

Brain Network Dynamics are Hierarchically Organised in Time

Diego Vidaurre¹, Stephen Smith¹, Mark Woolrich¹

¹University of Oxford

Submitted to Proceedings of the National Academy of Sciences of the United States of America

The brain recruits neuronal populations in a temporally coordinated manner in task and at rest. However, the extent to which large-scale networks exhibit their own organised temporal dynamics is unclear. We use an approach designed to find repeating network patterns in whole-brain resting fMRI data, where networks are defined as graphs of interacting brain areas. We find that the transitions between networks are non-random, with certain networks more likely to occur after others. Further, this non-random sequencing is itself hierarchically organized, revealing two distinct sets of networks, or metastates, that the brain has a tendency to cycle within. One metastate is associated with sensory and motor regions and the other involves areas related to higher-order cognition. Moreover, we find that the proportion of time that a subject spends in each brain network and metastate is a consistent subject-specific measure, is heritable and shows a significant relationship with cognitive traits.

Resting-state networks | Metastates | Dynamic functional connectivity | Hidden Markov Model

Introduction

Resting brain activity is far from random and has been shown to organize into a number of large-scale networks with characteristic spatial architectures (1, 2, 3). It is also known that many of the networks of activity found in the resting brain are also observed in tasks (4). However, despite evidence at the microscale that the brain deploys populations of neurons in a temporally coordinated manner both in tasks (5) and at rest (6), relatively little is known about the temporal organisation of large-scale resting state networks. This has been limited to evidence of, for example, anti-correlations between certain networks (7), long-range temporal dependencies (8), and time-varying modularity (9,10). Non-random transitions have also been found between brain states defined as EEG microstates (11), attractors in metastable systems (12) or local minima in energy landscape analysis (13). In this work we hypothesise that there must exist more specific patterns of temporal organization than have been described so far, and propose that these transitions are not only non-equiprobable, but are themselves organised in a hierarchical manner.

To study the temporal organization of the brain's large-scale network dynamics, we used a method designed to discover networks that repeat over time (referred to as brain states). Importantly, we define networks as probability distributions representing graphs, with not only distinct patterns of activation but, crucially, also of functional connectivity. This approach has access to temporal scales as fast as the data modality allows and overcomes the statistical limitations of the sliding windows technique when applied to the analysis of brain dynamics (14, 15, 16, 17). Applying this method to whole-brain resting-state fMRI data from 820 subjects from the Human Connectome project (HCP) (18) reveals a clear temporal organization. We find that transitions between different brain networks do not occur in complete randomness; instead, certain networks are much more likely to follow others in time. Further, this non-random sequencing of brain networks is itself hierarchically organized, and strikingly reveals two sets of brain networks, or *metastates*, that the brain has a tendency to cycle within, with sporadic cycling between them.

One metastate is associated with sensorimotor and perceptual (visual and auditory) regions. The other involves areas related to higher-order cognition, including regions of the default mode network, language and extensive prefrontal.

We also examine the most basic temporal characteristic of these brain networks, or states, which is their *occupancy*. Occupancy is defined as the proportion of time that each subject spends in each brain state. Notably, these occupancies show a strongly dominant mode of variation over subjects that correspond to the same two metastates that were apparent in the sequencing of brain networks. We then demonstrate that the proportion of time spent in each metastate is very subject-specific and significantly heritable. Not only can we predict the metastate occupancy in a new session using the occupancy from previous sessions, but we can also significantly predict the occupancy using behavioural information of the subject. In particular, more occupancy of the cognitive metastate tends to relate to positive traits, in particular to cognitive performance and satisfaction. We analyse the relation of the metastates' distribution to sleep and motion, concluding that the high-level temporal dynamics are not an artefactual result of these factors.

In summary, this paper highlights the hierarchically organised temporal nature of resting-state networks of interacting brain areas, characterising its properties in a large cohort of subjects, and relating its cross-subject variability to behaviour and heritability.

Results

Dynamic switching between brain networks is not random

We used resting state fMRI data from 820 subjects in the Human Connectome Project (18). Data across the multiple subjects

Significance

We address the important question of the temporal organization of large-scale brain networks, finding that the spontaneous transitions between networks of interacting brain areas are predictable. More specifically, the network activity is highly organised into a hierarchy of two distinct metastates, such that transitions are more probable within, than between, metastates. One of these metastates represents higher-order cognition, and the other the sensorimotor systems. Furthermore, the time spent in each metastate is subject specific, heritable, and relates to behaviour. Although evidence of non-random state transitions has been found at the microscale, this finding at the whole-brain level, together with its relation to behaviour, is a completely new finding and has wide implications regarding the cognitive role of large-scale resting-state networks.

