



HHS Public Access

Author manuscript

Neuroimage. Author manuscript; available in PMC 2018 August 15.

Published in final edited form as:

Neuroimage. 2017 August 15; 157: 364–380. doi:10.1016/j.neuroimage.2017.05.067.

The energy landscape underpinning module dynamics in the human brain connectome[★]

Arian Ashourvan^{a,b}, Shi Gu^{a,c}, Marcelo G. Mattar^{a,d}, Jean M. Vettel^{a,b,f}, and Danielle S. Bassett^{a,e,*}

^aDepartment of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA

^bU.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005, USA

^cApplied Mathematics and Computational Science Graduate Program, University of Pennsylvania, Philadelphia, PA 19104, USA

^dDepartment of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA

^eDepartment of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA

^fDepartment of Psychological & Brain Sciences, University of California, Santa Barbara, CA 93106, USA

Abstract

Human brain dynamics can be viewed through the lens of statistical mechanics, where neurophysiological activity evolves around and between local attractors representing mental states. Many physically-inspired models of these dynamics define brain states based on instantaneous measurements of regional activity. Yet, recent work in *network neuroscience* has provided evidence that the brain might also be well-characterized by time-varying states composed of locally coherent activity or functional modules. We study this network-based notion of brain state to understand how functional modules dynamically interact with one another to perform cognitive functions. We estimate the functional relationships between regions of interest (ROIs) by fitting a pairwise maximum entropy model to each ROI's pattern of allegiance to functional modules. This process uses an information theoretic notion of energy (as opposed to a metabolic one) to produce an energy landscape in which local minima represent attractor states characterized by specific patterns of modular structure. The clustering of local minima highlights three classes of ROIs with similar patterns of allegiance to community states. Visual, attention, sensorimotor, and subcortical ROIs are well-characterized by a single functional community. The remaining ROIs affiliate with a

[★]A.A., J.M.V., and D.S.B developed the project. M.G.M. acquired the data. A.A. analyzed the data. S.G. contributed computational tools and expertise. A.A., J.M.V., and D.S.B. wrote the paper.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

*Corresponding author at: Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA., dsb@seas.upenn.edu (D.S. Bassett).

Supplementary information

See SI for complementary analyses, results, and discussions that were omitted from the manuscript for brevity. Information regarding the goodness of fit of the MEM, additional information on asymmetries of basin transition frequencies and probabilities, spatial maps of ICN components, calculation of band-passed wavelet-coherence, and alternative parcellation methods are discussed in depth in SI 1,2,3,4, and 5 respectively.

putative executive control community or a putative default mode and salience community. We simulate the brain's dynamic transitions between these community states using a random walk process. We observe that simulated transition probabilities between basins are statistically consistent with empirically observed transitions in resting state fMRI data. These results offer a view of the brain as a dynamical system that transitions between basins of attraction characterized by coherent activity in groups of brain regions, and that the strength of these attractors depends on the ongoing cognitive computations.

Keywords

Energy landscape; Maximum entropy model; Community structure; Modularity; Functional brain network; Graph theory

Introduction

The human brain is a complex dynamical system comprised of billions of neurons that continuously communicate with one another. Although the vast number of processing units challenges exact prediction of single neuron activity, recently developed statistical models reveal a characteristic meso-scale structure whereby sets of larger-scale brain regions display coherent activity at rest. These sets form putative functional modules characterized by locally dense functional connectivity, and include the default mode, salience, attention, fronto-parietal, cingulo-opercular, motor, visual, auditory, and subcortical systems (Salvador et al., 2005; Meunier et al., 2009; Yeo et al., 2011; Power et al., 2011). Interestingly, although within-module functional connectivity is in general higher than between-module functional connectivity, these patterns fluctuate dynamically over short periods of time (Ma et al., 2014; Kiviniemi et al., 2011; Watanabe et al., 2013), both at rest and during task performance (Cole et al., 2014; Mattar et al., 2015; Bassett et al., 2011, 2013b, 2015; Braun et al., 2015).

The existence of functional modules – as measured using graph theory – was first reported nearly a decade ago (Salvador et al., 2005; Meunier et al., 2009), and these initial reports were swiftly corroborated by studies using complementary methodologies ranging from co-activation analysis to independent components analysis. A natural question following the observation of these modules was “What do they do? And how are they recruited as we go through life performing a variety of functions?” To address these questions, dynamic community detection methods were developed and applied to neuroimaging data, revealing the fact that modules reconfigure in support of working memory (Braun et al., 2015, 2016), reinforcement learning (Gerraty et al., 2016), visuo-motor learning (Bassett et al., 2011, 2013b, 2015), and linguistic processing (Chai et al., 2017; Doron et al., 2012a). Module reconfiguration at rest has also been reported as a marker of aging and development (Betzel et al., 2015). These studies collectively demonstrate that module reconfiguration is a hallmark of healthy brain function (Telesford et al., 2016), and recent evidence suggests that it is a marker that is altered in psychiatric disease, even providing an intermediate phenotype of schizophrenia (Braun et al., 2016).

Despite the extensive work demonstrating the presence and cognitive utility of dynamic module reconfiguration, fundamental insights into the mechanisms or rules by which modules interact with one another over time have remained elusive (Mattar et al., 2015; Khambhati et al., 2016). Evidence suggests that module reconfiguration may be driven by glutamatergic signaling (Braun et al., 2016), affect and arousal (Betzel et al., 2016a), and may provide a substrate for cognitive control (Khambhati et al., 2016), supporting a delicate balance between domain-general and domain-specific function (Fedorenko and Thompson-Schill, 2014a). Although these observations support the biophysical relevance of the phenotype, they do not provide computational theories for its existence. Indeed, the development of a theory explaining and predicting module dynamics would be particularly important in informing studies of its cognitive specificity, transmitter-level drivers, and role in psychiatric disease.

One potential route towards a mechanistic theory of brain network dynamics is to consider probabilistic models that were originally developed in the field of statistical mechanics. Pair-wise maximum entropy models (MEM), for example, have proven very useful in estimating and predicting spiking activity in neurons (Shlens et al., 2006), local field potentials from neuronal assemblies (Tang et al., 2008), and blood oxygen level dependent signals (BOLD) from brain regions using functional magnetic resonance imaging (fMRI) (Watanabe et al., 2013, 2014b, 2014c). When a pair-wise MEM accurately fits empirical data, it implies that the observed activation pattern can be described as a combination of each unit's independent activation rate plus the units' joint activation rates. When a pair-wise MEM does not accurately fit empirical data, it implies that higher order interactions (such as triplets) or nonlinearities contribute to the observed dynamics. Importantly, pair-wise MEMs can be used to infer an *energy* landscape of brain activity during task performance. Here the term energy is used in an information theoretic sense, and it does not have a direct relation to metabolic energy or to a formal Hamiltonian. The energy landscape is characterized by basins of attraction representing common brain states, as well as the paths or trajectories along which the brain moves as it transitions from one basin to another. The architecture of this energy landscape has proven useful in predicting individual differences in human perception and behavior (Watanabe et al., 2014c).

In traditional applications of MEMs to neurophysiological data, a brain state is defined as a pattern of activity across brain regions (or similarly, a neural state is defined as a pattern of activity across neurons). However, these notions of brain state are agnostic to the patterns of communication or synchronization linking brain regions, and therefore are unable to address the question of how one pattern of coherent activity could evolve into another pattern of coherent activity. To address this question, we explicitly define a *network* state as the pattern of module allegiance across brain regions, and we use this definition to examine transitions between network states. We test the hypothesis that mesoscale structure in dynamic functional connectivity patterns is well explained by pairwise interactions between communities. The null hypothesis that we seek to reject is that such patterns cannot be explained without contributions from higher-order interactions between communities. To test this hypothesis, we construct a time-dependent network by linking 10 regions of interest by the low frequency (0.06–0.19 Hz) wavelet coherence between their time series in a given time window. We use a community detection algorithm to identify groups of brain regions

that show stronger coherence with one another than they do to other groups. We refer to these groups as network communities, and we fit the MEM to each ROI's time series of the state of co-occurrence in the same community with other ROIs. This approach enables us to identify network states that form local energy minima, as well as features of the energy landscape surrounding these minima. More generally, this approach highlights the dynamic functional roles that different ROIs play in network states and the transitions between them.

Our results reveal the presence of local minima on the energy landscape, many of which are characterized by the activation of a single community. Interestingly, different ROIs show different patterns of membership to these single community states. Visual, attention, sensorimotor, and subcortical ROIs tend to form a single functional community (Class-I). The remaining ROIs form the putative executive control network (Class-II) and the putative default mode and salience network (Class-III). To further study these dynamics, we modeled the transitions of single community states over the landscape of the states's energy via a random walk process. Our numerical simulations of basin transitions using an MCMC random walk predict empirical frequencies and probabilities of state transitions with high fidelity for Class-I and Class-III ROIs, and with lower fidelity for executive control (Class-II) ROIs. In addition, empirically the executive control ROIs also display higher entropy energy landscapes, linking diverse state classes, and utilizing uniform transition probabilities across basins. These features support the unique role of executive control regions in diversifying the brain's dynamic functional repertoire across many cognitive processes via their rich and flexible dynamic functional fingerprint.

Results

Distillation of drivers of resting state dynamics

Maximum entropy models are optimally constructed to fit patterns of interactions between *relatively few* brain regions. We therefore sought to distill the drivers of resting state dynamics to a few well-chosen regions of interest. Specifically, in resting state fMRI data acquired from 20 healthy adult individuals in a multiband imaging sequence, we extract 10 regions of interest in a data-driven fashion as centroids of independent components (see Methods). These regions include the cuneus, precuneus, precentral gyrus, caudate, right and left rostral middle frontal cortex, dorsomedial prefrontal cortex, medial orbitofrontal cortex, and *pars triangularis* (see Table 1). We use these regions as proxies of their respective cognitive systems, spanning visual, dorsal attention, sensorimotor, basal ganglia, executive control, dorsomedial prefrontal cortex, default mode, and salience systems, respectively. For explicit maps of each independent component, and the representative region chosen, see SI3.

Maximum entropy model of network states

Our goal is to understand how the brain transitions between network states. We focus our attention on the transitions characterized by changes in the community structure of the network, or the organization of putative functional modules. This focus is motivated by a growing literature demonstrating (i) the presence of network communities at rest, which map on to known cognitive systems (Salvador et al., 2005; Meunier et al., 2009; Power et al., 2011; Yeo et al., 2011), and (ii) changes in the integration or segregation of these

communities during task performance (Bassett et al., 2011, 2013b, 2015; Cole et al., 2014; Braun et al., 2015). Based on these emerging lines of research, we define network states based on regions' congruent community allegiance. Specifically, we fit a MEM to the binarized community allegiance probability of region i where ($i = 1, 2, \dots, N$) with the other 9 ROIs. This approach assesses the community-based interactions between a single ROI and all others, thus significantly reducing the space of possible states from the total number of partitions of an n -element set (the Bell number 115975) to only 2^{N-1} (=512) local single community states, and thereby increasing model accuracy (Grimaldi, 2006).

In SII we discuss the goodness of fit of the model in detail. Briefly, we observe that the pairwise MEM offers a statistically supported model of observed dynamics, as it accounts for $\approx 70\%$ of the entropy that is not captured by the first-order model. Interestingly, the accuracy of the pairwise MEM plateaus as a function of the amount of data (i.e., the number of subjects) used to fit the model.

Local minima in the brain's functional energy landscape

After fitting the MEM to regional congruence in community allegiance across all subjects, we characterized the resultant energy landscapes of all ROIs. We identified 3–5 local minima states from the landscape of a single ROI, for a total of $M = 25$ unique local minima states across all ROIs. We note that each state represents the set of brain regions that are commonly allied together in a single community (see Fig. 2A). Because ROIs with strong functional interactions are expected to display congruent membership in local minima states, we performed hierarchical clustering on the pattern of ROI allegiance to local minima states. Interestingly, we observed that ROIs divided neatly into three separate classes. Class-I was composed of occipital, parietal, and subcortical ROIs in the visual, attention, sensorimotor, and basal ganglia systems. Class-II was composed of fronto-parietal, and frontal ROIs in the right and left executive control network and the rostral middle frontal systems. Finally, Class-III was composed of medial and opercular ROIs in the dorsal medial prefrontal, default mode, and salience systems.

Although we identified a total of 25 unique local minima states across all ROI landscapes, the majority of these states approximated single communities consisting largely of ROIs from one or two identified classes. For example as see in Fig. 2A, local minima states 1–3 are characterized by congruent community allegiance of Class-II and Class-III ROIs. Hierarchical clustering of local minima states based on their similarity (as measured by the Hamming distance between the local minima states) underlines several groups of very similar local minima states, where each group is characterized by a few common ROIs. Interestingly, these results highlight the tendency of Class-I and Class-III ROIs to display incongruent (or dissimilar) community allegiance, whereas Class-II ROIs form communities with ROIs across both classes.

