns of brain activity in young carriers of the APOE -ε4 allele. *Proc. Natl. Acad. Sci.* 106, 7209–7214. <https://doi.org/10.1073/pnas.0811879106>
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nat. Neurosci.* 18, 1664–1671. <https://doi.org/10.1038/nn.4135>
- Fox, M.D., Greicius, M., 2010. Clinical applications of resting state functional connectivity. *Front. Syst. Neurosci.* <https://doi.org/10.3389/fnsys.2010.00019>
- Garrett, D.D., Epp, S.M., Kleemeyer, M., Lindenberger, U., Polk, T.A., 2020. Higher performers upregulate brain signal variability in response to more feature-rich visual input. *Neuroimage* 217, 116836. <https://doi.org/10.1016/j.neuroimage.2020.116836>
- Garrett, D.D., Kovacevic, N., McIntosh, A.R., Grady, C.L., 2010. Blood oxygen level-dependent signal variability is more than just noise. *J. Neurosci.* 30, 4914–4921. <https://doi.org/10.1523/JNEUROSCI.5166-09.2010>
- Greene, A.S., Gao, S., Noble, S., Scheinost, D., Constable, R.T., 2020. How Tasks Change Whole-Brain Functional Organization to Reveal Brain-Phenotype Relationships. *Cell Rep.*

- 32, 108066. <https://doi.org/10.1016/j.celrep.2020.108066>
- 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., Moeller, S., Xu, J., Yacoub, E., Baselli, G., Ugurbil, K., Miller, K.L., Smith, S.M., 2014. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* 95, 232–247. <https://doi.org/10.1016/j.neuroimage.2014.03.034>
- Hasson, U., Nusbaum, H.C., Small, S.L., 2009. Task-dependent organization of brain regions active during rest. *Proc. Natl. Acad. Sci. U. S. A.* 106, 10841–6. <https://doi.org/10.1073/pnas.0903253106>
- Haynes, J.-D., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523–534. <https://doi.org/10.1038/nrn1931>
- Hindriks, R., Adhikari, M.H.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *Neuroimage* 127, 242–256. <https://doi.org/10.1016/j.neuroimage.2015.11.055>
- Hoerl, A.E., Kennard, R.W., 1970. American Society for Quality Ridge Regression: Biased Estimation for Nonorthogonal Problems. Source: *Technometrics* 12, 55–67.
- Huang, N.E., Shen, Z., Long, S.R., Wu, M.C., Snin, H.H., Zheng, Q., Yen, N.C., Tung, C.C., Liu, H.H., 1998. The empirical mode decomposition and the Hubert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. A Math. Phys. Eng. Sci.* 454, 903–995. <https://doi.org/10.1098/rspa.1998.0193>
- Huang, N.E., Wu, Z., Long, S.R., Arnold, K.C., Chen, X., Blank, K., 2009. ON INSTANTANEOUS FREQUENCY. *Adv. Adapt. Data Anal.* 01, 177–229. <https://doi.org/10.1142/S1793536909000096>
- Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. *Neuroimage* 62, 782–790. <https://doi.org/10.1016/j.neuroimage.2011.09.015>
- Lu, H., Zou, Q., Gu, H., Raichle, M.E., Stein, E.A., Yang, Y., 2012. Rat brains also have a default mode network. *Proc. Natl. Acad. Sci.* 109, 3979–3984. <https://doi.org/10.1073/pnas.1200506109>
- Lurie, D.J., Kessler, D., Bassett, D.S., Betzel, R.F., Breakspear, M., Kheilholz, S., Kucyi, A., Liégeois, R., Lindquist, M.A., McIntosh, A.R., Poldrack, R.A., Shine, J.M., Thompson, W.H., Bielczyk, N.Z., Douw, L., Kraft, D., Miller, R.L., Muthuraman, M., Pasquini, L., Razi, A., Vidaurre, D., Xie, H., Calhoun, V.D., 2020. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw. Neurosci.* 4, 30–69. https://doi.org/10.1162/netn_a_00116
- 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. *Proc. Natl. Acad. Sci. U. S. A.* 113, 12574–12579. <https://doi.org/10.1073/pnas.1608282113>
- McIntosh, A.R., Kovacevic, N., Itier, R.J., 2008. Increased brain signal variability accompanies lower behavioral variability in development. *PLoS Comput. Biol.* 4. <https://doi.org/10.1371/journal.pcbi.1000106>
- Munkres, J., 1957. Algorithms for the Assignment and Transportation Problems. *J. Soc. Ind. Appl. Math.* 5, 32–38. <https://doi.org/10.1137/0105003>
- Nienborg, H., Cumming, B.G., 2009. Decision-related activity in sensory neurons reflects more than a neurons causal effect. *Nature* 459, 89–92. <https://doi.org/10.1038/nature07821>