Reserved for Publication Footnotes

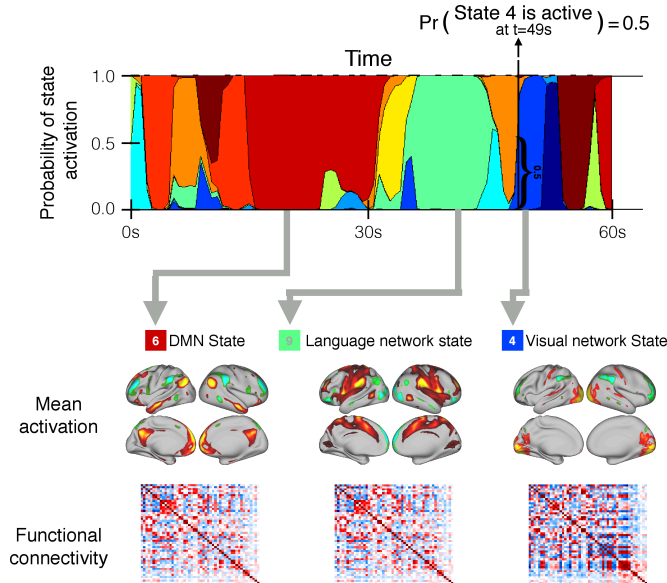


Fig. 1. The HMM estimates a number of brain networks (or states) that are common to all subjects, together with a specific state time course for each subject indicating when each state is active. The states are characterised by their mean activation and functional connectivity matrix. The top panel shows a 60 seconds section of the state time course for one example subject. Underneath, mean activation maps (projected from 50-dimensional ICA space to brain space) are shown for three out of the 12 inferred states, along with their corresponding functional connectivity matrices.

were temporally concatenated resulting in a single group data matrix from which a Hidden Markov Model (HMM) with 12 states was inferred (see Methods). The HMM states represent unique brain networks of distinct activity and functional connectivity (where functional connectivity is defined as Pearson correlation over time between brain areas) that repeat at different points in time (19, 20). Other approaches (typically based on sliding-windows estimations and subsequent clustering) have been proposed to find repeating states (21); the HMM, however, offers a probabilistic (generative) model that, through a single process of Bayesian inference (see Methods), models the time series in a self-contained manner. Here, whereas the states are estimated at the group level, each subject has a characteristic *state time course* that indicates the points in time at which that state is active. **Figure 1** shows an illustrative 1min section of the state time course for one example subject. Mean activation maps are shown for three out of the 12 inferred states, along with their corresponding functional connectivity matrices. The mean activity maps for all states are shown in **Figure S1**.

We used the temporal structure of the inferred brain states to reveal that the temporal organization of brain activity is, while stochastic, not completely random. In particular, we examined the HMM's estimate of the *transition probability matrix*, which specifies the probability of transitioning from any state to another (see Methods for mathematical details). **Figure 2a** shows the estimated transition probabilities in two different formats: a matrix and a graph where nodes represent states and arrow thickness represents (thresholded) state transition probability. Community discovery analysis using the transition probability matrix as an input (22) reveals two sets of states or *metastates* (coloured as blue and red), indicating that the brain tends to cycle between brain networks within a metastate, with less frequent transitions between them (the average number of metastate switches is ~30 per scanning session, and 96% of the sessions have at least one switch). To gain confidence in the result, we performed bootstrapped subject resampling (1000 draws of 820 subjects with

repetition) and applied the community discovery algorithm (22) on each bootstrap sample separately. We obtained the exact same separation in all bootstrap samples, confirming that the separation between these sets is highly significant. Further, it was found that the probability of transition between states heavily depends on how similar the brain networks in those states are (estimated using the correlation between the brain state functional connectivity matrices), but very little on the correlation between the brain state activation patterns, suggesting that changes at the level of connectivity are smoother than changes at the level of absolute activation (see **Figure S2**).

Time spent visiting distinct brain networks is not random

Another summary measure of the temporal characteristics of the states, or brain networks, is the proportion of time spent visiting each state. This is referred to as the *fractional occupancy (FO)* for each state and for each subject. This allows us to examine the FO correlations for each pair of states across subjects; that is, if matrix R refers to the FO for all subjects, such that $R_{i,k}$ is the proportion of time spent by subject i in state k , then the (no. of states by no. of states) FO correlation matrix C has elements $C_{k,l} = \text{corr}(R_{:,k}, R_{:,l})$, where $R_{:,k}$ represents the column vector of R that corresponds to state k . **Figure 2b** illustrates the result of correlating the FO of each pair of states across subjects, and performing hierarchical clustering (using Ward's algorithm) on the resulting correlation matrix. This analysis reveals between-state correlations that are extremely high (with absolute values over 0.5 and often close to 1.0), indicating an exceptionally clear hierarchical organisation of the brain dynamics, with the states divided into two sets, or *metastates*, corresponding to those in **Figure 2a**. More explicitly, here the metastates correspond to sets of states whose FOs are strongly correlated across subjects (red blocks in **Figure 2b**); conversely, the between-metastate FO correlations are negative (blue areas in **Figure 2b**). State 5, which was included in one of the metastates by the previous community detection analysis (in **Figure 2a**), is however uncorrelated to the rest of the states in terms of its FO. This state, with a relatively strong correlation with motion (see **Figure S5a**) and a much higher signal variance for all regions than the rest of the states (**Figure S5c**), may well be related to head movements and is thus left out of subsequent analyses. As we show in the **SI (Figure S3)**, the metastates organisation in the FO correlation matrix can largely be explained by between-subject differences of the Markovian dynamics encoded in the transition probability matrix. To illustrate the extent to which transition probabilities (though estimated in the HMM at the group-level) can vary across subjects, we recalculated the transition probability matrix for the 100 subjects with the highest occupancy for metastate 1 and the 100 subjects with the highest occupancy for metastate 2. These are shown in **Figure 2c**, depicting much larger transition probabilities from metastate 2 to metastate 1 in the first group, and the contrary in the second group.