Basins surrounding local energy minima

In the previous section, we identified minima of the energy landscape underpinning module dynamics, and we further described the clustering of these minima states into groups with similar regional profiles of allegiance to dynamic modules. These minima were located at

the locally least-energetic position on the energy landscape. However, it is also intuitively of interest to study the group of states neighboring each local minimum since the shape of the surrounding low-energy basin effects the state transition dynamics over the basin. The brain states are expected to rapidly converge to the local minima in a smooth, steep, funnel-like basin, whereas over the larger rough basins the systems displays frustrated dynamics; i.e. it is difficult to reach local minima because of the large number of peaks and troughs. To study this larger group of states, we defined *basin states* of local minimum *A* to be any state from which one can travel down the steepest gradient of the energy landscape (using a steep search algorithm) to reach the local minimum *A* (see Methods for details). Using this definition, we estimated any basin's size as the ratio of the number of basin states to the total number of possible states. Since the energy landscapes of the community allegiance states are estimated locally, the estimated landscape for each ROI naturally differs from all the other landscapes due to its reduced dimension ($N-1$). Therefore the size of the analogous basins of attraction can appear marginally different across ROIs. Importantly, this estimate of basin size is not the same as an estimate of basin depth. We observed that the largest basin – making up to >60% of the total state space of Class-I ROIs – surrounded minima state #17, which was characterized by congruent allegiance of visual, sensorimotor, basal ganglia, and attention systems (see Fig. 2B). The large size of this basin suggested that these systems had stronger than average intrinsic functional interactions, perhaps due to the ongoing visual fixation that is characteristic of the resting state with eyes open.

To better understand the anatomical drivers of observed module dynamics, we next studied which brain regions contributed most to basin states surrounding each minimum. To address this goal, we began by counting the number of times that each basin state appeared in a single region's profile. Using this information, we defined for each ROI the number of basin states seen by that ROI as a fraction of the total number of possible states. In other words, associated with each ROI was a vector that represented the fraction of its states that were identified in the basin of each of the minima states. Interestingly, the identified classes of ROIs seemed to show similar profiles, as measured by a Pearson correlation coefficient between any two ROIs' vectors (see Fig. 2C). Moreover, with the exception of a few states (e.g., states 7 and 8), the ROIs from the same class appear with the same frequency between different basin states; in Fig. 2D, this effect is evident through the comparable average allegiance values of ROIs from the same class.

To understand how local minima states with large basin sizes effect the state transition dynamics, we calculated the *dwel time* of each ROI in a given basin. Intuitively, a dwell time is the length of time in which an ROI remains in a given local minima's basin. We observed that the size of the basin was positively related to dwell time within the basin (see Fig. 2E). We note that on average the dwell time of each basin is only a few seconds (< 10 *secs*), except for state #17 which displays twice the average dwell time (21.46 *secs*) although the variance is large (*std*= 31.41 *secs*). In addition, Supplementary figures S11–2 also reveal the relatively rough surface of the state #17 basin, marked by the presence of several prominent peaks and troughs. Together these results suggest that close functional relationships between ROIs (especially between Class-I ROIs) promotes frustrated dynamics over large basins, marked by coherent activity in Class-I ROIs.

The frequency of transitioning between basins

After identifying local minima states and the times spent dwelling within their basin, we next turned to examining the frequencies with which the brain transitioned between single community states to estimate an ROI's transition profile. Because each time point of the subjects' dataset is often associated with two or more single community states, we estimated the frequency of transitioning from one basin to another, separately for each ROI's energy landscape (see Fig. 3). Specifically, after establishing the dynamic functional community organization and the local community allegiance states from the empirical data, we identified the basins that each state belongs to. Next for each ROI, the empirical basin transition frequency from basin *A* to *B* was calculated as the number of state transitions from basin states *A* to *B* across the scan, and the empirical basin transition probability was determined by normalizing the transition frequencies: dividing by the total number of transitions from *A* to all basins. Importantly, we refer to this quantity as the basin transition probability, rather than the probability of transitioning between local minima states, because the measure counts transitions between peripheral areas of basins, rather than only counting transitions from basin centers. For Class-I ROIs that include occipital, parietal, and subcortical areas in the visual, attention, sensorimotor, and basal ganglia systems, we observed that the basin of minima state #17 (characterized by the largest basin size and by congruent allegiance of visual, sensorimotor, basal ganglia, and attention systems) was also the most frequently visited basin. These results once again highlight the strength of this attractor state, which in turn echoes the close functional interactions between Class-I ROIs. For Class-II ROIs including fronto-parietal and frontal areas in the right and left executive control network and the rostral middle frontal systems, we observed a more uniform distribution of transition frequencies between basins. This more uniform transition probability architecture is also characteristic of Class-III ROIs, which are composed of medial and opercular areas in the dorsal medial prefrontal, default mode, and salience systems.

To have a better understanding of the fluidity of the state transitions for each ROI, we measured the level of unpredictability (i.e., entropy) of each ROI's community allegiance states. We calculated the entropy of allegiance state probabilities for each ROI separately. The significantly higher entropy of the Class-II ROI landscapes ($p < 0.01$ bootstrap) provides converging evidence of the highly dynamic community organization of ECN ROIs where their single community states more fluidly transition between local minima state basins as highlighted in Fig. 3 (right).

Next, we asked whether we could use the maximum entropy model results to predict the empirically observed transition frequencies, as a confirmation of the modeling framework. To derive theoretically expected transition frequencies, we simulated the transition dynamics between basins using a Markov Chain Monte Carlo (MCMC) method over the energy landscape of each ROI using the Metropolis-Hastings algorithm (see Methods for details). A secondary benefit of this approach was that it allowed us to identify groups of brain regions whose module dynamics were not well-predicted by a random walk model, and therefore might require consideration of more highly constrained walk dynamics. In general, we observed that simulated and empirical transition frequencies were positively correlated with

one another (average slope $\beta = 1.09$, $p < 0.0002$ except LECN for which $p = 0.02$), offering initial validation of the MEM approach. Interestingly, however, this relationship displayed differential strength across the three classes of ROIs (see as Fig. 4A and SI2 for more details), indicating that some transition frequencies were less well characterized by pairwise interactions (the assumption of the MEM) and/or by random walks across the underlying energy landscape (the assumption of the MCMC). In particular, Class-I ROIs showed the strongest relationships between the simulated and empirical transition frequencies; Class-III ROIs showed significant but weaker correlations and Class-II in general showed insignificant correlations after FDR correction for multiple comparisons. In the earlier MEM analysis, the functional activity of Class-II ROIs were well fit, yet the transition frequencies here find low similarity between the model and empirical data. This suggests that a simple random walk model fails to capture the state transition behavior of putative areas in the executive control networks, which form the bulk of Class-II ROIs.

Finally, we asked whether the empirical transition frequencies could be predicted by simpler statistics drawn from the maximum entropy model (and associated energy landscape) and not requiring the full MCMC modeling approach. In general, we observe a positive relationship between the size of a basin and the transition frequency of that basin: the brain tends to transition into and out of large basins. This effect is strongest in Class-I and Class-III ROIs, and weaker for Class-II ROIs (Fig. 4B). Intuitively, while the basin size is likely to be a strong predictor of transition frequency, another important consideration lies in the energy barriers between basins. That is, are two basins separated by a low hill or by a high mountain on the energy landscape? To clarify the relative predictive power of basin size *versus* barriers between basins, we estimated the energy barrier between pairs of basins by identifying saddle nodes on the energy landscape (see Methods). Importantly, we could estimate these barriers either by considering symmetric transition frequency estimates (averaging both transitions into and out of a state), or by considering the full transition frequency matrix with small asymmetries. We did not observe any consistent trend linking the size of the energy barrier and the empirically observed transition frequencies (Fig. 4C), particularly in the case of the asymmetric estimates (Fig. 4D). These results suggest that module dynamics are best explained by basin size rather than by barriers between basins.

Non-stationarity of BOLD FC timeseries

In this study we assume that statistical interdependence between regions is a reflection of their functional connectivity, and that the fluctuations in FC (or dynamic FC) are a true reflection of changes in the underlying functional relationships between brain regions. However it is possible that these changes are in fact only artifactual (e.g. noise) and the true underlying FC process is a non-stationary one. Several studies have proposed statistical approaches to test for the presence of non-stationarity, including the assessment of the variance of the FC time series (Sakuma et al., 2010), test statistics based on the FC time series' Fourier-transform (Handwerker et al., 2012), linear (e.g., variance of correlation series (Hindriks et al., 2016)) and nonlinear test statistics (Zalesky et al., 2014), among others (Chang and Glover, 2010; Keilholz et al., 2013; Laumann et al., 2016). The bulk of the evidence points to the non-stationary nature of BOLD FC, a conclusion that is consistent with recent work suggesting that non-stationarity in BOLD functional connectivity reflects

changes in ongoing cognitive processes supporting learning, working memory function, linguistic processing, and executive function (Bassett et al., 2011; Braun et al., 2015; Chai et al., 2016, 2017; Hutchison et al., 2013a). Nonetheless, some of these reports are inconclusive, mainly because test statistics are commonly compared against that of null (stationary) time series and creating such time series with matching covariance structure, spectral properties, and stationary FC to this day remains a challenge (Hindriks et al., 2016).

Here we explored an alternative approach to answer this question, namely by comparing the estimated interaction matrices from the original data to that of an appropriate null model. Specifically, we created surrogate time series with preserved covariance structure using the method proposed by (Laumann et al., 2016), which builds the time series from stationary white noise while matching the covariance structure and spectral profile of the empirical data. Next, we estimated the energy landscape of the community allegiance states from the surrogate data. In Fig. 5B, we display estimated local minima states of this surrogate data as well as the corresponding size of their basins, and the basin similarity correlation matrix. Visually, we can note that these results look quite different from the estimated empirical results reported in the manuscript. To statistically test this observation, we performed a non-parametric statistical test (bootstrap $N = 600$) to measure (i) the correlations between the J_{ij} interaction matrices (all possible pairs) estimated from the null time series, and (ii) the correlations between the J_{ij} interaction matrices estimated from a single set of null time series and that of the true data. We observed that the J_{ij} interaction matrices derived from separate sets of null time series were significantly ($p \approx 0$) more similar to one another than J_{ij} interaction matrices derived from the true time series (Fig. 6).

The observed dissimilarity between the J_{ij} matrices estimated from the null and empirical time series suggests that there are fundamental differences in the null and empirical energy landscapes. This suggestion is supported by additional tests examining simple summary statistics including state existence, basin size, and basin dwell time. First, we observe that 6 out of the 25 local minima states identified from the BOLD time series are present less than expected (one-tailed, $p < 0.05$, $N = 1046$ iterations) in the null energy landscapes (Fig. 7A). Importantly, minima states 14 and 15 are significant at an $\alpha = 0.05$, with an FDR correction for multiple comparisons (see Fig. 7A); they also remain significant after a stricter Bonferroni correction. Second, we observe significant (one-tailed, $p < 0.05$, $N = 1046$ iterations) changes in the dwell times as well as in the basin size of several local minima states (Fig. 7B–C). Minima states 9 and 13 in Fig. 7B and minima state 9 in Fig. 7C are also significant at an $\alpha = 0.05$, with an FDR correction for multiple comparisons; they also remain significant after a stricter Bonferroni correction. These data support the notion that the null and empirical energy landscapes are fundamentally different from one another.

It is worth noting that the presence of fluctuations in FC and the presence of functional states are not definitive signs of non-stationarity in the underlying process. In fact, a weak-sense stationary process, such as a stationary stochastic Markov process, can display dynamic FC and state transitions over short timescales, while appearing stationary when studied over long timescales (Liegeois et al., 2017). Therefore, when testing for stationarity of a given timeseries, it is advised to create a null with properties such as linearity and Gaussianity that are analogous to the original timeseries. Here we demonstrated that

empirical interactions estimated using a pairwise MEM were significantly different than those observed in a null with matching covariance structure and spectral profile (Laumann et al., 2016). These differences are consistent with the presence of non-stationarity in the empirical data, but they could also be explained by differences in linearity and Gaussianity of the null model. Further work is necessary to determine whether the empirical time series are in a weak sense stationary or indeed non-stationary.

Discussion

In this work we aimed to characterize the dynamic organization of large-scale brain networks and to provide a probabilistic model of the manner in which the brain transitions between large-scale functional states. Drawing on recent work demonstrating the fundamental nature of modular organization in large-scale functional dynamics (Sporns and Betzel, 2016), we defined these states via local patterns of brain regions' allegiance to network communities (Bassett et al., 2015). We utilized the maximum entropy model (MEM) framework to estimate the probability of occurrence of these states as well as each ROI's co-occurrence with another ROI in the same community. We introduced a novel definition of brain state that allowed us to uncover functional modules that drive large-scale brain functional dynamics. In addition, we studied the inter-ICN relationship while prior work mainly focused on the intra-ICN landscape. Our results highlight the existence of three classes of ROIs with similar functional relationships. Visual, attention, sensorimotor, and subcortical ROIs tend to form a single functional community (Class-I). The remaining ROIs form the putative executive control network (Class-II) and the putative default mode and salience network (Class-III). In addition to identifying these distinct classes of ROIs that display inherently different dynamics within putative functional modules, we also studied the probabilities and frequencies with which the brain transitioned from one pattern of functional modules to another pattern of functional modules. By modeling basin transitions using an MCMC random walk, we predicted empirical probabilities of state transitions with high fidelity for Class-I and Class-III ROIs, and with lower fidelity for executive control (Class-II) ROIs. Interestingly, executive control ROIs also displayed higher entropy energy landscapes, linking diverse state classes, and utilizing uniform transition probabilities across basins, consistent with their unique role in diversifying the brain's dynamic functional repertoire. More generally, the relatively good fit of the MEM suggests that the complex patterns of network module dynamics can be described simply by pairwise interactions between regional allegiances to communities. In contrast to the widespread interest in network statistics, our results provide the critical first steps towards a probabilistic model of brain networks dynamics.