- Nirenberg, S., Latham, P.E., 2003. Decoding neuronal spike trains: How important are correlations? *Proc. Natl. Acad. Sci. U. S. A.* 100, 7348–7353. <https://doi.org/10.1073/pnas.1131895100>
- Ponce-Alvarez, A., Deco, G., Hagmann, P., Romani, G.L., Mantini, D., Corbetta, M., 2015. Resting-State Temporal Synchronization Networks Emerge from Connectivity Topology and Heterogeneity. *PLOS Comput. Biol.* 11, e1004100. <https://doi.org/10.1371/journal.pcbi.1004100>
- Quinn, A.J., Lopes-dos-Santos, V., Huang, N., Liang, W.-K., Juan, C.-H., Yeh, J.-R., Nobre, A.C., Dupret, D., Woolrich, M.W., 2021. Within-cycle instantaneous frequency profiles report oscillatory waveform dynamics. *bioRxiv* 2021.04.12.439547. <https://doi.org/10.1101/2021.04.12.439547>
- Rabiner, L.R., 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77, 257–286. <https://doi.org/10.1109/5.18626>
- Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P.A., Carlsson, G., Glover, G., Reiss, A.L., 2018. Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nat. Commun.* 9, 1399. <https://doi.org/10.1038/s41467-018-03664-4>
- Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M., 2014. Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468. <https://doi.org/10.1016/j.neuroimage.2013.11.046>
- Shine, J.M., Breakspear, M., Bell, P.T., Ehgoetz Martens, K., Shine, R., Koyejo, O., Sporns, O., Poldrack, R.A., 2019a. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nat. Neurosci.* 22, 289–296. <https://doi.org/10.1038/s41593-018-0312-0>
- Shine, J.M., Hearne, L.J., Breakspear, M., Hwang, K., Müller, E.J., Sporns, O., Poldrack, R.A., Mattingley, J.B., Cocchi, L., 2019b. The Low-Dimensional Neural Architecture of Cognitive Complexity Is Related to Activity in Medial Thalamic Nuclei. *Neuron* 104, 849–855.e3. <https://doi.org/10.1016/j.neuron.2019.09.002>
- Smith, S.M., Nichols, T.E., Vidaurre, D., Winkler, A.M., Behrens, T.E.J., Glasser, M.F., Ugurbil, K., Barch, D.M., Van Essen, D.C., Miller, K.L., 2015. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nat. Neurosci.* <https://doi.org/10.1038/nn.4125>
- Snoek, L., Miletić, S., Scholte, H.S., 2019. How to control for confounds in decoding analyses of neuroimaging data. *Neuroimage* 184, 741–760. <https://doi.org/10.1016/j.neuroimage.2018.09.074>
- Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E., 2012. The Human Connectome Project: A data acquisition perspective. *Neuroimage* 62, 2222–2231. <https://doi.org/10.1016/j.neuroimage.2012.02.018>
- Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A.J., Alfaro-Almagro, F., Smith, S.M., Woolrich, M.W., D, V., R, A., R, B., AJ, Q., F, A.-A., SM, S., MW, W., Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A.J., Alfaro-Almagro, F., Smith, S.M., Woolrich, M.W., 2018. Discovering dynamic brain networks from big data in rest and task, *NeuroImage*. Academic Press Inc. <https://doi.org/10.1016/j.neuroimage.2017.06.077>