The emerging metastates have distinctive spatial features

Figure 3a shows, for each metastate, a spatial map of the average *absolute* amplitude within the metastate. Regions with high values in these maps can be interpreted as having brain activity that deviates substantially from their average activity levels, during the time the metastate is active. **Figure 3b** shows a spatial map of the functional connectedness (or degree) of each brain area, computed as the sum of the *absolute* value of functional connectivity with the rest of the brain. Both the amplitude and connectivity reflect a clear functional distinction between the first metastate, composed of sensory (somatic, visual and auditory) and motor regions, and the second metastate, covering higher-order cognitive regions that include the default mode network (DMN), language and prefrontal areas. The correlation between the absolute amplitude and the connectedness maps is significantly positive ($r=0.52$, $p\text{-value}<0.001$).

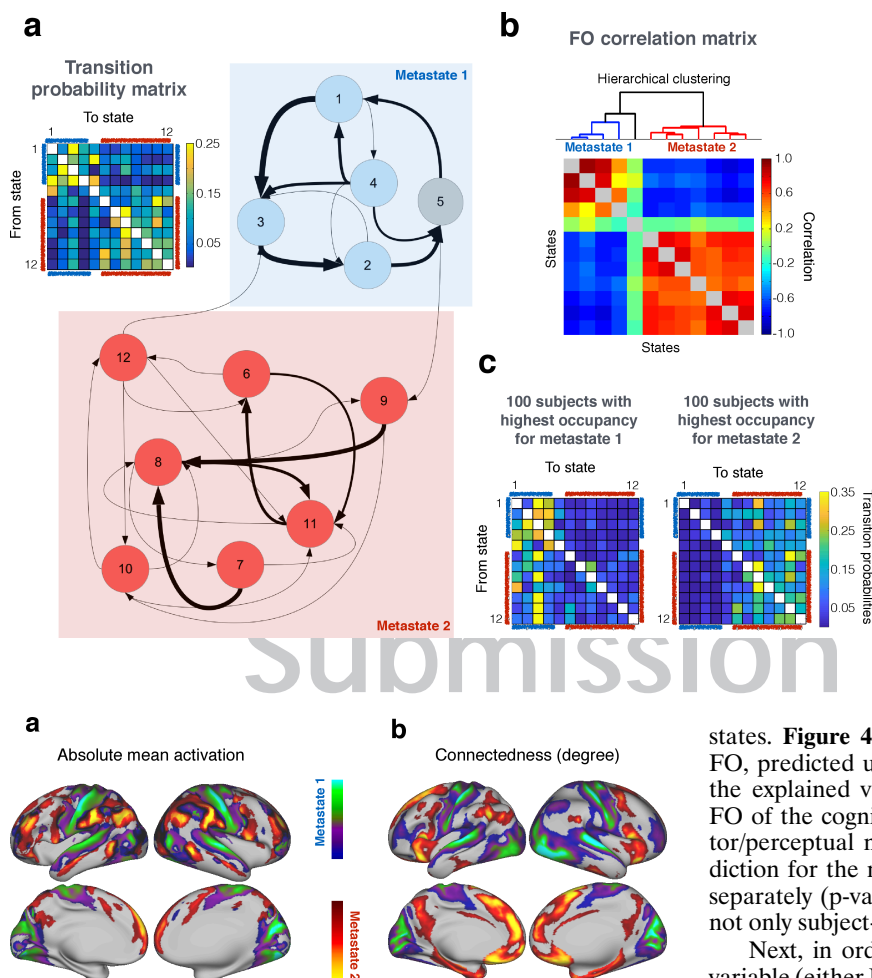


Fig. 3. The two metastates contain distinctive functional areas. Whereas the first metastate is associated with sensory (somatosensory, visual and auditory) and motor regions, the second metastate involves areas related to higher-order cognition (including regions of the default mode network, language and extensive prefrontal). This is apparent in both the activation level and connectivity. In the first case (a), we look at the average absolute amplitude of each region within the metastate, which can be interpreted as a measure of the amount of deviation from average activity levels. In the second case (b), we compute the connectedness, or degree, defined as the sum of functional connectivity of each region with the rest of the brain.

Time spent visiting brain networks and metastate predicts behavioural traits

If the temporal organisation of the brain network dynamics is functionally relevant, then it should link to behaviour. We next assessed how strongly the time spent (i.e., FO) in the states or metastates can be predicted by subject-specific behavioural traits. Motion, sex, age, body mass index, and a measure of surface registration differences (using the surface warping Jacobian obtained from the Multimodal Surface Matching process, and taking the mean of the absolute logarithm of the area ratio images) were regressed out as confounds from all the considered variables (FO and traits). Furthermore, we discarded the 10% of subjects with the highest measures of motion. Then, using (Bayesian) Partial Least Squares combined with permutation testing (23), the relationship between FO and behavioural traits was found to be statistically significant (p -value = 0.003, see Methods).