Model-based versus data-driven approaches to studying dynamic functional connectivity

One of our fundamental aims in performing this work was to understand the dynamic functional interactions between large-scale brain networks driven by underlying neurophysiological processes at smaller spatial scales (Logothetis, 2008). While many empirical and data-driven approaches are currently being developed and utilized (Hutchison et al., 2013b), model-based approaches comprise a relatively smaller literature largely including efforts in the Virtual Brain (Roy et al., 2014) and dynamic causal modeling

(Stephan and Roebroek, 2012) communities. Development of such approaches is imperative to our understanding of the role that network dynamics play in attention (Kucyi et al., 2016), learning (Bassett et al., 2015), language (Doron et al., 2012b), and memory (Braun et al., 2015), and their evolution through development (Chai et al., 2017) or alteration in psychiatric disease (Siebenhuhner et al., 2013; Weiss et al., 2011; Du et al., 2016; Yu et al., 2015) and neurological disorders (Khambhati et al., 2015; Burns et al., 2014). Here we build on a maximum entropy based modeling framework that has been previously utilized in the context of functional activation profiles (rather than functional connectivity profiles) (Watanabe et al., 2013, 2014b, 2014c). We adapt this approach to study the functional energy landscape of network *states* and their associated attractor state dynamics. Functional brain states are defined based on regional allegiance to dynamic functional communities, providing insight into the physiological patterns of synchronization between groups of brain regions. This model-based approach revealed notable reductions in the distance ($\approx 70\%$) between the estimated and empirical distributions of patterns of ICN functional modules from resting state data by considering interactions between all pairs of regions. Therefore, the pairwise MEM suggests that the observed patterns of BOLD-derived ICNs dynamic functional communities are partly due to their intrinsic tendency to synchronize with one another, likely across known structural connections (Watanabe et al., 2013).

Attractor communities in dynamic brain networks

Our modeling framework is explicitly based on the role that single regions play in the meso-scale organization of dynamic communities. We fit a maximum entropy model to each ROI's time series of co-occurrence with another ROI in the same community. This approach enables us to identify the dynamic functional roles that different ROIs play in network state dynamics. When combining information from all ROIs, this approach identifies network states that form local energy minima or in essence single community attractors. Indeed, while each region displays a distinct profile of activation across the energy minima, several groups of regions also show similarities in their activation profiles. We refer to these groups as classes, and observe that each single community attractor tended to contain all members of one or two ROI classes: Class-II ROIs (largely comprising executive control regions) coupled with Class-I or Class-III ROIs, whereas Class-I and Class-III ROIs coupled infrequently. This observation suggests that regions in the executive control network play unique roles in meso-scale functional dynamics (Andrews-Hanna et al., 2014), forming transient control hubs that can guide interactions between large-scale functional networks (Mattar et al., 2015; Bassett et al., 2015; Braun et al., 2015; Cole et al., 2013).

Dimensionality of mesoscale brain dynamics

Any assessment of brain states is faced with the question of "So, how many are there?" Most prior studies suggests that brain dynamics can be distilled into between 4 and 7 states at the coarsest level of inquiry (Allen et al., 2012; Britz et al., 2010; Khanna et al., 2015; Shirer et al., 2012b). Yet, evidence from electrophysiology points to the presence of so-called *microstates*, of which there may be many more and which can last for very short periods of time (Khanna et al., 2015; Vakorin et al., 2011). The repertoire of states available to the brain is therefore arguably more accurately characterized as a hierarchy, with a few coarse states composed of multiple levels of more transient (temporally localized) and focal

(spatially localized) states. This complex organization requires computational and data-science approaches (Turk-Browne, 2013) such as the one that we develop in this work. Here, we uncover 25 states, as defined by local minima in the energy landscape of module dynamics during the resting state, and the fact that basin size is positively correlated with transition probability and with transition frequency indicates that the landscape tends to be dominated by large basins. Once the brain transitions out of the large basin, there is a much higher probability that the system will transition back to the large basin than to a small basin. Each state is characterized by a distinct pattern of module allegiance embodied by different brain regions. The identification of these states offers a complementary view to that provided by prior work, which has described changes in module allegiance of brain regions over time (Bassett et al., 2011, 2013b, 2015; Braun et al., 2015; Mattar et al., 2015) and speculated on the cognitive drivers of these changes as regions critical for domain-general processing (Fedorenko and Thompson-Schill, 2014b), associative processing (Bassett et al., 2013b), cognitive control (Bassett et al., 2015), and cognitive flexibility (Braun et al., 2015). While regional roles are important, the states themselves may also offer insights into what cognitive processes are occurring, either in parallel or in series (Mattar et al., 2015). It would be interesting in future work to manipulate cognitive processes via task performance and determine the direct relationship between local minima states and mental states (Andrews-Hanna et al., 2014; Kucyi et al., 2016; Betzel et al., 2016b; Gu et al., 2017).

Cognitive control, flexibility, and task-switching

Executive control networks in fronto-parietal cortices play a unique role in the pattern of results that we uncover here. Executive ROIs display states of module allegiance that are unlike the states displayed by other brain areas. Moreover, our results suggest that the brain transitions in and out of these states in a manner that is not as well-fit by an MCMC random walk on the observed energy landscape, suggesting a peculiar complexity of dynamics. Interestingly, executive ROIs also display higher entropy energy landscapes, link diverse state classes, and utilize uniform transition probabilities across basins. These findings are particularly interesting when viewed in the context of executive control function, and its instantiation in brain network architecture. Executive function supports the ability to link information to solve problems, inhibit inappropriate behaviors, and transition between tasks and states (Royall et al., 2002). Recent work suggests that these capabilities occur via dynamic interactions between large-scale neural circuits (Dajani and Uddin, 2015; Cole et al., 2013), often taking the form of competitive or cooperative dynamics (Cocchi et al., 2013) between putative functional modules (Mattar et al., 2015). Indeed, recent evidence points to fronto-parietal cortices as hubs of flexible modular reconfiguration during task states (Bassett et al., 2013b), that are directly correlated with individual differences in learning (Bassett et al., 2011, 2015), memory (Braun et al., 2015), and cognitive flexibility (Braun et al., 2015). Our current results complement these findings by suggesting that the baseline functional architecture of executive regions supports their role during task performance: (i) the high entropy energy landscapes of these regions can support highly transiently dynamics, and (ii) the uniform transition probabilities can support the integration (cooperative) and segregation (competitive) of many other cognitive systems. It is interesting to speculate that these unique features of executive region dynamics observed at rest may in

part be driven by white matter microstructure, consistent with recent evidence pointing towards a structural driver of executive function (Gu et al., 2015).

Methodological considerations

For computational reasons, we have reduced the dimensionality of the data in two ways. First, we chose a small number of ROIs ($n = 10$) whose time series offer a reasonable representation of the resting state networks from which they are derived (see Methods). However, because of this down-sampling procedure, we are not sensitive to dynamic fluctuations *within* a network. Future work using larger datasets could explore whether higher order ICA decompositions offer additional insights into finer-scale dynamics underlying the functional hierarchy that we observe here. Second, we fit the MEM to each ROI's community allegiance time series, rather than to all ROI pairs' allegiance at the same time. Future work could aim to develop novel optimization-based methods to characterize the energy landscape of the global structure.

Here in the main manuscript, we present the results based on ROI time series to maximize interpretability and anatomical specificity of claims. Yet, a single estimated time series cannot exactly capture the activity of an entire network. To address this limitation, it is useful to consider the utility of examining the back-projection of the component itself. In the supplementary section entitled 'Energy Landscape of ICA Components', we examine the ICA back-reconstructed time series and observe results that are consistent with those we obtained from the ROI time series.

Moreover, recent work has demonstrated that dynamic FC displays significant covariance with head motion after standard motion correction (Burgess et al., 2016). Therefore it is critical to understand the robustness of the reported results to any residual noise. The converging results from the ROI-based and ICA back-reconstructed time series (presented in SI-5) suggest that any residual effect of head motion does not simply explain the results and interpretations provided in this paper.

A distinct set of important methodological considerations relates to the dynamic community structure that we estimate as the input to the regional MEMs. We estimate dynamic community structure using a modularity maximization approach (Lancichinetti and Fortunato, 2011), which has an implicit structural resolution parameter that can be used to tune the number of communities. Following prior work, we employ the default parameter value of unity (Bassett et al., 2013a), and it would be interesting in future to study the changes in energy landscapes that occur at different spatial and temporal scales (Mucha et al., 2010). Moreover, following the extraction of dynamic community structure, we binarize the pair-wise allegiance probabilities, necessarily losing sensitivity to fine-scale network perturbations of nodes that are loosely associated with a single community. However, since we mainly focus on the most robust attractor communities, the computational benefits of the reduced state space following binarization outweighs the cost associated with this lack of sensitivity.

It is important to note that while we apply the maximum entropy model to study the relationships between functional connections estimated with a Pearson correlation, the

results do not address the question of whether higher-order (>2nd) interactions exist in functional brain networks. This question requires a different statistical approach, such as explicitly fitting higher-order MEMs to activity time series rather than functional connections. Prior work on this question has suggested that while there may be a few higher-order relationships (Ganmor et al., 2011), most of the variance in neurophysiological activity can be accounted for by pairwise MEMs both at the level of BOLD (Watanabe et al., 2013) and at the level of neuronal activity (Schneidman et al., 2006).

We performed an additional test to assess the likely affect of head motion on the modeled dynamics. We examined the Spearman correlation coefficient between the mean framewise displacement and the total basin transition frequency between all pairs of basins, assessed over subjects. We observed no statistically significant relationship between total basin transition frequency and mean framewise displacement, either when considering all sessions for all subjects, or when considering across-session averages (all $p > 0.05$, FDR corrected for multiple comparisons across regions). These data suggest that the frame-wise displacement is a poor predictor of basin transition dynamics.

Conclusions and future directions

Here we present a viable probabilistic model of dynamic reconfiguration in functional brain networks estimated from resting state fMRI. By representing a brain state as a pattern of functional interactions between brain regions, we reveal structured transitions between a finite number of brain states that act as basins of attraction. Critically, each basin is characterized by a specific set of functional modules: groups of brain areas that display coherent BOLD time series. By characterizing the energy landscape surrounding these basins, we accurately predict the manner in which the brain transitions between states, and we uncover novel markers of the functional role that executive regions play in guiding these transitions. These efforts lay the groundwork for empirical investigations into how these energy landscapes change during task performance, over normative neurodevelopment, throughout healthy aging, or in the context of psychiatric disease or neurological disorders. Moreover, they lay important theoretical groundwork in the critical development of probabilistic models of brain network dynamics subserving cognitive function.

Methods

Participants

Twenty participants (nine female; ages 19–53 years; mean age = 26.7 years) with normal or corrected vision and no history of neurological disease or psychiatric disorders were recruited for this experiment. All participants volunteered and provided informed consent in writing in accordance with the guidelines of the Institutional Review Board of the University of Pennsylvania (IRB #801929).

Human fMRI Data collection

Magnetic resonance images were obtained at the Hospital of the University of Pennsylvania using a 3.0 T Siemens Trio MRI scanner equipped with a 32-channel head coil. T1-weighted structural images of the whole brain were acquired on the first scan session using a three-

dimensional magnetization-prepared rapid acquisition gradient echo pulse sequence (repetition time (TR) 1620 ms; echo time (TE) 3.09 ms; inversion time 950 ms; voxel size 1 mm × 1 mm × 1 mm; matrix size 190 × 263 × 165). A field map was also acquired at each scan session (TR 1200 ms; TE1 4.06 ms; TE2 6.52 ms; flip angle 60°; voxel size 3.4 mm × 3.4 mm × 4.0 mm; field of view 220 mm; matrix size 64 × 64 × 52) to correct geometric distortion caused by magnetic field inhomogeneity. In all experimental runs with a behavioral task, T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a slice accelerated multiband echo planar pulse sequence (TR 2000 ms; TE 25 ms; flip angle 60°; voxel size 1.5 mm × 1.5 mm × 1.5 mm; field of view 192 mm; matrix size 128 × 128 × 80). In all resting state runs, T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a slice accelerated multiband echo planar pulse sequence (TR 500 ms; TE 30 ms; flip angle 30°; voxel size 3.0 mm × 3.0 mm × 3.0 mm; field of view 192 mm; matrix size 64 × 64 × 48).

fMRI Preprocessing

We preprocessed the resting state fMRI data using FEAT (FMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, <http://www.fmrib.ox.ac.uk/fsl>). Specifically, we applied: EPI distortion correction using FUGUE (Jenkinson, 2004); motion correction using MCFLIRT (Jenkinson et al., 2002); slice-timing correction using Fourier-space timeseries phase-shifting; non-brain removal using BET (Smith, 2002); grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor; highpass temporal filtering (Gaussian-weighted least-squares straight line fitting, with sigma = 50.0 s). Our sample had an average relative framewise displacement of 0.047 mm (standard deviation of 0.015 mm).

Nuisance timeseries were voxelwise regressed from the preprocessed data. Nuisance regressors included (i) three translation (X, Y, Z) and three rotation (Pitch, Yaw, Roll) timeseries derived by retrospective head motion correction ($R = [X, Y, Z, pitch, yaw, roll]$), together with expansion terms ($[RR^2R_{t-1}R_{t-1}^2]$), for a total of 24 motion regressors (Friston et al., 1996); (ii) the five first principal components calculated from timeseries derived from regions of non-interest (white matter and cerebrospinal fluid), using the anatomical CompCor method (aCompCor) (Behzadi et al., 2007) and (iii) the average signal derived from white matter voxels located within a 15 mm radius from each voxel, following the ANATICOR method (Jo et al., 2010). Global signal was not regressed out of voxel time series (Murphy et al., 2009; Saad et al., 2012; Chai et al., 2012). Finally, the mean functional image and the 125-scale Lausanne parcellation template (Cammoun et al., 2012a) were coregistered using Statistical Parametric Mapping software (SPM12; Wellcome Department of Imaging Neuroscience, <http://www.fil.ion.ucl.ac.uk/spm>) in order to extract ROIs' mean timeseries.