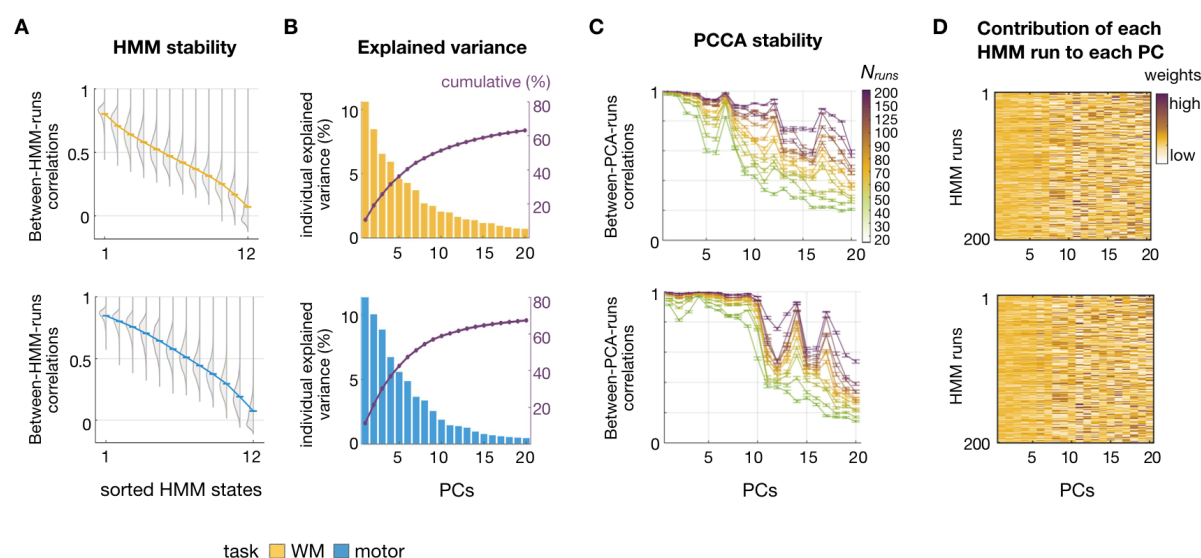
- Vidaurre, D., Llera, A., Smith, S.M., Woolrich, M.W., 2021. Behavioural relevance of spontaneous, transient brain network interactions in fMRI. *Neuroimage* 229, 117713. <https://doi.org/10.1016/j.neuroimage.2020.117713>
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci.* 114, 12827–12832. <https://doi.org/10.1073/pnas.1705120114>
- Vidaurre, D., Woolrich, M.W., Winkler, A.M., Karapanagiotidis, T., Smallwood, J., Nichols, T.E., 2019. Stable between-subject statistical inference from unstable within-subject functional connectivity estimates. *Hum. Brain Mapp.* 40, 1234–1243. <https://doi.org/10.1002/hbm.24442>
- Wang, H.T., Bzdok, D., Margulies, D., Craddock, C., Milham, M., Jefferies, E., Smallwood, J., 2018. Patterns of thought: Population variation in the associations between large-scale network organisation and self-reported experiences at rest. *Neuroimage* 176, 518–527. <https://doi.org/10.1016/j.neuroimage.2018.04.064>
- Winkler, A.M., Webster, M.A., Vidaurre, D., Nichols, T.E., Smith, S.M., 2015. Multi-level block permutation. *Neuroimage* 123, 253–268. <https://doi.org/10.1016/j.neuroimage.2015.05.092>

Supplementary Information


Supplementary Table 1. Canonical Correlation Analysis between IMF-specific PC time series and both WM performance and WM load.


	WM performance			WM load		
	r	p -Perm	p -FDR	r	p -Perm	p -FDR
IMF1	0.26	0.767	0.767	0.31	0.0001	0.0003
IMF2	0.34	0.046	0.155	0.36	0.0001	0.0003
IMF3	0.45	0.004	0.045	0.29	0.0001	0.0003
IMF4	0.28	0.397	0.662	0.10	0.3313	0.8282
IMF5	0.25	0.624	0.722	0.07	0.7004	0.9999
IMF6	0.36	0.017	0.087	0.03	0.9168	0.9999
IMF7	0.44	0.075	0.188	0.04	0.9180	0.9999
IMF8	0.27	0.650	0.722	0.004	0.9999	0.9999
IMF9	0.33	0.187	0.374	0.002	0.9998	0.9999
IMF10	0.25	0.626	0.722	0.001	0.9999	0.9999

IMF: intrinsic mode function; p -Perm: p-value of 10000 permutation test; p -FDR: false discovery rate adjusted p-value; WM: working memory.



Supplementary Fig. 1. Time-varying FC trajectories in task span multiple temporal dimensions that can be reliably measured. For WM (top) and motor (bottom) task conditions: **A** Pearson correlation coefficient between states for each pair of HMM runs (aligned to coincide in their ordering), for $K = 12$ states and 200 HMM runs (i.e. 19900 pairs). **B** Plot representing the percentage of individual (bars) and cumulative (line) explained variance by the top 20 PCs that resulted from applying PCA to the $K = 12$ states from 200 HMM runs (i.e., $12 \times 200 = 2400$ states). **C** Pearson correlation coefficient (y-axis) of the 20 PCs (x-axis) between each pair of 30 PCA runs, for several numbers of HMM runs (from 20 to 200); error bars represent the standard error of the mean correlation across all pairs of PCA runs. **D** Average state contribution (PCA weights) of each HMM run to the final PCCA model for the top 20 PCs.





Supplementary Fig. 3. Distributions of instantaneous frequencies for each IMF, for the top five PC time series. Each IMF occupies a specific frequency band with very little overlap between them. For each IMF, dots and bars show the median and the 90% frequency range for each PC (shown in colour) as well as their average (in grey). The instantaneous frequencies of each IMF were obtained using the Hilbert-Huang Transform.