We then questioned whether the metastates' FO would be better predicted by the behavioural traits than the individual states' FO, as this would show that the metastates are behaviourally relevant above and beyond the underlying brain

Fig. 2. The transitions between brain networks are not random, giving rise to two very distinct sets of states, which are referred to as metastates. (a) The transition probability matrix (top left) indicates the probability of transitioning from any state to another, showing that some transitions are much more likely than others. This is apparent when shown in graph format, where the nodes represent brain states, and the thickness of the arrows represents the state transition probability (transitions were thresholded for readability). (b) The fractional occupancy matrix, which contains the total time spent in each state per subject, across subjects, exhibits exceptionally strong correlations between states. Even more strongly than the transition probability matrix, these correlations indicate a clear hierarchical metastate structure. Hierarchical clustering (illustrated above the fractional occupancy correlation matrix) confirms this result. (c) Transition probability matrix for the 100 subjects with the highest occupancy for metastate 1 and the 100 subjects with the highest occupancy for metastate 2.

states. **Figure 4a** shows the explained variance of each state's FO, predicted using cross-validation (see Methods), along with the explained variance of the *metastate profile*, defined as the FO of the cognitive metastate minus the FO of the sensorimotor/perceptual metastate. This analysis illustrates that the prediction for the metastate is higher than any of the single states separately (p -value < 0.001), suggesting that the metastates are not only subject-specific but also relevant to behaviour.

Next, in order to assess the individual significance of each variable (either FO or trait), we use canonical correlation analysis (CCA) (24). Motion, sex, age, body mass index and registration differences (see above) were again regressed out as confounds from all the considered variables (FO and traits). We then looked at the individual correlation of each behavioural trait or each state's FO against the opposite canonical covariate (i.e. the canonical covariate of the state FO was used for the behavioural traits, and vice versa; see Methods). **Figure 4b** shows the behavioural traits and brain states separated into two groups according to the sign of the correlation, where the four coloured boxes (red and blue) indicate statistical significance (permutation testing, significance level of 0.05). As observed, there is a tendency of the positive traits to correlate with the FO of the higher-order cognitive metastate.

Time spent in each metastate is a subject-specific, heritable measure

Given that we have seen that the proportion of time-spent in each metastate is associated with behavioural traits (**Figure 4ab**), we could expect them to also be very subject-specific measures. In order to assess this, we performed cross-validated predictions of the FO for each session, either for the states or the metastates, as a function of the FO of the other three sessions for the same subject. **Figure 4c** shows that, despite the fact that the subjects were scanned on different days, the metastate FOs are consistent across sessions and can be reliably predicted.

It has been recently reported that functional connectivity within and across large-scale brain networks has a heritable component (25, 26). Here, we found that both the states' and the metastates' FO are also highly heritable. The dataset contains twin structure, including a combination of non-identical twins and non-twin siblings. That allowed us to compare whether related subjects had a more similar metastate profile than non-related

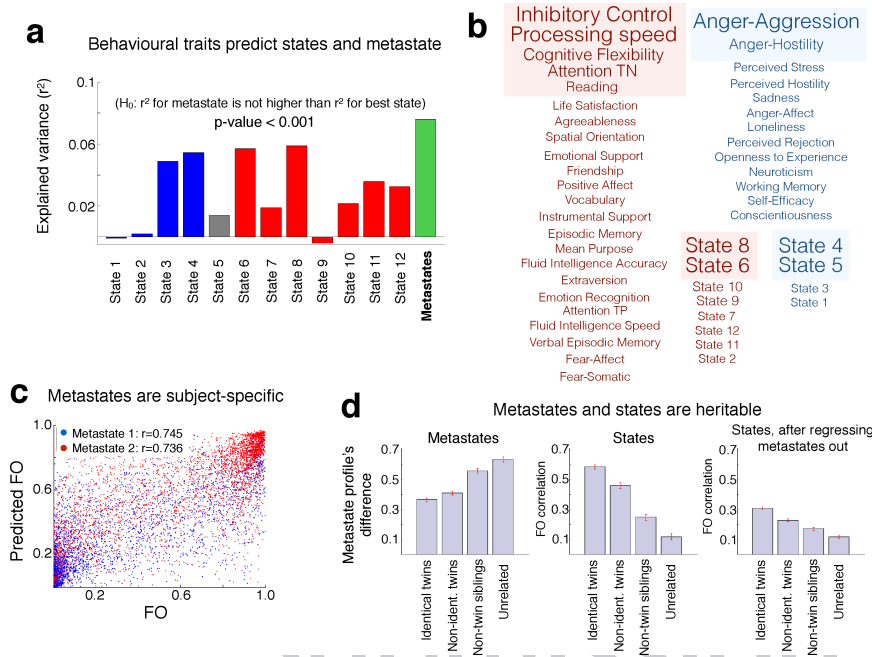


Fig. 4. The brain network dynamics and the metastate profile are subject-specific, relate to behaviour and are heritable. (a) Using behavioural traits (wellbeing, intelligence and personality) as regressors, we can significantly predict the states' fractional occupancy and (even more accurately) the metastate profile. (b) The metastate profile separates "positive" from "negative" traits, which correlate to the metastate profile with opposite signs; statistical significance of this correlation is highlighted (variables correlate when they have the same colour and anticorrelate otherwise) and the font size corresponds to the magnitude of the correlation. (c) The fractional occupancy of the metastates for a given session can be accurately predicted for each subject using information from the other sessions, suggesting that the metastate profile is a very specific subject fingerprint. (d) The metastate (left) and state (centre) distributions are strongly heritable; the state distribution is heritable (but to a lesser extent) if we regress out the metastates information (right).