ICA-informed identification of regions of interest

We used group-ICA (GIFT toolbox (Calhoun et al., 2001)) to identify ten large-scale intrinsic connectivity networks (ICN) characteristic of the resting state. Computational considerations preclude us from studying a larger number of networks in the context of the maximum entropy model approach. Next, we identified the regions of interest (ROIs) by

choosing the ROI from a commonly used anatomical atlas (the 125 scale Lausanne parcellation (Cammoun et al., 2012b)) in which we observed the maximum weighting of an ICA component. A list of identified ICNs and their corresponding Lausanne atlas ROI are provided in Table 1. In Supplementary Fig. SI5, we also discuss the effect of anatomical *versus* functional parcellation methods on our results.

For computational reasons which we will discuss in more details shortly, we only included ROIs from the left hemisphere for the majority of components. The one exception to this rule was the right executive control network (RECN). Prior work has demonstrated that the putative fronto-parietal executive control network is identified as two bilateral ICNs (Shirer et al., 2012a; Laird et al., 2011), which separately play critical roles in executive function through their dynamical interactions with the DMN and attention systems (Andrews-Hanna et al., 2014; Menon and Uddin, 2010). Thus we included time series from both the right and left executive control networks in all following analysis.

A note on statistics and computational considerations

The fit of the maximum entropy model to small datasets can be subject to a sampling bias. More exactly, the estimation of the entropy suffers severely from downward bias (Treves and Panzeri, 1995) such that the estimated entropy from the observations is lower than the actual entropy of the underlying model (Macke et al., 2011). Ultimately the amount of data needed to accurately fit the model is exponential to the number of ROIs. Therefore even a small number of ROIs (on the order of ten) requires access to extremely large datasets. This issue of computational complexity is especially critical for fMRI data where the slow sampling rate prohibits collection of individual subject data with a large number of observations. Consequently, in this manuscript we focus on group concatenated multiband data (with a quarter-second TR) and do not discuss results from individual subjects. In Supplementary Figure SI1, we briefly discuss the model fit at the subject- and group-level where we demonstrate that the goodness of fit and accuracy of the model drops considerably when considering individual subjects as opposed to the group.

The total number of partitions of an n -element set is the Bell number nB_n (Grimaldi, 2006). In combinatorial mathematics, the Bell numbers count the number of partitions of a set. Thus the total number of possible community states of the 10 ROIs equals 115975. To accurately model this large number of states would require a large amount of data. However, local analysis (through the lens of a single ROI) of congruent community allegiance deals with a much smaller state space of 2^{N-1} (=512). At this level, it is computationally feasible to fit a MEM with a relatively large multiband fMRI dataset such as the one we use here.

It is worth noting that ICA assumes the BOLD data can be linearly decomposed into several spatially independent sources and the activation of each voxel is a mixture of all the ICA sources with varying contribution to each component. Therefore in order to find an ROI within the component that is least diluted with other sources and represents the estimated ICA time series, we need to pick the ROI with the highest contribution to each component. In fact our complementary analysis (SI5) of the energy landscape of the ICA time series shows the close similarity between the results, which provides additional verification that the chosen ROI correctly represents the estimated signal sources. Nevertheless as shown

recently by Watanabe et al. (2013), the resting state networks such as the DMN and executive control network also contain several basins of attractions of their own. Therefore we expect to see the reported landscape change if we choose ROIs other than the peak voxel, and these landscapes will be less representative of the average large-scale dynamics. To support this assumption, we have estimated the energy landscape of a randomly chosen set of ROIs from each ICA component. The results are presented in Fig. 11. B, where we observe that the energy landscape differs from that characteristic of the ROIs most highly contributing to the independent component. Therefore, choosing ROIs other than the peak ROI will alter the results, either due to mixing of sources or due to the true internal state space dynamics. Our choice to use the ROI with the highest contribution to the component ensures that our results are most representative of the average state dynamics, but finer-scale state dynamics could be explored by choosing ROIs that are not the highest contributors to the ICs.

Functional network construction

Following prior work (Bassett et al., 2011), we estimated the dynamic functional connectivity between all pairs of ROIs using wavelet coherence (WTC (Grinsted et al., 2004)). In Supplementary SI4 Fig. 1, we show that we observe two distinct bands of high WTC: 0.64–0.2 Hz and 0.19–0.06 Hz. In the main manuscript, we focus on the 0.19–0.06 Hz band due to known sensitivity to underlying neural activity, and we relegate discussion of the higher frequency band to the Supplement. WTC amplitudes were averaged over all frequencies within the selected band to construct the timecourse of the bandpassed WTC for each pair of ROIs resulting in a total of $T (= 1190 \text{ (TRs)} \times 20 \text{ (subjects)} \times 4 \text{ (runs)} = 95200)$ unique functional connectivity patterns, which we represent in $N \times N$ adjacency matrices \mathbf{A} (see Fig. 8).

Community detection and module allegiance estimation

In the maximum entropy framework, it is critical that data points are as temporally distinct from one another as possible. In the context of our study, this requires that we reduce the dependence of community structure in the neighboring time slices. To do so, we identified the community structure (Fortunato, 2010; Porter et al., 2009) of the each time slice adjacency matrix independently using a Louvain-like (Blondel et al., 2008) locally greedy heuristic algorithm to maximize the modularity quality function (Newman, 2006) with a structural resolution parameter of $\gamma = 1$ (Bassett et al., 2013a). The method partitions ROIs into communities based on the optimization of the following function:

$$Q_0 = \sum_{ij} [A_{ij} - \gamma P_{ij}] \delta(g_i, g_j)$$

where \mathbf{A} is a weighted adjacency matrix, ROI i and j are assigned respectively to community g_i and g_j , the Kronecker delta $\delta(g_i, g_j) = 1$ if $g_i = g_j$ (and zero otherwise), γ is the structural resolution parameter, and P_{ij} is the expected weight of the edge between ROI i and j under some null model. We used the Newman-Girvan null model (Girvan and Newman, 2002)

$$P_{ij} = \frac{k_i k_j}{2m}$$

where $k_i = \sum_j A_{ij}$ is the strength of ROI i and $m = \frac{1}{2} \sum_{ij} A_{ij}$. Briefly, this method partitions the ROIs into groups such that the total connection strength within communities is more than expected in the null model.

Importantly, the algorithm we use is a heuristic that implements a non-deterministic optimization (Good et al., 2010). Consequently we repeated the optimization 100 times (Bassett et al., 2013a), and we report results summarized over those iterations by forming a module allegiance matrix (Bassett et al., 2015; Mattar et al., 2015).

The allegiance matrix for each time slice represented the probability that ROI i and j were assigned to the same community over all iterations of the community detection algorithm (Fig. 9). For use in the maximum entropy model, we binarized the allegiance matrix by subtracting a random null model allegiance matrix from the original allegiance matrix: elements in the binarized allegiance matrix were 1 when the true allegiance was greater than the null, and 0 otherwise (Bassett et al., 2013a). The null allegiance matrices were generated by shuffling the ROI community assignments for each individual time point uniformly at random (Fig. 10).

Maximum entropy model fitting

Here we hypothesize that the brain transitions between different functional community states. To obtain an unbiased estimate of these states and their probabilities, we fit a pair-wise maximum entropy model. The principle of maximum entropy states that when estimating the probability distribution, given the constraints, one should find the distribution that maximizes the uncertainty (i.e., entropy). Choosing any other distribution that lowers the entropy would assume information that we do not possess; therefore the only reasonable distribution is the maximum entropy distribution. Fitting the MEM entails tuning the first and second-order interaction parameter between regions so that the predicted activation rates and co-activation rates match that of the empirically observed values. An accurate pair-wise MEM fit suggests that the observed dynamics of the communities can be simply explained as a combination of each region's independent activation rate plus the regions' joint activation rates. In other words, the MEM allows us to establish a model of brain functional dynamics as a probabilistic process shaped by intrinsic relationships between brain regions.

We fit the pairwise MEM to the binarized community allegiance of ROI pairs. To reduce the size of the state space and therefore ensure less error in our estimates, we fit the MEM to each row of the allegiance matrix independently for each ROI, effectively reducing the dimensionality to $N-1$ ($= 9$) and reducing the total number of possible congruent community membership states to $2^9 = 512$. For ROI i at time t the congruent community membership state is defined as $V^t = [\sigma_1^t, \sigma_2^t, \dots, \sigma_{i-1}^t, \sigma_{i+1}^t, \dots, \sigma_N^t]$, where σ_j^t is the binarized community allegiance of i and j at time point t ('1' for congruent community membership and '0' otherwise), and N is the total number of ROIs ($= 10$). Because it is not possible to

directly compare the local community allegiance states due to the unmatched cardinality of the set of ROIs, we assume that the ROI being examined will always have an allegiance state of one. For ROI i the empirical congruent community membership rate of ROI j , $\langle \sigma_j \rangle$, is

given by $(\frac{1}{T}) \sum_{t=1}^T \sigma_j^t$, where T is the number of time slices, which in this case is equal to the number of TRs. Likewise the empirical pairwise congruent community membership rate of

ROIs j and l , $\langle \sigma_j \sigma_l \rangle$, is defined as $(\frac{1}{T}) \sum_{t=1}^T \sigma_j^t \sigma_l^t$.

Here our only constraints were that the model $\langle \sigma_j \rangle_m$ and $\langle \sigma_j \sigma_l \rangle_m$ matched the empirical values of $\langle \sigma_j \rangle$ and $\langle \sigma_j \sigma_l \rangle$ respectively. It is known that given these constraints the probability distribution that maximizes the entropy is the Boltzman distribution (Jaynes, 1957)

$$P(V_k) = e^{-E(V_k)} / \sum_{q=1}^{2^{N-1}} e^{-E(V_q)} \quad (1)$$

where $P(V_k)$ is the probability distribution of k th state V_k , and $E(V_k)$ the energy of that state given by

$$E(V_k) = - \sum_{j=1}^{N-1} h_j \sigma_j(V_k) - \frac{1}{2} \sum_{j=1}^{N-1} \sum_{l=1, j \neq l}^{N-1} j_{jl} \sigma_j(V_k) \sigma_l(V_k) \quad (2)$$

where $\sigma_j(V_k)$ is the value of σ_j for state V_k , h_j represents the expected base allegiance of ROI j (with respect to ROI i) in isolation, and j_{jl} represents the functional interaction between ROI j and l . Fitting the MEM entails iterative adjustment of h_j and j_{jl} with a gradient ascent algorithm (similar to (Watanabe et al., 2014c)) until the empirical $\langle \sigma_j \rangle$ and

$\langle \sigma_j \sigma_l \rangle$ values approximately match the model $\langle \sigma_j \rangle_m = \sum_{q=1}^{2^{n-1}} \sigma_j(V_q) P(V_q)$ and

$\langle \sigma_j \sigma_l \rangle_m = \sum_{q=1}^{2^{n-1}} \sigma_j(V_q) \sigma_l(V_q) P(V_q)$. In depth analysis of the goodness of MEM fits (provided in S11), allows us to conclude that the pair-wise MEM can account for a large portion of the observed functional module dynamics. Nevertheless, higher-order and/or nonlinear interactions likely contribute to smaller yet non-negligible portions of the observed brain dynamics.

It is important to note that in our presented framework, the concept of energy is not a literal one, but is rather metaphorical. Indeed, we cannot yet claim that the calculations are related to biological notions of energy (e.g., ATP) or physical notions of energy (e.g., an exact Hamiltonian for BOLD fluctuations). Instead, we use the term because it is the technical term associated with the landscapes that can be extracted from a maximum entropy model, and is therefore a relevant function based on the region's community allegiance that allows us to capture the observed spatial correlation patterns. It should be treated simply as a useful mathematical quantity.

Defining an energy landscape

The energy landscape for each ROI is defined separately by the network of congruent community membership states V_k and their corresponding energy values $E(V_k)$. In this landscape the adjacent states' vectors are only one hamming distance apart, which means that all ROIs except one have the same binary values as in adjacent states' vectors. (Watanabe et al., 2014a).

An interesting question to ask of this landscape is the location and nature of energy minima. To address this question, we exhaustively searched the entire landscape using a steep search algorithm to find the states with energies lower than all their neighboring states (i.e. local minima). Next in order to identify the states that belong to the basin of each local minima we first start at a given state V_k (one of the $2^{N-1}(=2^9 = 512)$ possible states) and iteratively move to the neighboring state V_w in the landscape if $E(V_w) < E(V_k)$. We continue tracing out this path until we reach a local minima state where no neighboring states exist with smaller energy values (similar to (Watanabe et al., 2014c, 2013)). We consider this final state V_k the basin state of the local minima. We also define the basin size of that local minima state as the fraction of the number of basin states to the total number of possible states.

We were next interested in understanding the predicted barriers between states. We estimated the energy barrier opposing the transition between all the local minima states in the following way. (1) We removed the state (node) with the highest energy from the energy landscape along with the edges connecting that state to its neighbors. (2) We assessed whether each pair of local minima were connected by a path in the reduced landscape. We repeated steps (1) and (2) until we found the saddle state where removing the highest energy node disconnects one or more local minima from the rest of the landscape. We continued this process until we obtained a reduced landscape where all of the local minima are isolated and we identified all saddle states. (3) We calculated the symmetric energy barrier (Zhou, 2011) between all pairs of minima states as the minimum of $[E^S(V_k, V_w) - E(V_k), E^S(V_k, V_w) - E(V_w)]$, where $E^S(V_k, V_w)$ is the energy of the saddle point between minima states V_k, V_w and $E(V_k)$ and $E(V_w)$ are the energies at these states, respectively. If the energy barrier between to minima states was high then the model predicts that the rate of transition between them is low, at least in one direction (Watanabe et al., 2013). We also calculated the asymmetric energy barrier between minima states V_k, V_w as $E^S(V_k, V_w) - E(V_k)$ and $E^S(V_k, V_w) - E(V_w)$, where the former indicated the $V_k \rightarrow V_w$ and the latter the $V_w \rightarrow V_k$ energy barriers. Overall, our results did not show any relationship between estimated energy barriers (symmetric and asymmetric) between local minima states and the empirical basin transition probabilities and frequencies. We speculate that the energy barrier mimics the basin transition probabilities only when the basins are smooth and funnel-like. We speculate that other unaccounted factors such as the shape of the basins may contribute more to the observed basin transition probabilities than the energy barriers between the basins.

Simulation of state transitions

To better understand the dynamic patterns of functional communities at rest, we simulated these dynamics as a random walk process over the estimated local energy landscapes using a Markov chain Monte Carlo with Metropolis-Hastings (MCMC) algorithm (Metropolis et al.,

1953; Hastings, 1970; Zhou, 2011). In this model, local community allegiance state V_i is allowed an isometric transition to one of $N-1$ neighboring state with uniform probability. Next the actual transition from V_i to V_j occurs with probability $P_{ij} = \min [1, e^{E(V_i)-E(V_j)}]$. For each ROI, we repeated a 4×10^7 step (plus 3^4 initial steps) walk with randomly chosen initial states 2^4 times. Next, we removed the initial steps to ensure independence of results from the initial conditions and decreased the sampling rate by 500 to reduce the correlation between the samples. Since each state in the energy landscape belongs to the basin of a single local minima, we can construct a trajectory of local minima states' basin transitions from the down-sampled state transitions patterns. Comparing the empirical and simulated basin transition probability and frequency patterns allows us to evaluate the resemblance of the proposed random walk model's dynamics to that of the brain.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was directly supported by the Army Research Laboratory through contract number W911NF-10-2-0022, a collaborative mechanism between J.M.V and D.S.B. In addition, D.S.B., S.G., M.G.M, and A.A. would also like to acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the Army Research Office through contract number W911NF-14-1-0679, the National Institute of Health (2-R01-DC-009209-11, 1R01HD086888-01, R01-MH107235, R01-MH107703, and R21-MH-106799), the Office of Naval Research, and the National Science Foundation (BCS-1441502, CAREER PHY-1554488, and BCS-1631550). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies

References

- Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD. Tracking whole-brain connectivity dynamics in the resting state. *Cereb Cortex*. 2012; 24:663–676. [PubMed: 23146964]
- Andrews-Hanna JR, Smallwood J, Spreng RN. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Ann New Y Acad Sci*. 2014; 1316:29–52.
- Bassett DS, Porter MA, Wymbs NF, Grafton ST, Carlson JM, Mucha PJ. Robust detection of dynamic community structure in networks. *ChaosL: an Interdisciplinary. J Nonlinear Sci*. 2013a; 23:013142.
- Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST. Dynamic reconfiguration of human brain networks during learning. *Proc Natl Acad Sci USA*. 2011; 108:7641–7646. [PubMed: 21502525]
- Bassett DS, Wymbs NF, Rombach MP, Porter MA, Mucha PJ, Grafton ST. Task-based core-periphery organization of human brain dynamics. *PLoS Comput Biol*. 2013b; 9:e1003171. [PubMed: 24086116]
- Bassett DS, Yang M, Wymbs NF, Grafton ST. Learning-induced autonomy of sensorimotor systems. *Nat Neurosci*. 2015; 18:744–751. [PubMed: 25849989]
- Behzadi Y, Restom K, Liao J, Liu TT. A component based noise correction method (compcor) for bold and perfusion based fmri. *NeuroImage*. 2007; 37:90–101. [PubMed: 17560126]
- Betzel RF, Miši B, He Y, Rumschlag J, Zuo XN, Sporns O. Functional brain modules reconfigure at multiple scales across the human lifespan. 2015 arXiv:1510.08045.
- Betzel RF, Satterthwaite TD, Gold JI, Bassett DS. A positive mood, a flexible brain. 2016a arXiv: 1601.07881.
- Betzel RF, Satterthwaite TD, Gold JI, Bassett DS. A positive mood, a flexible brain. 2016b arXiv: 1601.07881.

- Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *J Stat Mech*. 2008:P10008.
- Braun U, Schäfer A, Bassett DS, Rausch F, Schweiger JI, Bilek E, Erk S, Romanczuk-Seiferth N, Grimm O, Geiger LS, et al. Dynamic brain network reconfiguration as a potential schizophrenia genetic risk mechanism modulated by nmda receptor function. *Proc Natl Acad Sci*. 2016; 113:12568–12573. [PubMed: 27791105]
- Braun U, Schafer A, Walter H, Erk S, Romanczuk-Seiferth N, Haddad L, Schweiger JI, Grimm O, Heinz A, Tost H, Meyer-Lindenberg A, Bassett DS. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proc Natl Acad Sci USA*. 2015; 112:11678–11683. [PubMed: 26324898]
- Britz J, Van De Ville D, Michel CM. BOLD correlates of EEG topography reveal rapid resting-state network dynamics. *NeuroImage*. 2010; 52:1162–1170. [PubMed: 20188188]
- Burgess GC, Kandala S, Nolan D, Laumann TO, Power JD, Adeyemo B, Harms MP, Petersen SE, Barch DM. Evaluation of denoising strategies to address motion-correlated artifacts in resting-state functional magnetic resonance imaging data from the human connectome project. *Brain Connect*. 2016; 6:669–680. [PubMed: 27571276]
- Burns SP, Santaniello S, Yaffe RB, Jouny CC, Crone NE, Bergey GK, Anderson WS, Sarma SV. Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proc Natl Acad Sci USA*. 2014; 111:E5321–E5330. [PubMed: 25404339]
- Calhoun V, Adali T, Pearlson G, Pekar J. A method for making group inferences from functional mri data using independent component analysis. *Human Brain Mapp*. 2001; 14:140–151.
- Cammoun L, Gigandet X, Meskaldji D, Thiran JP, Sporns O, Do KQ, Maeder P, Meuli R, Hagmann P. Mapping the human connectome at multiple scales with diffusion spectrum mri. *J Neurosci Methods*. 2012a; 203:386–397. [PubMed: 22001222]
- Cammoun L, Gigandet X, Meskaldji D, Thiran JP, Sporns O, Do KQ, Maeder P, Meuli R, Hagmann P. Mapping the human connectome at multiple scales with diffusion spectrum MRI. *J Neurosci Methods*. 2012b; 203:386–397. [PubMed: 22001222]
- Chai LR, Khambhati AN, Ciric R, Moore T, Gur RC, Gur RE, Satterthwaite TD, Bassett DS. Evolution of brain network dynamics in neurodevelopment. *Network Neuroscience*. 2017
- Chai XJ, Berken JA, Barbeau EB, Soles J, Callahan M, Chen JK, Klein D. Intrinsic functional connectivity in the adult brain and success in second-language learning. *J Neurosci*. 2016; 36:755–761. [PubMed: 26791206]
- Chai XJ, Castañón AN, Öngür D, Whitfield-Gabrieli S. Anticorrelations in resting state networks without global signal regression. *NeuroImage*. 2012; 59:1420–1428. [PubMed: 21889994]
- Chang C, Glover GH. Time-frequency dynamics of resting-state brain connectivity measured with fmri. *NeuroImage*. 2010; 50:81–98. [PubMed: 20006716]
- Cocchi L, Zalesky A, Fornito A, Mattingley JB. Dynamic cooperation and competition between brain systems during cognitive control. *Trends Cogn Sci*. 2013; 17:493–501. [PubMed: 24021711]
- Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE. Intrinsic and task-evoked network architectures of the human brain. *Neuron*. 2014; 83:238–251. [PubMed: 24991964]
- Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, Braver TS. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat Neurosci*. 2013; 16:1348–1355. [PubMed: 23892552]
- Dajani DR, Uddin LQ. Demystifying cognitive flexibility: implications for clinical and developmental neuroscience. *Trends Neurosci*. 2015; 38:571–578. [PubMed: 26343956]
- Doron KW, Bassett DS, Gazzaniga MS. Dynamic network structure of interhemispheric coordination. *Proc Natl Acad Sci*. 2012a; 109:18661–18668. [PubMed: 23112199]
- Doron KW, Bassett DS, Gazzaniga MS. Dynamic network structure of interhemispheric coordination. *Proc Natl Acad Sci USA*. 2012b; 109:18661–18668. [PubMed: 23112199]
- Du Y, Pearlson GD, Yu Q, He H, Lin D, Sui J, Wu L, Calhoun. Interaction among subsystems within default mode network diminished in schizophrenia patients:a dynamic connectivity approach. *Schizophr Res*. 2016; 170:55–65. [PubMed: 26654933]
- Fedorenko E, Thompson-Schill SL. Reworking the language network. *Trends Cogn Sci*. 2014a; 18:120–126. [PubMed: 24440115]

- Fedorenko E, Thompson-Schill SL. Reworking the language network. *Trends Cogn Sci*. 2014b; 18:120–126. [PubMed: 24440115]
- Fortunato S. Community detection in graphs. *Phys Rep*. 2010; 486:75–174.
- Friston KJ, Williams S, Howard R, Frackowiak RS, Turner R. Movement-related effects in fmri time-series. *Magn Reson Med*. 1996; 35:346–355. [PubMed: 8699946]
- Ganmor E, Segev R, Schneidman E. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc Natl Acad Sci*. 2011; 108:9679–9684. [PubMed: 21602497]
- Gerraty RT, Davidow JY, Foerde K, Galvan A, Bassett DS, Shohamy D. Dynamic flexibility in striatal-cortical circuits supports reinforcement learning. 2016:094383. bioRxiv.
- Girvan M, Newman MEJ. Community structure in social and biological networks. *Proc Natl Acad Sci USA*. 2002; 99:7821–7826. [PubMed: 12060727]
- Good BH, de Montjoye YA, Clauset A. Performance of modularity maximization in practical contexts. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2010; 81:046106. [PubMed: 20481785]
- Grimaldi, RP. *Discrete and Combinatorial Mathematics, 5/e*. Pearson Education India; 2006.
- Grinsted A, Moore JC, Jevrejeva S. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlin Process Geophys*. 2004; 11:561–566.
- Gu S, Betzel RF, Cieslak M, Delio P, Grafton ST, Pasqualetti F, Bassett BS. Optimal trajectories of brain state transitions. *Neuroimage*. 2017; 148:305–317. [PubMed: 28088484]
- Gu S, Pasqualetti F, Cieslak M, Telesford QK, Yu AB, Kahn AE, Medaglia JD, Vettel JM, Miller MB, Grafton ST, Bassett DS. Controllability of structural brain networks. *Nat Commun*. 2015; 6:8414. [PubMed: 26423222]
- Handwerker DA, Roopchansingh V, Gonzalez-Castillo J, Bandettini PA. Periodic changes in fmri connectivity. *NeuroImage*. 2012; 63:1712–1719. [PubMed: 22796990]
- Hastings WK. Monte carlo sampling methods using markov chains and their applications. *Biometrika*. 1970; 57:97–109.
- Hindriks R, Adhikari MH, Murayama Y, Ganzetti M, Mantini D, Logothetis NK, Deco G. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fmri? *NeuroImage*. 2016; 127:242–256. [PubMed: 26631813]
- Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J, Handwerker DA, Keilholz S, Kiviniemi V, Leopold DA, de Pasquale F, Sporns O, Walter M, Chang C. Dynamic functional connectivity: promise, issues, and interpretations. *NeuroImage*. 2013a; 80:360–378. [PubMed: 23707587]
- Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J, et al. Dynamic functional connectivity: promise, issues, and interpretations. *NeuroImage*. 2013b; 80:360–378. [PubMed: 23707587]
- Jaynes ET. Information theory and statistical mechanics. *Phys Rev*. 1957; 106:620.
- Jenkinson, M. Improving the registration of b0-distorted epi images using calculated cost function weights. *Proceedings of the Tenth International Conference on functional mapping of the human brain*; 2004.
- Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*. 2002; 17:825–841. [PubMed: 12377157]
- Jo HJ, Saad ZS, Simmons WK, Milbury LA, Cox RW. Mapping sources of correlation in resting state fmri, with artifact detection and removal. *Neuroimage*. 2010; 52:571–582. [PubMed: 20420926]
- Keilholz SD, Magnuson ME, Pan WJ, Willis M, Thompson GJ. Dynamic properties of functional connectivity in the rodent. *Brain Connect*. 2013; 3:31–40. [PubMed: 23106103]
- Khambhati AN, Bassett DS, Oommen BS, Chen SH, Lucas TH, Davis KA, Litt B. Recurring functional interactions predict network architecture of interictal and ictal states in neocortical epilepsy. 2016:090662. bioRxiv.
- Khambhati AN, Davis KA, Oommen BS, Chen SH, Lucas TH, Litt B, Bassett DS. Dynamic network drivers of seizure generation, propagation and termination in human neocortical epilepsy. *PLoS Comput Biol*. 2015; 11:e1004608. [PubMed: 26680762]