subjects. For each pair of subjects, we computed the absolute difference of their metastate profile values, gathering a distribution of differences for each type of familiar relation. **Figure 4d** (left) shows that identical twins, followed by non-identical twins and then non-twin siblings, all have a closer metastate profile than unrelated individuals. We also calculated the correlation of the state FO (a vector of 12 elements per subject) for each pair of individuals. **Figure 4d** (middle) reveals a huge difference between unrelated pairs of subjects and pairs of subjects that are twins or even non-twin siblings. We then questioned whether these differences in the state FO were solely due to the differences in the metastate profile. To answer this, we regressed the metastate profile out of the FO and repeated the analysis. **Figure 4d** (right) indicates that, whereas much of the differences were actually due to the heritability of the metastate profile, the specific states FO contain further information that is also heritable

Are the metastates related to motion and/or sleep?

One possible concern about the cognitive interpretation of the metastates is whether they are driven by motion. However, the correlation between head motion (measured as the session average) and the metastate profile is relatively low ($r=0.05$) in comparison with other subject measures (**Figure S4a**). Also, the correlation between motion and the metastate profile is very low in comparison with the individual state correlations (**Figure S4b**). In particular, state 5 (the one associated with the highest signal variance – see **Figure S4c**) has the highest correlation with motion in absolute value, followed by state 6, which corresponds to the DMN and anticorrelates with motion significantly. Furthermore, if we regress motion out of the FO matrix and then repeat the calculation of the correlation between the states' FO across subjects (as in **Figure 2b**), the two metastates (and the unrelated state 5) are still evident - see **Figure S4d**. Therefore, although there seems to be certain influence of motion on the individual states, the metastate distribution appears to be relatively free of motion confounds.

If we consider the metastates to represent a high-level state of cognition, then it is important to consider how the metastates relate to *macrostates* such as sleep. As reported by Tagliazucchi and Laufs (27), many subjects transit from wakefulness to sleep stages N1 and N2 during the first ten minutes of the resting-state

session, and that is certainly a major cause of genuine neural activity change over time. Moreover, if we divide the 15min sessions in four blocks, the occurrence of the metastates varies significantly between the four blocks, with an increase in the occurrence of the sensorimotor/perceptual metastate as the session progresses (**Figure S5a**).

Hence, do the sensorimotor/perceptual and the cognitive metastates simply represent sleep and wakefulness, respectively? Although it is possible that the sensorimotor/perceptual metastate is higher in states of drowsiness and sleep (given **Figure S5a** and the results from (27)), we argue that there is more information in the metastates than just sleep, and that both metastates coexist in pure wakefulness. First, the HCP resting-state protocol corresponds to subjects having their eyes open with fixation, which reduces the probability of falling asleep. Second, the amount of time spent on average in the sensorimotor/perceptual metastate during the first minute of scanning (when sleep is highly unlikely) is around 20s (**Figure S5b**). Third, if we look at the state dwell times of the metastates (after removing metastate visits shorter than 4s to rule out noisy, artefactual transitions), **Figure S5c** shows that most of the transitions are too short to capture sleep or wakefulness periods (for example, around 50% of the metastate visits are shorter than 15s even when ignoring the shortest visits). Finally, if the metastates are primarily capturing sleep vs. wakefulness, we might expect the metastate profile to correlate highly with the Pittsburgh Sleep Quality Assessment (PSQI) score provided within the HCP set of variables, which essentially measures the quality and pattern of sleep of the subjects. The PSQI is however uncorrelated to the metastate profile ($r=-0.0056$). **Figure S5d** shows that this correlation is lower in absolute value than the majority of the behavioural variables considered in this study. Also, as illustrated in **Figure S5e**, this correlation is also lower than most of the states when considered separately. Altogether, our results suggest that the metastates organization goes beyond separating wakefulness and sleep, and arguably represents a more general and ubiquitous temporal pattern in spontaneous brain activity.