- Khanna A, Pascual-Leone A, Michel CM, Farzan F. Microstates in resting-state eeg: current status and future directions. *Neurosci Biobehav Rev.* 2015; 49:105–113. [PubMed: 25526823]
- Kiviniemi V, Vire T, Remes J, Elseoud AA, Starck T, Tervonen O, Nikkinen J. A sliding time-window ica reveals spatial variability of the default mode network in time. *Brain Connect.* 2011; 1:339–347. [PubMed: 22432423]
- Kucyi A, Hove MJ, Esterman M, Hutchison RM, Valera EM. Dynamic brain network correlates of spontaneous fluctuations in attention. *Cereb Cortex Feb.* 2016; 13:bhw029.
- Laird AR, Fox PM, Eickhoff SB, Turner JA, Ray KL, McKay DR, Glahn DC, Beckmann CF, Smith SM, Fox PT. Behavioral interpretations of intrinsic connectivity networks. *J Cogn Neurosci.* 2011; 23:4022–4037. [PubMed: 21671731]
- Lancichinetti A, Fortunato S. Limits of modularity maximization in community detection. *Phys Rev E Stat Nonlinear Soft Matter Phys.* 2011; 84:066122.
- Laumann TO, Snyder AZ, Mitra A, Gordon EM, Gratton C, Adeyemo B, Gilmore AW, Nelson SM, Berg JJ, Greene DJ, et al. On the stability of bold fmri correlations. *Cerebral Cortex.* 2016
- Liegeois R, Laumann TO, Snyder AZ, Zhou HJ, Yeo BT. Interpreting temporal fluctuations in resting-state functional connectivity mri. 2017:135681. bioRxiv.
- Logothetis NK. What we can do and what we cannot do with fmri. *Nature.* 2008; 453:869–878. [PubMed: 18548064]
- Ma S, Calhoun VD, Phlypo R, Adali T. Dynamic changes of spatial functional network connectivity in healthy individuals and schizophrenia patients using independent vector analysis. *NeuroImage.* 2014; 90:196–206. [PubMed: 24418507]
- Macke JH, Murray I, Latham PE. How biased are maximum entropy models? *Adv. Neural Inf Process Syst.* 2011:2034–2042.
- Mattar MG, Cole MW, Thompson-Schill SL, Bassett DS. A functional cartography of cognitive systems. *PLoS Comput Biol.* 2015; 11:e1004533. [PubMed: 26629847]
- Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct.* 2010; 214:655–667. [PubMed: 20512370]
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of state calculations by fast computing machines. *J Chem Phys.* 1953; 21:1087–1092.
- Meunier D, Lambiotte R, Fornito A, Ersche KD, Bullmore ET. Hierarchical modularity in human brain functional networks. *Front Neuroinformatics.* 2009; 3:37.
- Mucha PJ, Richardson T, Macon K, Porter MA, Onnela JP. Community structure in time-dependent, multiscale, and multiplex networks. *Science.* 2010; 328:876–878. [PubMed: 20466926]
- Murphy K, Birn RM, Handwerker DA, Jones TB, Bandettini PA. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *NeuroImage.* 2009; 44:893–905. [PubMed: 18976716]
- Newman MEJ. Modularity and community structure in networks. *Proc Natl Acad Sci USA.* 2006; 103:8577–8696. [PubMed: 16723398]
- Porter MA, Onnela JP, Mucha PJ. Communities in networks. *Not Am Math Soc.* 2009; 56(1082–1097):1164–1166.
- Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, Petersen SE. Functional network organization of the human brain. *Neuron.* 2011; 72:665–678. [PubMed: 22099467]
- Roy D, Sigala R, Breakspear M, McIntosh AR, Jirsa VK, Deco G, Ritter P. Using the virtual brain to reveal the role of oscillations and plasticity in shaping brain's dynamical landscape. *Brain Connect.* 2014; 4:791–811. [PubMed: 25131838]
- Royall DR, Lauterbach EC, Cummings JL, Reeve A, Rummans TA, Kaufer DI, LaFrance WCJ, Coffey CE. Executive control function: a review of its promise and challenges for clinical research. A report from the committee on research of the american neuropsychiatric association. *J Neuropsychiatry Clin Neurosci.* 2002; 14:377–405. [PubMed: 12426407]
- Saad ZS, Gotts SJ, Murphy K, Chen G, Jo HJ, Martin A, Cox RW. Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. *Brain Connect.* 2012; 2:25–32. [PubMed: 22432927]

- Sako lu Ü, Pearlson GD, Kiehl KA, Wang YM, Michael AM, Calhoun. A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magn Reson Mater Phys, Biol Med*. 2010; 23:351–366.
- Salvador R, Suckling J, Schwarzbauer C, Bullmore E. Undirected graphs of frequency-dependent functional connectivity in whole brain networks. *Philos Trans R Soc Lond BBiol Sci*. 2005; 360:937–946. [PubMed: 16087438]
- Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*. 2006; 440:1007–1012. [PubMed: 16625187]
- Shirer W, Ryali S, Rykhlevskaia E, Menon V, Greicius M. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb cortex*. 2012a; 22:158–165. [PubMed: 21616982]
- Shirer WR, Ryali S, Rykhlevskaia E, Menon V, Greicius MD. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb Cortex*. 2012b; 22:158–165. [PubMed: 21616982]
- Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky E. The structure of multi-neuron firing patterns in primate retina. *J Neurosci*. 2006; 26:8254–8266. [PubMed: 16899720]
- Siebenhuhner F, Weiss SA, Coppola R, Weinberger DR, Bassett DS. Intra-and inter-frequency brain network structure in health and schizophrenia. *PLoS One*. 2013; 8:e72351. [PubMed: 23991097]
- Smith SM. Fast robust automated brain extraction. *Human Brain Mapp*. 2002; 17:143–155.
- Sporns O, Betzel RF. Modular brain networks. *Annu Rev Psychol*. 2016; 67:613–640. [PubMed: 26393868]
- Stephan KE, Roebroeck A. A short history of causal modeling of fMRI data. *NeuroImage*. 2012; 62:856–863. [PubMed: 22248576]
- Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, Prieto A, Petrusca D, Grivich MI, Sher A, et al. A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci*. 2008; 28:505–518. [PubMed: 18184793]
- Telesford QK, Lynall ME, Vettel J, Miller MB, Grafton ST, Bassett DS. Detection of functional brain network reconfiguration during task-driven cognitive states. *NeuroImage*. 2016; 142:198–210. [PubMed: 27261162]
- Treves A, Panzeri S. The upward bias in measures of information derived from limited data samples. *Neural Comput*. 1995; 7:399–407.
- Turk-Browne NB. Functional interactions as big data in the human brain. *Science*. 2013; 342:580–584. [PubMed: 24179218]
- Vakorin VA, Lippe S, McIntosh AR. Variability of brain signals processed locally transforms into higher connectivity with brain development. *J Neurosci*. 2011; 31:6405–6413. [PubMed: 21525281]
- Watanabe T, Hirose S, Wada H, Imai Y, Machida T, Shirouzu I, Konishi S, Miyashita Y, Masuda N. A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nat Commun*. 2013; 4:1370. [PubMed: 23340410]
- Watanabe T, Hirose S, Wada H, Imai Y, Machida T, Shirouzu I, Konishi S, Miyashita Y, Masuda N. Energy landscapes of resting-state brain networks. *Front Neuroinformatics*. 2014a; 8:12. <http://dx.doi.org/10.3389/fninf.2014.00012>.
- Watanabe T, Kan S, Koike T, Misaki M, Konishi S, Miyauchi S, Miyahsita Y, Masuda N. Network-dependent modulation of brain activity during sleep. *NeuroImage*. 2014b; 98:1–10. [PubMed: 24814208]
- Watanabe T, Masuda N, Megumi F, Kanai R, Rees G. Energy landscape and dynamics of brain activity during human bistable perception. *Nat Commun*. 2014c:5.
- Weiss SA, Bassett DS, Rubinstein D, Holroyd T, Apud J, Dickinson D, Coppola R. Functional brain network characterization and adaptivity during task practice in healthy volunteers and people with schizophrenia. *Front Hum Neurosci*. 2011; 5:81. [PubMed: 21887140]
- Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zollei L, Polimeni JR, Fischl B, Liu H, Buckner RL. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol*. 2011; 106:1125–1165. [PubMed: 21653723]

- Yu Q, Erhardt EB, Sui J, Du Y, He H, Hjelm D, Cetin MS, Rachakonda S, Miller RL, Pearlson G, Calhoun VD. Assessing dynamic brain graphs of time-varying connectivity in fMRI data: application to healthy controls and patients with schizophrenia. *NeuroImage*. 2015; 107:345–355. [PubMed: 25514514]
- Zalesky A, Fornito A, Cocchi L, Gollo LL, Breakspear M. Time-resolved resting-state brain networks. *Proc Natl Acad Sci*. 2014; 111:10341–10346. [PubMed: 24982140]
- Zhou Q. Random walk over basins of attraction to construct ising energy landscapes. *Phys Rev Lett*. 2011; 106:180602. [PubMed: 21635078]

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.neuroimage.2017.05.067>.

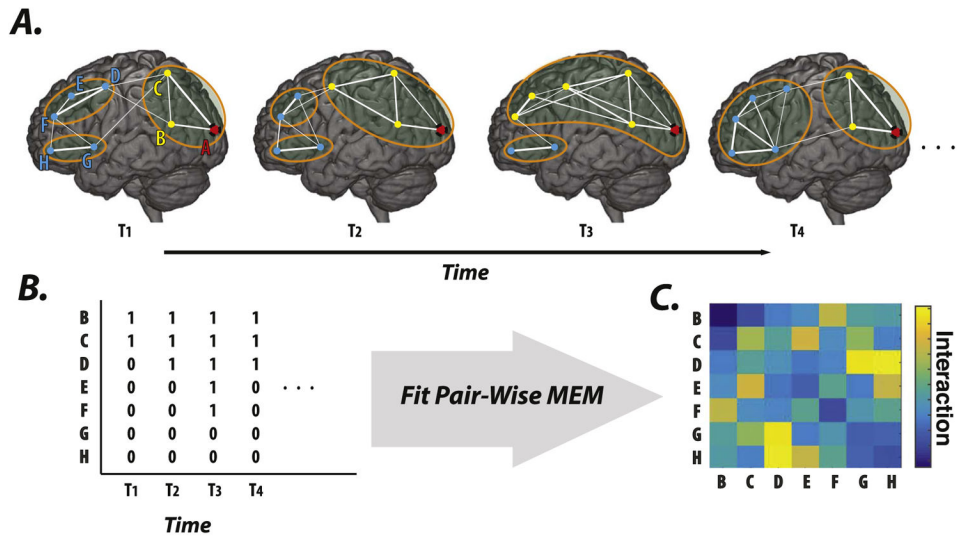
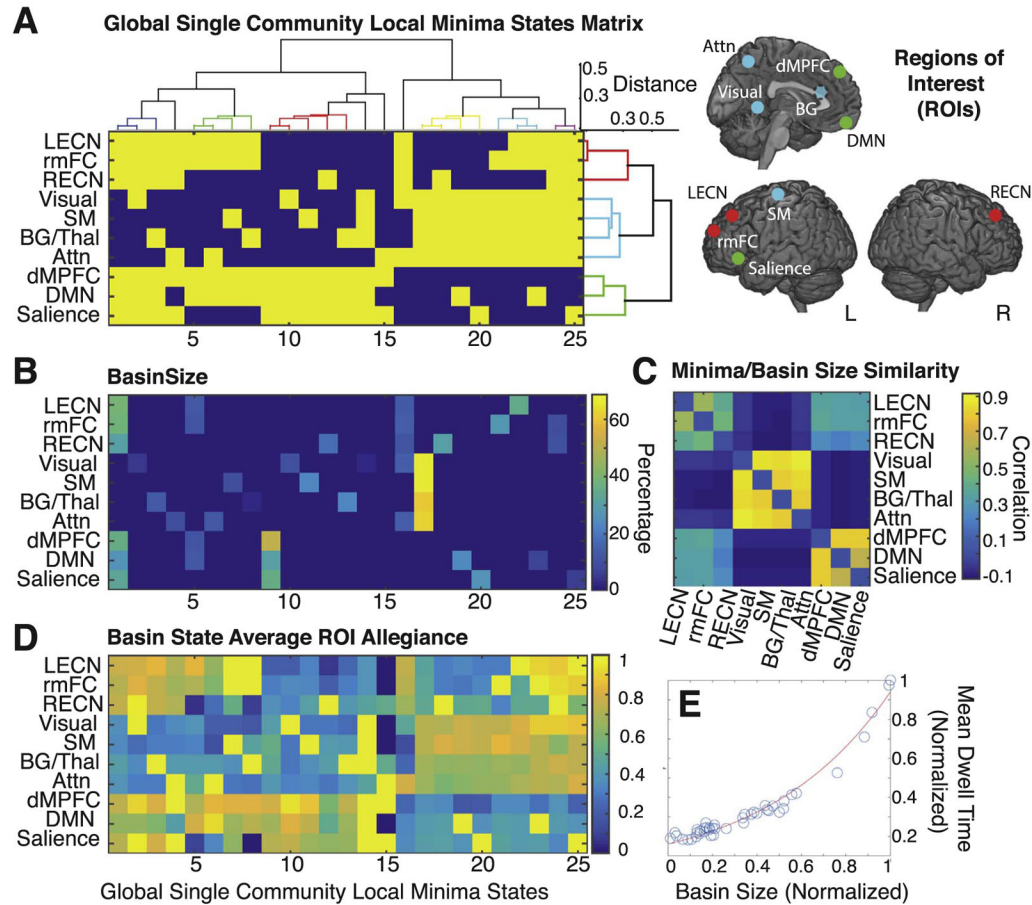
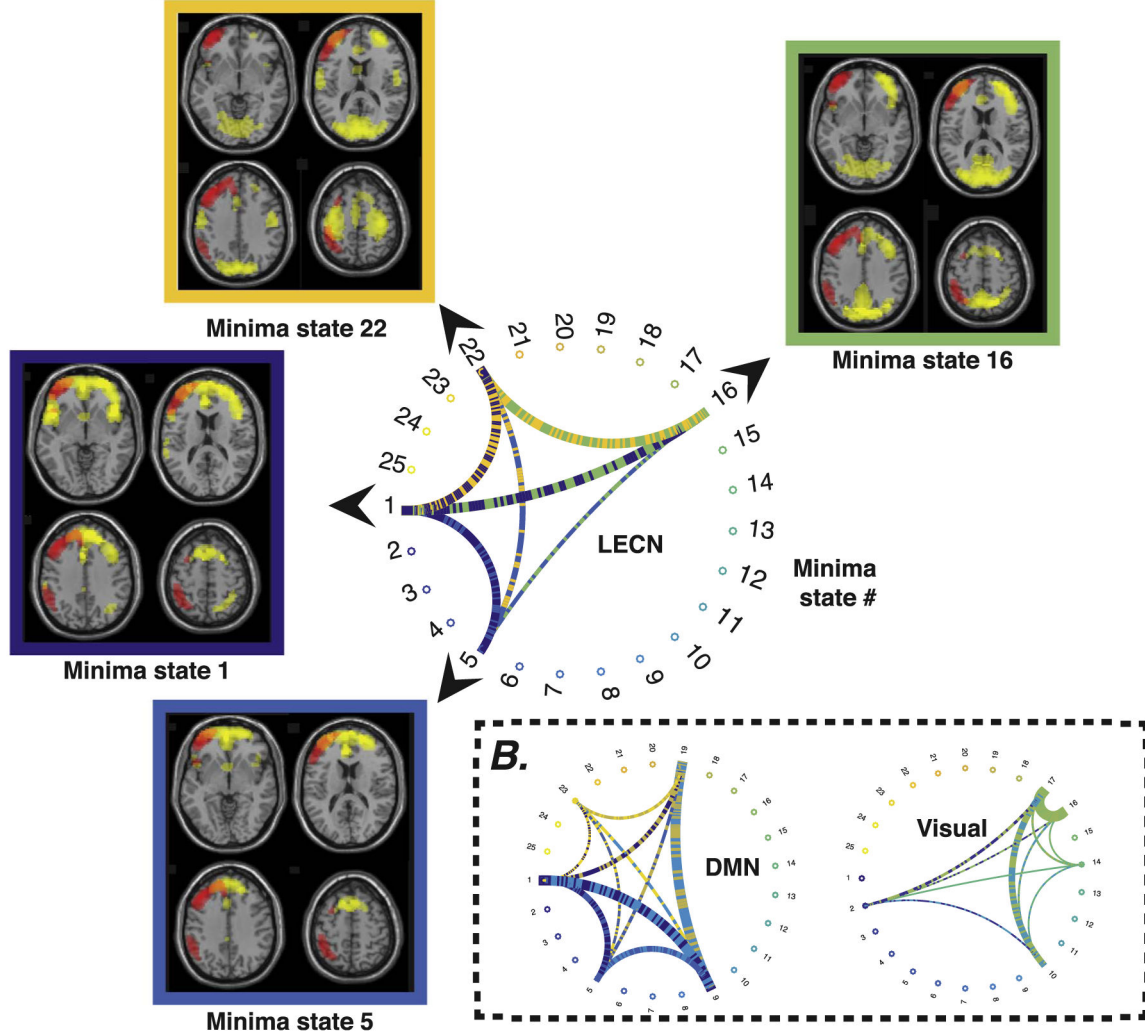


Fig. 1. Schematic of local community allegiance energy landscape estimation. (A) Fluctuations of the strength of functional connectivity between brain regions over time manifests as reconfiguration of the brain’s functional modules. (B) The state of the community allegiance of a single region (e.g., node A) with the rest of the brain regions (‘1’ when both pairs appear in the same community (yellow nodes) and ‘0’ otherwise (blue nodes)) are then used to establish local functional module allegiance states. (C) We fit a MEM using these allegiance state vectors and estimate the functional interaction strength between brain regions to construct an energy landscape of regional community allegiance states.

**Fig. 2.**

Local Minima in the Brain's Functional Energy Landscape. (A) We identified $M = 25$ unique local minima states characteristic of the time series of a region's allegiance to putative functional modules over time. Here, each state represents the set of regions that are commonly allied together in a single community. Using hierarchical clustering, we identified three classes of ROIs with similar patterns of community allegiance across these local minima. Class-I (blue), Class-II (red), and Class-III (green) branches are shown in the dendrogram on the right side of the matrix and visualized on the brain images as ROIs. Using the same clustering technique applied to the matrix transpose, we identified several classes of minima states with common ROI members, denoted by the colored branches on the top of the matrix. ICN labels and their corresponding Lausanne atlas ROI are provided in Table 1. (B) In addition to identifying the unique minima states, for each ROI we also calculated the *basin size* of local minima states as the fraction of the number of basin states to the number of total possible states. Note that the sizes of the analogous basins slightly vary across some ROIs since each local energy landscape is estimated using only 9 out of the 10 ROIs. (C) Calculating the Pearson correlation coefficient between any two ROIs' vectors of basin size across minima states revealed groups of ROIs with similar energy landscapes. (D) The basin states' average vectors highlight the unifying features of the basin states, i.e. omni-present core ROIs (average value 1), commonly present core ROIs (average value between 0.5 and 1) that tend to be among the minima states' member ROIs, and finally the

ROIs with incongruent membership (average value <0.5). (*E*) The average dwell time of minima state's basins (all ROIs combined) grows exponentially with respect to the basins' size. The close exponential fit (red curve) highlights this relationship.

A. ROIs' empirical state transition frequencies**Fig. 3.**

Empirically Estimated Transition Frequencies Between Basins. (A) The circular graph represents the empirical basin transition frequency pattern of a sample Class-II ROI representing the left hemisphere executive control network (LECN). Each color coded dot along the circumference of the circular graph represents one of the 25 local minima states. Lines linking local minima states indicate empirically estimated transition frequencies between local minima states's basins; the width of the line is proportional to the empirically estimated frequency of that transition. The length of the lines and the width of the vertical colored bars contain no quantitative information and they serve only to enhance the visualization of the connected minima states. Transition frequencies are separately normalized for ROIs by dividing by the largest transition frequencies calculated for each ROI. All ICNs with community allegiance congruent with LECN (red overlay) are represented with yellow brain overlays ($z > 1.5$) for all 4 local minima states. (B) The empirical basin transition frequency pattern for two sample ROIs from Class-I (Visual) and Class-III (DMN). Note that unlike LECN, the basin transition frequency pattern of the visual

ROI is heavily skewed towards a single state (that is, state 17, which we describe in greater detail in the body of the text).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

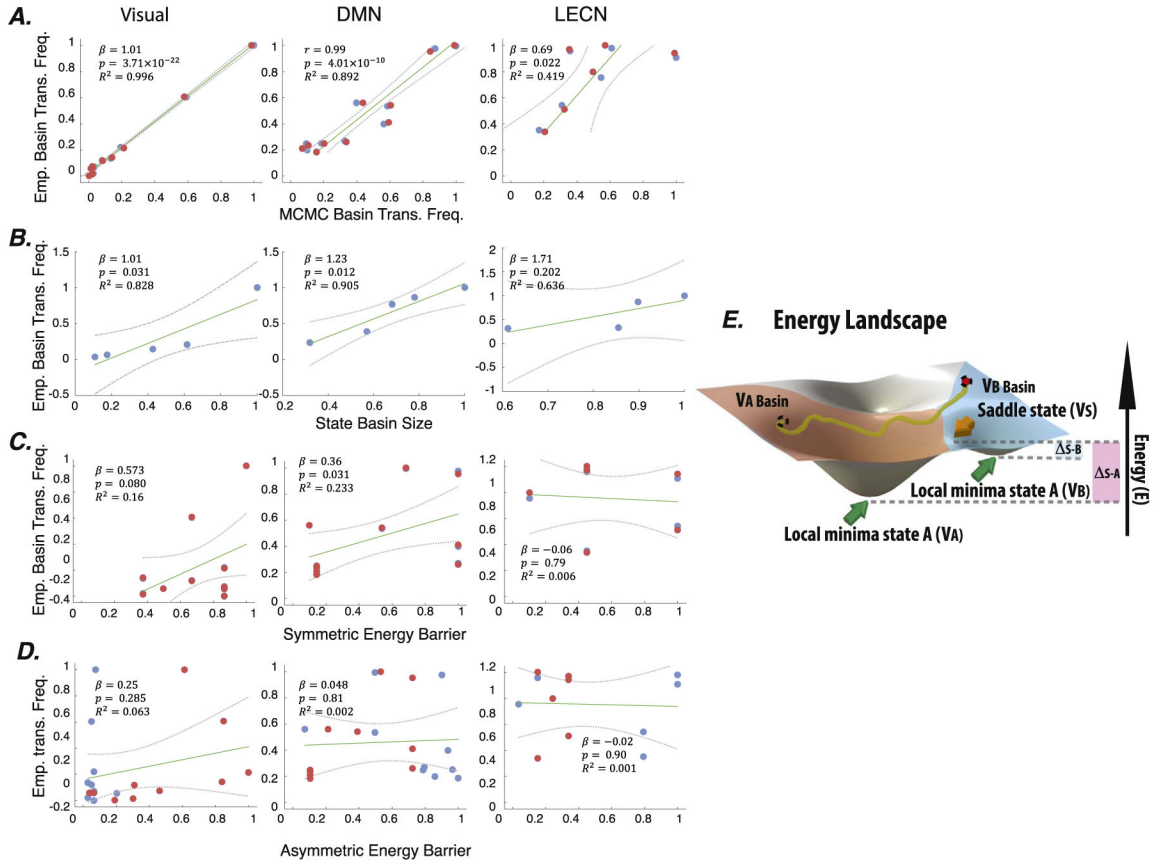


Fig. 4. Empirical versus predicted frequencies of transitioning between basins. (A) We used a Markov Chain Monte Carlo (MCMC) simulation over the energy landscape of each ROI via the Metropolis-Hastings algorithm to estimate theoretically predicted frequencies of transitioning between basins of local community allegiance dynamics of ROIs. The normalized empirical and model transition frequencies were significantly correlated for Class-I(left) and Class-III (middle) ROIs but not for Class-II(right) ROIs. Colored dots represent values on the lower (red) and upper (blue) triangles of the transition matrices (See SI2 for details). The β -value represents the slope of the fitted line, the p -value represents the significance of the fit, and R^2 represents the corrected R^2 values for all of the regressions (B) We define the net empirical basin transition frequencies of minima states as the total in and out transition frequencies. These net frequencies are strongly correlated with the size of the basin surrounding each minima state, particularly for Class-I and Class-III ROIs. (C) Relationship between the empirical basin transition frequency and the predicted energy barriers estimated from a symmetrized transition frequency matrix. (D) Relationship between the empirical basin transition frequency and the predicted energy barriers estimated from the complete asymmetric transition frequency matrix. (E) A two dimensional cartoon of an energy landscape with two local minima states V_A and V_B . We also show their basins (color coded in red and blue, respectively) as well as the saddle point between the two basins (V_S). Any state transition trajectory from basin A to B travels through at least one state with energies higher than the saddle point state (unless the trajectory includes the saddle point

state). Therefore the energy barrier for getting out of the local minima state is calculated by subtracting the energy values between the saddle state and the local minima state. The symmetric energy barrier between two local minima states is calculated as the minimum of the two energy barriers. Although we have only examined the basin transition frequency in ($A-D$), the results are similar for basin transition probabilities since the two are related to one another by a simple normalization factor (see SI2 Fig. 2).

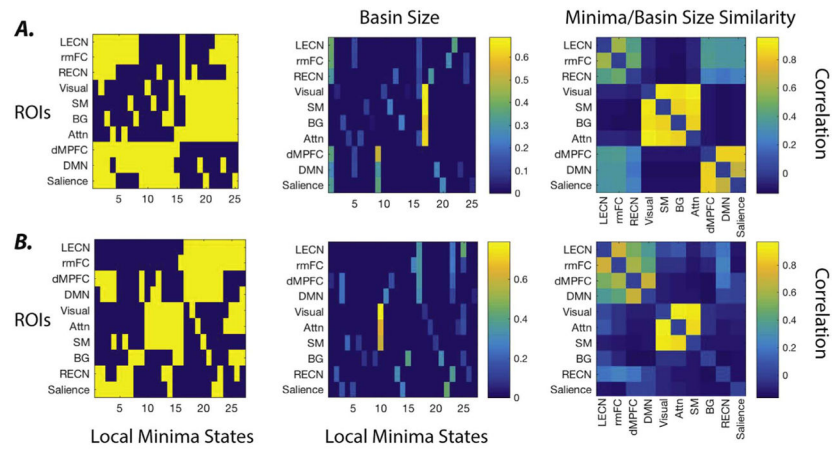
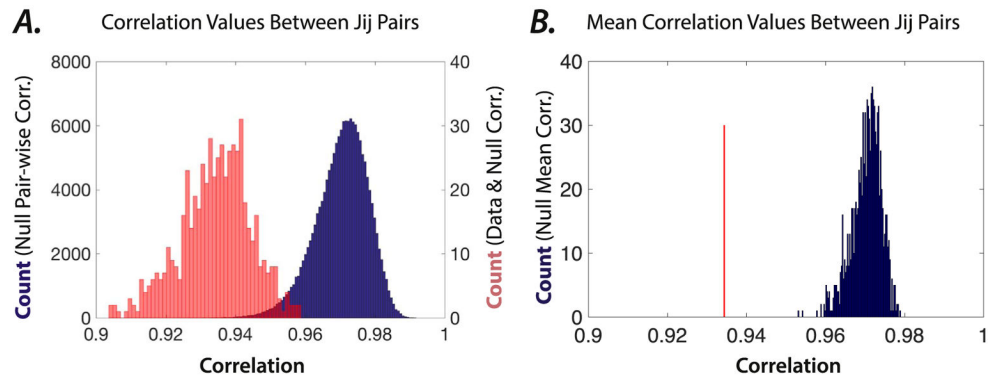


Fig. 5. Energy landscape of the surrogate time series. Here we show (Left) the local minima, (Middle) the basin size of local minima states as the fraction of the number of basin states to the number of total possible states, and (Right) the Pearson correlation coefficient between any two ROIs' vectors of basin size across minima states of the original time series. (B) Similar results are presented for the simulated time series created using the method proposed by (Laumann et al., 2016).