Discussion

Robust patterns of synchronized activity can be measured in the brain in both task and rest. The spatial organisation of these patterns at rest has been well-studied. In this paper, we focus on the temporal organisation of the dynamics of the resting state networks, and discover a temporal hierarchy in which brain networks are organised into two distinct sets, or metastates. These metastates have a clear functional separation (cognitive versus sensorimotor/perceptual) that is consistent across subjects, relate to behaviour, is heritable, and reproducible (see SI and **Figure S6**). Although other work has reported hierarchical aspects of the fMRI resting-state networks (28, 29, 30), these analyses assume the existence of a hierarchy in the first place, whereas the HMM neither imposes nor encourages it; that is, if a hierarchical structure arises from the HMM results, then this is purely driven by the structure in the data. More importantly, whereas previous work extracts the hierarchical structure from the zero-lag correlations of the (voxels' or ICA components') time series, we are here referring to the hierarchical structure of the *network dynamics*. This hierarchy manifests in two ways. First, in the hierarchical structure of the transition probability matrix; that is, the hierarchy emerges not from correlations but from transitions. Second, in the FO correlation across subjects; that is, it is not the correlation of the time series within subjects but the correlation between the (subject-specific) *averaged state probabilities* (of which the FO is an estimation) across subjects.

These results can be related to current theoretical perspectives on cognition. First, the hierarchy of brain networks (in the form of the metastates) and, more generally, the non-random transitions between the states, naturally imply that current and previous brain states can constrain future states in task-free conditions. This relates, and can potentially complement, current results on the persistence of activity in specific areas of the brain (31) to the whole-brain. Second, there is some overlap between the metastates and the two extremes of the principal axis of structural connectivity variation estimated by Margulies and colleagues (32), which is based on a mathematical decomposition of the structural connectivity elicited from the same 820 subjects used in this study. This coincidence of structural and functional hierarchies is meaningful from a theoretical perspective and deserves further investigation, perhaps also at other hierarchical levels. Third, from an evolutionary psychology perspective, it is interesting to observe that FO of the higher-order cognition metastate is in general positively correlated with positive traits, and so may indicate an adaptive value of intrinsically generated mental states (33).

Leveraging the common framework provided by the HMM to integrate different data modalities, future work will aim to investigate whether or not the same hierarchical structure and metastates can also be found in electrophysiological data such as magnetoencephalography and electroencephalography. To this end, we will use an observation model specifically tailored to this type of data (34).

Alternative representations of the data

It is worth noting that we are not claiming the HMM to be the ground truth or explain all aspects of the data, but to offer a useful perspective on the data. Other models and techniques, such as multivariate autoregressive modelling (35) or lag-based activity propagation approaches (36, 37) offer their own complementary and useful descriptions of the dynamics in the data. In this work, the HMM has been able to reveal a strong hierarchical structure in the data, represented by metastates. While it is possible that the same (or related) phenomena might be detected using other approaches, it is unclear how this would be achieved in practice. An advantage of the HMM over other techniques is that the HMM explicitly parametrises time-dependent information (in the form of the state time courses), such that it can be more easily

used to access network dynamics, for example in task responses (20). In the SI, we explore the relationship between the HMM and both the multivariate autoregressive model and lag-based activity propagation analysis in more detail.

Data stationarity

The HMM assumes that, within each state, the data is drawn from a single given HMM observation model, which is defined by a single set of parameters. Although more general statistical definitions of stationarity exist (38), we here refer to stationarity *with respect to* the observation model used by the HMM, which in this case is a multivariate Gaussian distribution. That is, we refer to data as being non-stationary when it has a time-varying mean and/or covariance within a given session. Now, the facts that (i) the HMM needs more than one state to model the data and does not collapse into a single state, and that (ii) multiple states are inferred with significantly different observation model parameters, suggest that the data is non-stationary in this sense. Further evidence supporting non-stationary can be gathered from the synthetic data used in the SI (see section 'Using the HMM to assess non-stationarity').

In relation to this question, recent studies have reasonably questioned whether observed changes in dynamical functional connectivity in resting-state fMRI are due to genuine brain transitions or, rather, are mostly explained by sampling variability (14, 15, 16, 17). However, these valid concerns are specific to techniques that measure dynamic function connectivity using sliding windows. The HMM, however, can bypass the problem of sampling variability by using, for each state, the entire set of subjects in order to provide an estimation of functional connectivity (see SI Methods). The amount of data used to infer the network characteristics for each brain state is thus orders of magnitude larger (on average, 820 subjects x 4 sessions x 1200 time points per session, divided by 12 states) than that which a sliding-window can possibly encompass, especially given the size of the data set used in this study. At the same time, and unlike long sliding windows, the model is able to capture quick changes in brain activity.

Conclusion

Altogether, our approach reveals an intriguing property in the resting-state temporal dynamics: the brain transitions between networks in a manner that is stochastic yet not completely random. More specifically, there is a hierarchical organization of the brain networks into two major metastates, whose functional attributes are well separated, such that one metastate covers regions that correspond to the sensorimotor systems and perception, whereas the other relates to high cognitive functions. The evidence of temporal organisation of large-scale network states in the resting brain adds to previous evidence at the microscale in both task (5) and rest (6), suggesting that this is a general characteristic of activity in human cognition that exists at a wide range of spatio-temporal scales. This new perspective, and the methods used to provide it, paves the way for future investigations into the cognitive role of the temporal dependency of brain network states, for example, in short-term memory and learning. Further, it contributes to the important debate about the feasibility of assessing dynamic functional connectivity in resting-state fMRI by showing the relation of our estimations to behaviour and heritability.

Methods

Data and preprocessing

We used resting-state fMRI data from N=820 subjects from the HCP, which provides the required ethics and consent needed for study and dissemination. These are all subjects with complete resting fMRI data from the 900-subject public data release, all healthy adults (ages 22–35 years, 453 females) scanned on a 3-T Siemens connectome-Skyra. For each subject, four 15-min runs of fMRI time series data with temporal resolution 0.73 s and spatial resolution 2-mm isotropic were available. The preprocessing pipeline followed (18, 39), and thus will be described only briefly here. Spatial preprocessing was applied using the procedure described in (40). After structured artefact removal using independent component analysis (ICA)

followed by FSL's FIX (41), which removed more than 99% of the artefactual ICA components in the data set, we used group spatial-ICA to obtain a "parcellation" of 50 components that covers both the cortical surfaces and the subcortical areas. We did not use global signal regression. Then, we used this parcellation to project the fMRI data into 50-dimensional time series. Such time series, of size (number of participants x number of scans x number of time points x number of ICA components = 820 x 4 x 1200 x 50), were finally standardised so that, for each scan, subject and ICA component, the data have mean 0 and standard deviation 1.

Hidden Markov Modelling

The Hidden Markov model (HMM) assumes that the time series data can be described using a sequence of a finite number of states. Each state is here represented by a multivariate Gaussian distribution, which is described by mean and covariance. The HMM is inferred using the (publicly available) *HMM-MAR* toolbox, which provides estimates of: the parameters of the state distributions, the (group-level) transition probability matrix, and the probabilities of each state to be active at each time point (20, 42). The FOs are computed as the aggregations of these probabilities for each subject. For full details see SI Materials and Methods.

Estimation and statistical testing on the metastates

The presence of the metastates were investigated in two different ways. First, using the Louvain community detection algorithm (22) on the transition

probability matrix. This method aims to find communities or nodes (here, metastates) in a graph (here, a directed, weighted graph representing the transition probability matrix) such that the connectivity between the nodes (here, states) within a community is strong with respect to the connectivity across communities. Second, we looked at the structure of the FOs correlation matrix, which contains the correlation for each pair of states' FO across subjects. For full details of this and the statistical testing with behaviour and heritability, see SI Materials and Methods.

Acknowledgements: We thank the WU-Minn HCP Consortium for generating and making publicly available the HCP data. DV is supported by a Wellcome Trust Strategic Award (098369/Z/12/Z). 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). We thank Janine Bijsterbosch for her valuable help with preprocessing details and James M. Shine for his useful comments on the manuscript.

Author contributions: DV, SMS and MWW designed the experiment. SMS preprocessed the data. DV performed the analysis and created the figures. DV and MWW wrote the manuscript, and all authors critically revised and approved it.

1. Damoiseaux J.S. et al. (2006). Consistent resting-state networks across healthy participants. *Proceedings of the National Academy of Sciences of the USA* **103**, 13848-13853.
2. Luca M., Beckmann C.F., De Stefano N., Matthews P.M., and Smith S.M. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage* **29**, 1359-1367.
3. Raichle M.E., et al. (2001) A default mode of brain function. *Proceedings of the National Academy of Sciences of the USA* **98**, 676-682.
4. Smith S.M., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences of the USA* **106**, 13040-13045.
5. Morcos A.S., and Harvey C.D. (2016) History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience* **19**, 1672-1681.
6. Berkes P., Orban G., Lengyel M., and Fiser J. (2011). Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science* **331**, 83-87.
7. Fox M.D., et al. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the USA* **27**, 9673-9678.
8. He B (2011). Scale-Free Properties of the Functional Magnetic Resonance Imaging Signal during Rest and Task. *Journal of Neuroscience* **39**, 13786-13795.
9. Zalesky A., Fornito A., Cocchi L., Gollo L.L., and Breakspear M. (2014). Time-resolved resting-state brain networks. *Proceedings of the National Academy of Sciences of the USA* **111**, 10341-10346.
10. Shine J.M., et al. (2016). The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Function. *Neuron* **92**, 544-554.
11. Ito J., Nikolaev A.R., and van Leeuwen C. (2007). Dynamics of spontaneous transitions between global brain states. *Human Brain Mapping* **28**, 904-913.
12. Kelso J.A. (2012). Multistability, and metastability: understanding dynamical coordination in the brain. *Philosophical Transactions of the Royal Society B* **367**, 906-918.
13. Watanabe T., et al. (2014). Energy landscapes of resting-state brain networks. *Frontiers in Neuroinformatics* **8**, 1-11.
14. Leonardi N., and van de Ville D. (2014). On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage* **104**, 430-436.
15. Zalesky A., and Breakspear M. (2015). Towards a statistical test for functional connectivity dynamics. *NeuroImage* **114**, 466-470.
16. Hindriks R., et al. (2016). Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *NeuroImage* **127**, 242-256.
17. Laumann T.O., et al. (2016). On the Stability of BOLD fMRI Correlations. *Cerebral Cortex*, 1-14.
18. Smith S.M. et al. (2013) Resting-state fMRI in the Human Connectome Project. *NeuroImage* **80**, 144-168.
19. Baker P., et al. (2014). Fast transient networks in spontaneous human brain activity. *eLife* **3**: e01867.
20. Vidaurre D., et al. (2016). Spectrally resolved fast transient brain states in electrophysiological data. *NeuroImage* **126**, 81-95.
21. Calhoun V.D., Miller R., Pearson G., and Adali T. (2016). The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* **84**, 262-274.
22. Blondel V.D., Guillaume J.L., Lambiotte R., and Lefebvre E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **10**, P10008 (2008).
23. Vidaurre D., van Gerven M.A.J., Biezla C., Larrañaga P., and Heskes T. (2013). Bayesian sparse partial least squares. *Neural Computation* **25**, 3318-3339.
24. Smith S.M., et al. (2015). A positive-negative mode of population covariation links brain connectivity, demographics and behaviour. *Nature Neuroscience* **18**, 1565-1567.
25. Ge T., Holmes A.J., Buckner R.L., Smoller J.W., and Sabuncu M.R. (2017). Heritability analysis with repeat measurements and its application to resting-state functional connectivity. *Proceedings of the National Academy of Sciences of the USA* **114**, 5521-5526.
26. Colclough, G.L., Smith, S.M., Nichols, T.E., Winkler, A.M., Sotiropoulos, S.N., Glasser, M.F., Van Essen, D.C., and Woolrich, M.W. (2017). The heritability of multi-modal connectivity in human brain activity. *eLife*, In press.
27. Tagliazucchi E., and Laufs H. (2014). Decoding Wakefulness Levels from Typical fMRI Resting State Data Reveals Reliable Drifts between Wakefulness and Sleep. *Neuron* **82**, 695-708.
28. Cordes D., Haughton V., Carew J.D., Arfanakis K., and Maravilla K. (2002). Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magnetic Resonance Imaging* **20**, 305-317.
29. Gleiser P.M., and Spormaker V.I. (2010). Modelling hierarchical structure in functional brain networks. *Philosophical Transactions of the Royal Society A* **368**, 5633-5644.
30. Doucet G., et al. (2011). Brain activity at rest: a multiscale hierarchical functional organization. *Journal of Neurophysiology* **105**, 2753-2763.
31. Curtis C.E., and D'Esposito M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences* **7**, 415-423.
32. Margulies D.S. et al. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences of the USA* **113**, 12574-12579.
33. Smallwood J. (2013). Distinguishing how from why the mind wanders: a process-occurrence framework for self-generated mental activity. *Psychological Bulletin* **139**, 519-535.
34. Vidaurre D., et al. (2017) Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. *BioRxiv*. doi: <https://doi.org/10.1101/150607>.
35. Marple S.L. (1986). Digital Spectral Analysis: with applications. Prentice-Hall Series in Signal processing. (Upper Saddle River, NJ, USA).
36. Mitra A., Snyder A.Z., Hacker C.D., and Raichle M.E. (2014). Lag structure in resting-state fMRI. *Journal of Neurophysiology* **111**, 2374-2391.
37. Mitra, A., and Raichle, M.E. (2016). How networks communicate: propagation patterns in spontaneous brain activity. *Philosophical Transactions of the Royal Society B* **371**.
38. Liegeois R., Laumann T., Snyder A.B., Zhou H.J., and Yeo B.T. (2017). Interpreting Temporal Fluctuations in Resting-State Functional Connectivity MRI. *BioRxiv*, doi: <http://dx.doi.org/10.1101/135681>.
39. Smith S.M., et al. (2013). Functional connectomics from resting-state fMRI. *Trends in Cognitive Sciences* **17**, 666-682.
40. Glasser M.F., et al. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* **80**, 105-124.
41. Griffanti L., et al. (2014). ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage* **95**, 232-247.
42. Bishop C.M. (2006). Pattern Recognition and Machine Learning. Springer New York (New York, USA).
43. Kringelbach M.L., and Berridge K.C. (2009). Towards a functional neuroanatomy of pleasure and happiness. *Trends in Cognitive Sciences* **13**, 479-487.
44. Deary I.J., Penke L., and Johnson W. (2010) The neuroscience of human intelligence differences. *Nature Reviews Neuroscience* **11**, 201-211.
45. Eisenberger N. I., Lieberman M.D., and Satpute A.B. (2005). Personality from a controlled processing perspective: an fMRI study of neuroticism, extraversion, and self-consciousness. *Cognitive, Affective & Behavioral Neuroscience* **5**, 169-181.
46. Vidaurre D., et al. (2017). Discovering dynamic brain networks from Big Data in rest and task. *NeuroImage*. In press.
47. Winkler A., Webster M.A., Vidaurre D., Nichols T.E., and Smith S.M. (2015). Multi-level block permutation. *NeuroImage* **123**, 253-268.
48. Linkenkaer-Hansen K., Nikouline V.V., Palva J.M., and Ilmoniemi R.J. (2001). Long-range temporal correlations and scaling behavior in human brain oscillations. *Journal of Neuroscience* **21**, 1370-1377.
49. Woolrich M.W., Behrens T.E.J., and Smith S.M. (2004). Constrained linear basis sets for HRF modelling using Variational Bayes. *NeuroImage* **21**, 1748-1761.
- 50.