**Fig. 6.**

Statistical comparison of the empirically observed interaction matrix to that expected in a null model. (A) Distribution of the correlation values between pairs of interaction matrices estimated from the null and original datasets. The blue histogram displays the distribution of the correlation values between all pairs of interaction matrices estimated from the null time series. The orange histogram displays the distribution of the correlation values between the empirical and all interaction matrices estimated from the null time series. Note that the empirical interaction matrix is notably less similar to the null matrices. (B) Distribution of the mean correlation values calculated between a single null and the rest of the null matrices (repeated for each individual null matrix). The red line shows the mean correlation value calculated between the empirical matrix and all null matrices. Note that the empirical mean is significantly smaller than that expected in the null distribution ($p < 0.00095$).

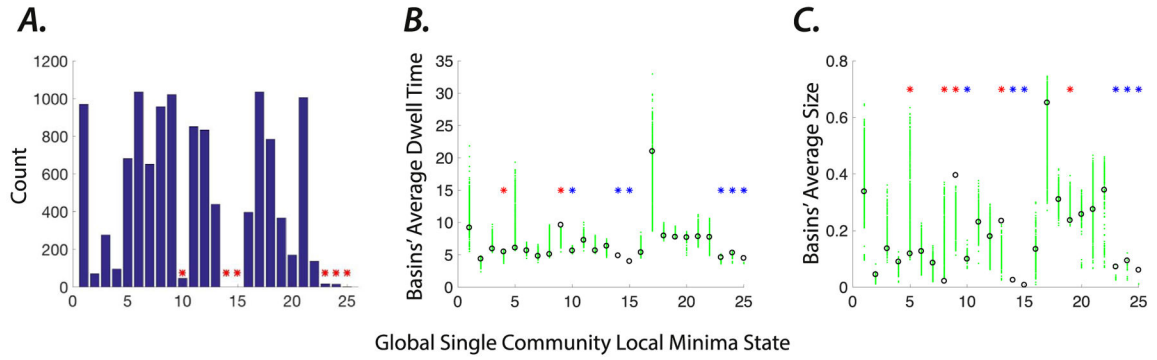


Fig. 7.

Comparison of summary statistics between the empirical and null energy landscapes. (A) The number of times that each local minima state identified from the empirical data is found in 1046 null energy landscapes. The minima states that are observed less frequently than chance (one-tailed, $p < 0.05$) in the null energy landscapes are marked by a red '*'. (B) The average dwell time of each local minima states' basin calculated from the empirical energy landscapes as well as from the null energy landscapes is marked by a black 'o' and a green '.', respectively. The minima states that display significantly (one-tailed, $p < 0.05$) larger or smaller dwell times compared to that of the null energy landscapes are marked by a red '*'. Similar to the results displayed in panel A, we mark the minima states that are observed less frequently than expected in the null energy landscapes by a blue '*'. (C) The average size of each local minima states' basin calculated from the empirical energy landscape as well as from the null energy landscapes is marked by a black 'o' and a green '.', respectively. The minima states that display significantly larger or smaller basin size compared to that of the null energy landscapes are marked by a red '*'. Similar to the results displayed in panels A and B, we mark the minima states that are observed less frequently than expected in the null energy landscapes by a blue '*'. These data support the notion that the null and empirical energy landscapes are fundamentally different from one another.

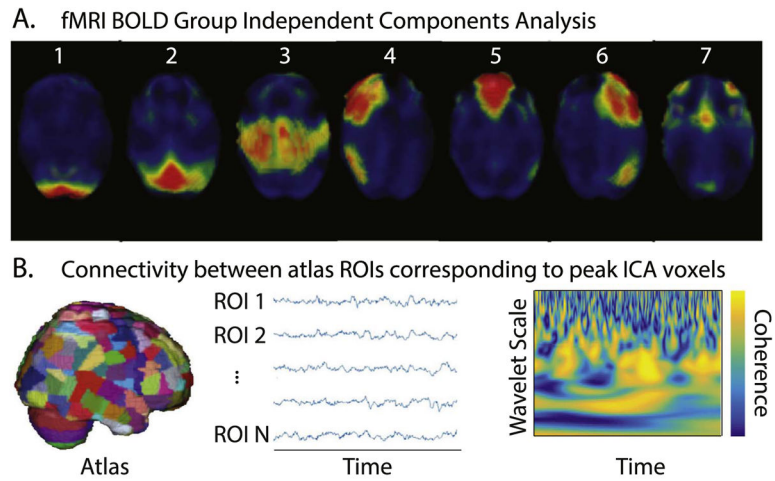


Fig. 8. Schematic of Methods. (A) We used a group-ICA decomposition to distill fMRI resting state BOLD into $N(=10)$ components representing putative baseline functional networks. For each ICN, we identified the voxel with the peak expression of that component. (B) Using the Lausanne 125 scale template (234 ROIs) (Cammoun et al., 2012b), we determined the atlas region corresponding to the peak expression of each component. After extracting BOLD time series from each ROI, we estimated the functional connectivity between pairs of ROIs using the wavelet coherence in the frequency interval 0.19–0.06 Hz (Bassett et al., 2011).

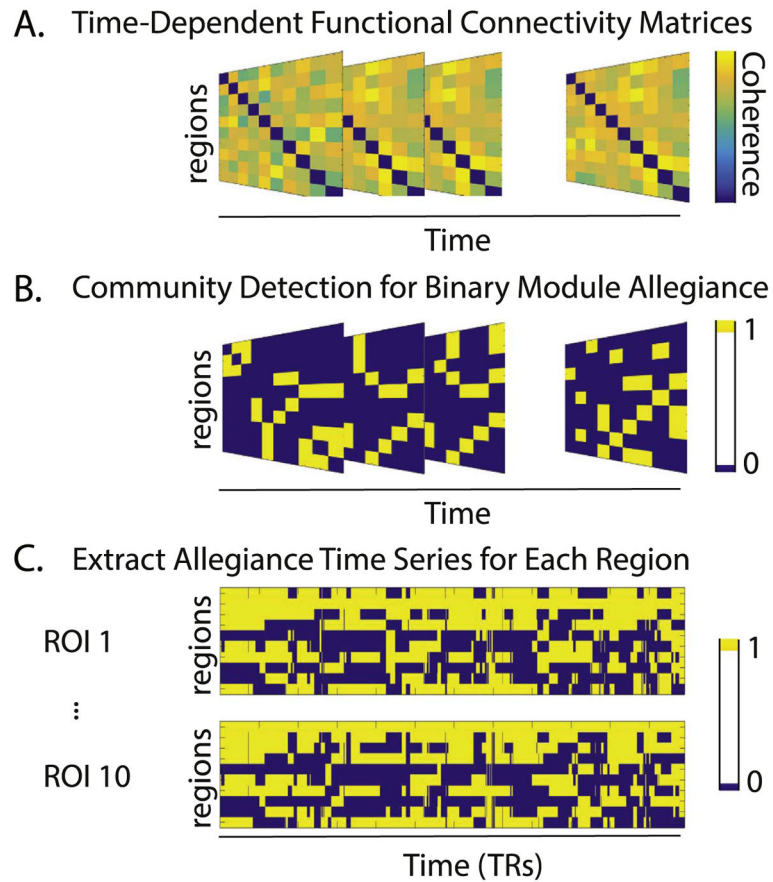



Fig. 9. Schematic of Methods for Extracting Dynamic Module Time Series. (A) We represent the T ($= 1190$ (TRs) $\times 20$ (subjects) $\times 4$ (runs) $= 95200$) unique functional connectivity patterns as $N \times N$ adjacency matrices \mathbf{A} . (B) Using community detection, we extract putative functional modules at each TR, and use a statistical comparison to a random null model to determine a region's binary module allegiance. More specifically, we identify the community organization of ROIs and calculate the probability of ROI pairs' congruent community allegiance for each time-point. The ROI pairs with higher than expected (via permutation tests) congruent community allegiance were thresholded to generate binarized pairwise allegiance matrices. (C) We reformat these data to separately store the allegiance time series of each ROI, which codes its co-allegiance with other ROIs to the same community (values of 1) as a function of time (TR).

A. Merge data from all subjects and fit a pair-wise MEM



B. ROI's Dynamic Community Allegiance Energy Landscapes

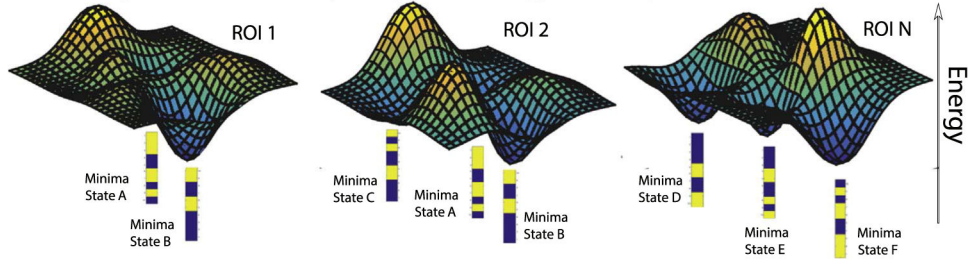


Fig. 10. Schematic of Maximum Entropy Model of Brain Network Dynamics. (A) Accurate fitting of a MEM requires large amounts of data. We therefore combined data across all subjects before fitting a pairwise MEM to each ROI's pattern of allegiance to functional modules. This fitting procedure produced an estimated interaction matrix for each ROI and each TR; colors indicate the strength of each element of the interaction matrix J_{ij} . (B) From the interaction matrices, we defined and characterized energy landscapes of the local community dynamics. Color indicates energy, with yellow indicating high energy and dark blue indicating low energy. Each minimum within each landscape is accompanied by an example network state, as defined by a binarized pattern of module allegiance (yellow indicating congruent module allegiance and dark blue indicating incongruent module allegiance) with other ROIs.

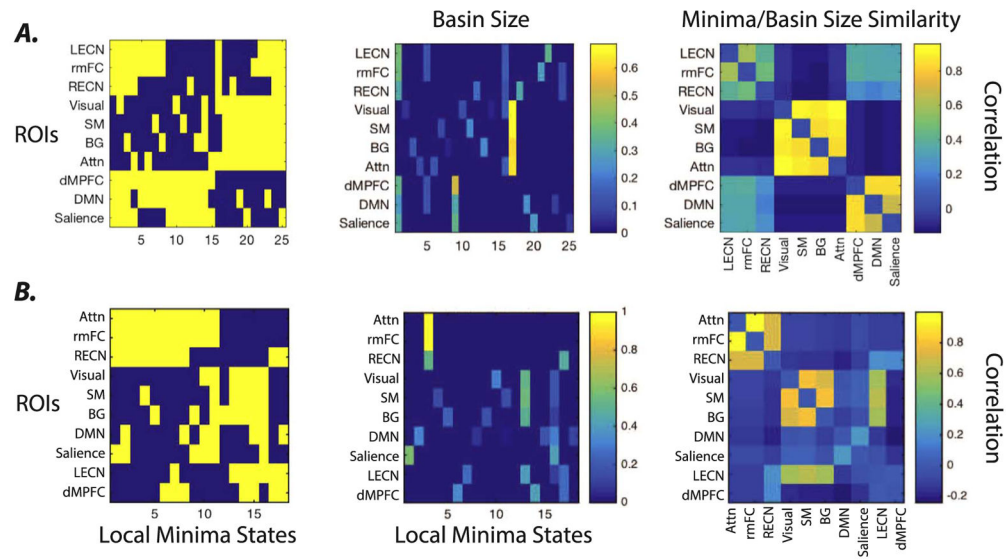


Fig. 11. Energy landscape of random ROIs time series. (A) Here we show (Left) the local minima, (Middle) the basin size of local minima states as the fraction of the number of basin states to the number of total possible states, and (Right) the Pearson correlation coefficient between any two ROIs’ vectors of basin size across minima states of the original time series. (B) Similar results are presented for the time series extracted from a random set of ROIs from ICA components.

Table 1

Regions of interest and their corresponding ICN.

ICN	Lausanne ROI (scale 125)
1. Visual	189 Cuneus.1
2. Dorsal Attention (Attn)	184 Precuneus.1
3. Sensory/Motor (SM)	147 Precentral.3
4. Basal Ganglia/Thalamus (BG)	228 Caudate
5. Left Executive Control Network (LECN)	128 Rostral middle frontal.2
6. Right Executive Control Network (RECN)	15 Rostral middle frontal.2
7. Rostral Middle Frontal Cortex (rmFC)	130 Rostral middle frontal.4
8. Dorsomedial Prefrontal Cortex (dMPFC)	135 Superior frontal.3
9. Default Mode Network (DMN)	122 Medial orbitofrontal.1
10. Salience	124 Pars triangularis.1

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript