

On co-activation pattern analysis and non-stationarity of resting brain activity

Teppei Matsui^{1,2}, Trung Quang Pham³, Koji Jimura⁴, Junichi Chikazoe^{3,5}

1. Department of Biology, Okayama University, Okayama, Japan
2. JST-PRESTO, Japan Science and Technology Agency, Tokyo, Japan
3. Section of Brain Function Information, Supportive Center for Brain Research, National Institute for Physiological Sciences, Okazaki, Japan
4. Keio University, Yokohama, Japan
5. Araya Inc., Tokyo, Japan

Corresponding Author:

Teppei Matsui (tematsui@okayama-u.ac.jp)

Acknowledgements

We thank T. Ichikawa and M. Taki for discussion. This study was supported by JSPS Kakenhi (20H05052 to TM, 19K20390 to TQP, 19H04914 and 20K07727 to KJ, 21H02806 to JC); a grant from Japan Agency for Medical Research and Development (AMED) to JC (grant number JP21dm0207086). a grant from Brain/MINDS Beyond (AMED) to TM (grant number JP20dm0307031); a grant from JST-PRESTO to TM; a grant from Narishige Neuroscience Research Foundation to TM.

Abstract

The non-stationarity of resting-state brain activity has received increasing attention in recent years. Functional connectivity (FC) analysis with short sliding windows and coactivation pattern (CAP) analysis are two widely used methods for assessing the non-stationary characteristics of brain activity observed with functional magnetic resonance imaging (fMRI). However, whether these techniques adequately capture non-stationarity needs to be verified. In this study, we found that the results of CAP analysis were similar for real fMRI data and simulated stationary data with matching covariance structures and spectral contents. We also found that, for both the real and simulated data, CAPs were clustered into spatially heterogeneous modules. Moreover, for each of the modules in the real data, a spatially similar module was found in the simulated data. The present results suggest that care needs to be taken when interpreting observations drawn from CAP analysis as it does not necessarily reflect non-stationarity or a mixture of states in resting brain activity.

Introduction

In contrast to initial studies of resting-state functional connectivity (FC), which assumed FC was stable throughout a relatively long scan-duration, recent studies are increasingly focusing on non-stationary aspects of resting brain activity (Hutchison et al., 2013). This dynamic view of resting brain activity hypothesizes that the brain switches between various states within which the activity of different sets of brain regions are coordinated. Sliding window correlation analysis and co-activation pattern (CAP) analysis have been developed to assess such non-stationary characteristics of resting brain activity (Allen et al., 2014; Liu et al., 2013; Liu and Duyn, 2013).

Sliding-window correlation analysis and CAP analysis detect momentary brain activity patterns by using FC within a short sliding-window (~40 sec) (Allen *et al.*, 2014) and brain activations in selected single volumes (Liu *et al.*, 2013; Liu and Duyn, 2013), respectively. The detected momentary brain activity patterns are often heterogeneous, even within the same run, and grouped into distinct clusters of modules. In both sliding-window correlation analysis and CAP analysis, these modules exhibit specific spatial patterns which are therefore interpreted as distinct states of the resting brain (Liu and Duyn, 2013; Liu et al., 2018). Numerous studies have applied these methods to characterize the non-stationary aspects of resting brain activity in healthy individuals, patients with neurological disorders and animals (see (Prete et al., 2016)(Liu *et al.*, 2018) for review).

Importantly, whether the results obtained from sliding-window correlation analysis or CAP analysis truly reflect the non-stationarity of resting brain activity is a crucial issue that needs careful study. In fact, several studies have shown that some results claimed to reflect the non-stationarity of resting brain activity can be replicated from simulated data with stationary null models (Cifer et al., 2017; Laumann et al., 2016; Novelli and Razi, 2021). Although these studies do not prove or claim that resting brain activity is stationary, they provide the necessary and important background required for appropriately interpreting the results obtained with these analyses designed to capture the non-stationarity of resting brain activity. In this report, we add to this discussion by examining whether simulated stationary data can reproduce the spatially heterogeneous modules obtained with CAP analysis (Liu *et al.*, 2018). To our surprise, we found that CAP analysis applied to a stationary null model yielded spatially heterogeneous modules, each of which closely approximated the modules found in the real data.

Materials & Methods

Dataset

We used publicly available data of resting-state fMRI distributed by the Human Connectome Project (HCP; <http://humanconnectomeproject.org/>). The data was a collection of region-of-interest (ROI) time series (1200 volumes \times 264 ROIs \times 64 individuals; Repetition time (TR), 0.72 sec) (Power et al., 2011). Each time course was normalized to have zero mean and unit variance. Global signal regression was conducted by following a standard procedure. For each scan, a global signal time course was obtained by averaging time courses across all ROIs. Then for each ROI, the global signal time course was regressed out from the ROI's time course. The data for group analysis was made by concatenating data of all the individuals in the volume dimension (Liu and Duyn, 2013).

For voxel-based analyses, we selected a single axial slice ($z = 27$ in the MNI coordinate) containing the posterior parietal cortex (PPC) in the HCP resting-state data to mitigate computational demand. The data of 60 individuals were used in the analysis. All data were preprocessed according to the HCP's standard pipeline. To select the voxels corresponding gray matter, we made a gray matter mask for each individual using the segmentation program implemented in SPM12 (<https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>; threshold set at 0.7) and then took a union of the masks. The obtained mask was applied to each axial slice to extract 2320 voxels corresponding to the gray matter.

Static FC

Static FC was calculated by taking the correlation coefficient between the time courses of two ROIs using all volumes in the concatenated runs. For the voxel-based analyses, an ROI centered at the PPC (6 mm by 6 mm centered at MNI coordinates [0, -53, 27]) was used, and the temporal correlation was calculated between the ROI and each voxel.

Generation of simulated data

The simulated data were generated by adapting previously described methods and codes (Laumann *et al.*, 2016). Briefly, random samples were drawn from a Gaussian distribution with dimensions matching the real data. These time courses were multiplied in the spectral domain by the spectrum derived from the real data. The time courses were then projected onto eigenvectors of the covariance matrix derived from the real data. This procedure enabled the construction of simulated time courses, stationary by construction,

with covariance structures and spectral contents equivalent to the real data. Simulated time courses were generated for each run for each individual. The results were similar when the data was generated for the real data concatenated across individuals (data not shown). Global signal regression was similarly performed for the real data. Group data was then obtained by concatenating all the simulated data across individuals.

CAP analysis

The CAP analysis was conducted according to standard procedures (Liu and Duyn, 2013; Liu *et al.*, 2018). For each ROI time course, time points exceeding the percentile threshold (top 15%) were collected. The set of volumes corresponding to these time-points were defined as CAPs (using the ROI as Seed). In the group analysis, for each chosen ROI, CAPs were selected using the concatenated time courses. For the voxel-based analyses, the same PPC ROI used to calculate the static FC was used. Average CAP (for a Seed ROI) was obtained by averaging across all detected CAPs. Modules were extracted by k-means clustering of the CAPs using the correlation distance. The number of clusters was set to eight as in the original study (Liu and Duyn, 2013). The similarity of the modules, both within and across data types, was measured by calculating the ROI-wise or voxel-wise correlation between two modules.

Statistical comparison of the real and simulated CAPs

For each pair of real and simulated CAPs obtained with a seed ROI, we tested the null hypothesis that the two sets of random multivariate variables, *i.e.*, the real and simulated CAPs, were drawn from the same distribution using energy statistics (Szekely and Rizzo, 2013). We used a Matlab implementation provided by Dr. Brian Lau (<https://github.com/brian-lau/highdim>) to analyze the energy statistics. Because of a limitation in computational power, the real and simulated CAPs were each subsampled to 1000 samples before being subjected to the energy test. Covariance matrices were compared by taking the correlation coefficient of the off-diagonal elements.

Results

A small number of volumes suffices to approximate static FC in both real and simulated data

A key observation in the CAP analysis was that static FC could be closely approximated with a small fraction of time points. Because this observation was the conceptual starting point of the CAP analysis (Liu and Duyn, 2013), we first examined if this property was genuine to the real fMRI data or reproducible with the simulated stationary data. Figure 1 shows comparisons between static FC, the average CAP of the real data, and the average CAP of the simulated data whose covariance structure and spectral contents were matched to the real data. Supplementary Figure 1 shows similar comparisons for two representative individuals. These results suggest that a small number of volumes can be used to approximate static FC for both real and simulated data. Consistent results were found for CAPs obtained using the top 5% as the threshold (Supplementary Figure 2) and for the analysis without global signal regression (Supplementary Figure 3). Thus, these results suggest that the stationary null model replicated the observation that a small fraction of time points was sufficient to approximate static FC.

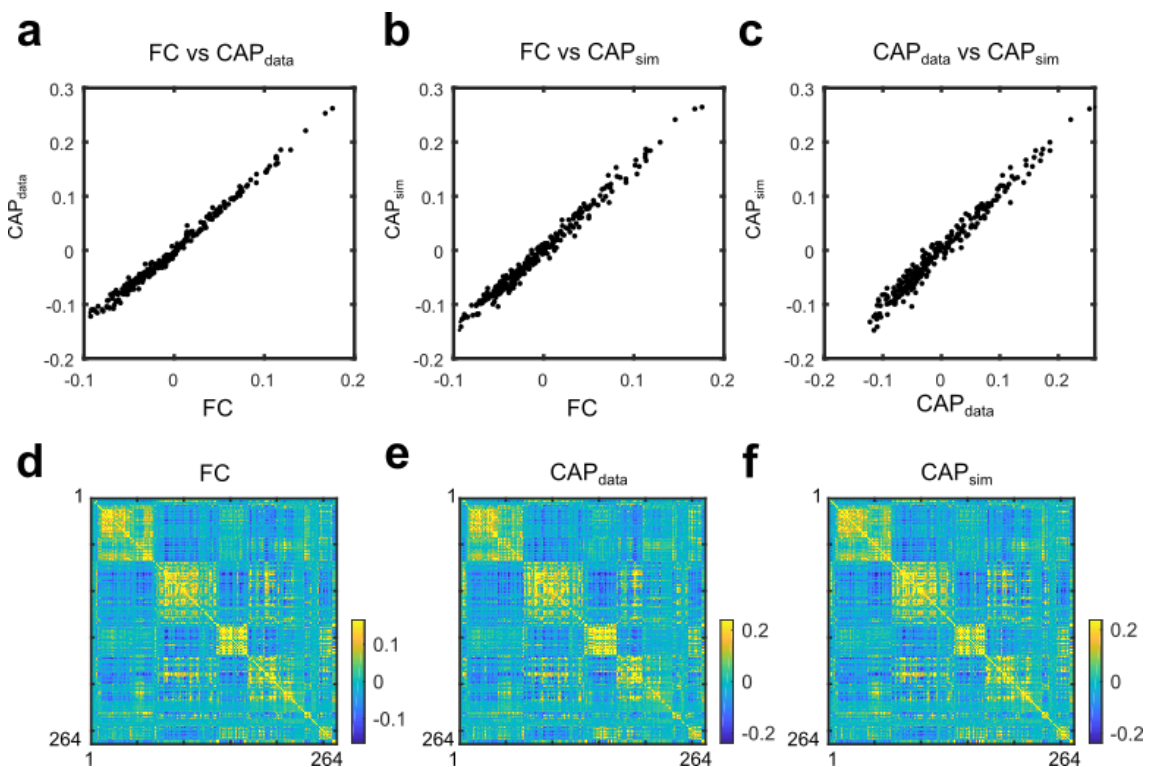


Figure 1 The average CAPs of the real and simulated data closely approximated

static FC. a-c. Scatter plots of a representative ROI showing close correspondence between FC, CAP_{data}, and CAP_{sim}. **d-e.** ROI by ROI matrices of FC, CAP_{data}, and CAP_{sim}.

Both real and simulated CAPs can be clustered into heterogeneous modules

Another key finding from the initial CAP analysis was that CAPs could be clustered into modules with distinct spatial patterns. The presence of these modules was interpreted as distinct states of resting-state brain dynamics and regarded them as an indication of non-stationarity (Liu and Duyn, 2013). If this interpretation were true, CAPs obtained in the stationary null model would not yield spatially heterogeneous modules. To test this, we clustered CAPs into modules for both the real and simulated data. Figure 2 shows the similarity between modules, within the same data type, for the real and simulated data. The clustering of CAPs resulted in spatially heterogeneous modules (*i.e.*, low spatial correlation between modules) for both real and simulated data. Thus, the presence of spatially heterogeneous modules cannot be taken as evidence for non-stationarity in resting brain activity.




Figure 2 Spatially heterogeneous modules were found both for the real and simulated CAPs. Colored matrices show similarities between modules obtained by clustering the real or simulated CAPs. Left, real data. Right, simulated data.

This result seemingly contradicts the original report by Liu and Duyn where they found no modular structure in CAPs obtained with control (simulated) data generated using normal distributions (Liu and Duyn, 2013). Importantly, unlike the simulation devised by Laumann and colleagues (Laumann *et al.*, 2016), the simulation used by Liu and Duyn did not impose spatial covariance among voxels (*i.e.*, each voxel treated as independent). We found that simulated data with zero spatial covariance among ROIs greatly increased the similarity between CAP modules (Supplementary Figure 4), replicating the results described by the previous study. Thus, the apparent contradiction

between the present simulation results (Figure 2) and Liu and Duyn's study is due to the difference in null model construction. Because the assumption of zero spatial covariance among voxels or ROIs is unlikely to hold for real brain activity, we only considered the Laumann-type null model in the following analyses.

Having seen that the stationary null model can produce heterogeneous modules of CAPs, we next proceeded to ask whether modules of CAPs obtained from the real and simulated data were similar. If the real and simulated modules were alike, it would mean that the spatial pattern of modules is primarily determined by static properties of the real data, *i.e.*, the covariance structure and spectral contents. Figure 3 shows (between datasets) a comparison of the modules obtained from the real and simulated data. Unexpectedly, for each module obtained from the real data, a similar module could be found from the simulated data (Figure 3a). To quantify the similarity of the two sets of modules, we matched the two sets to maximize the mean of the diagonal elements of the correlation matrix (Figure 3b). The similarity index defined by the mean of the diagonal elements was 0.83. Importantly, this value was within the range of the similarity index calculated by comparing two sets of modules derived from two independent simulations (mean \pm SD, 0.81 ± 0.045 , $n = 100$ pairs of simulations; Figure 3c), indicating that the similarity between the real and simulated modules was close to the noise ceiling set by statistical sampling error. A high degree of similarity was found between the real and simulated modules across all ROIs (mean \pm SD, 0.81 ± 0.037 , $n = 264$ ROIs; Figure 3d). Taken together, these results suggest that CAPs were clustered into spatially heterogeneous modules in both real and simulated data. Moreover, individual modules obtained from the simulated data were similar to modules obtained from the real data, suggesting that the spatial pattern of modules is largely determined by static properties of resting brain activity.




Figure 3 CAP modules derived from real and simulated data were very similar. The similarity between modules was measured by the correlation coefficient. **a.** original ordering as in Figure 2. **b.** same as *a* but the modules were reordered to maximize matching (mean of diagonal elements = 0.83). **c.** similarity distribution of modules between two independent simulations (same ROI as in *a, b*). For each pair of simulations, module matching was performed (as in *b*), and the mean of the diagonal elements was taken as the similarity value. The magenta line indicates the similarity value for the data shown in *a, b*. **d.** similarity distribution of the modules between the real and simulated data tested for all ROIs. The magenta line indicates the similarity value for the data shown in *a, b*.

We replicated the above analyses with voxel-based data to visualize each module and confirm that the spatial averaging for the ROI-based analysis did not cause the similarity between real and simulated data. To mitigate computational demand, we performed this analysis on a single slice containing PPC (see Methods). Overall, the results from the voxel-based analysis were similar to those of the ROI-based analysis described above. Like the results shown in Figure 1a-c, we found the average CAP_{data} and CAP_{sim} maps were similar to the static FC map (Figure 4a). For both the real and simulated data, modules of CAPs were dissimilar to each other within the same data type (Figure 4b,c; see also Figure 2 for comparable results in the ROI-based analysis). Finally and most importantly, simulated modules with similar spatial patterns were found for most of the

modules in the real data (Figure 4d,e). Figure 5 shows spatial maps of the modules shown in Figure 4d,e. These examples visually demonstrate the high degree of similarity between the real and simulated modules. Thus, the voxel-based results corroborate the findings of the ROI-based analyses and further show that the presence or absence of spatial averaging does not change the results.

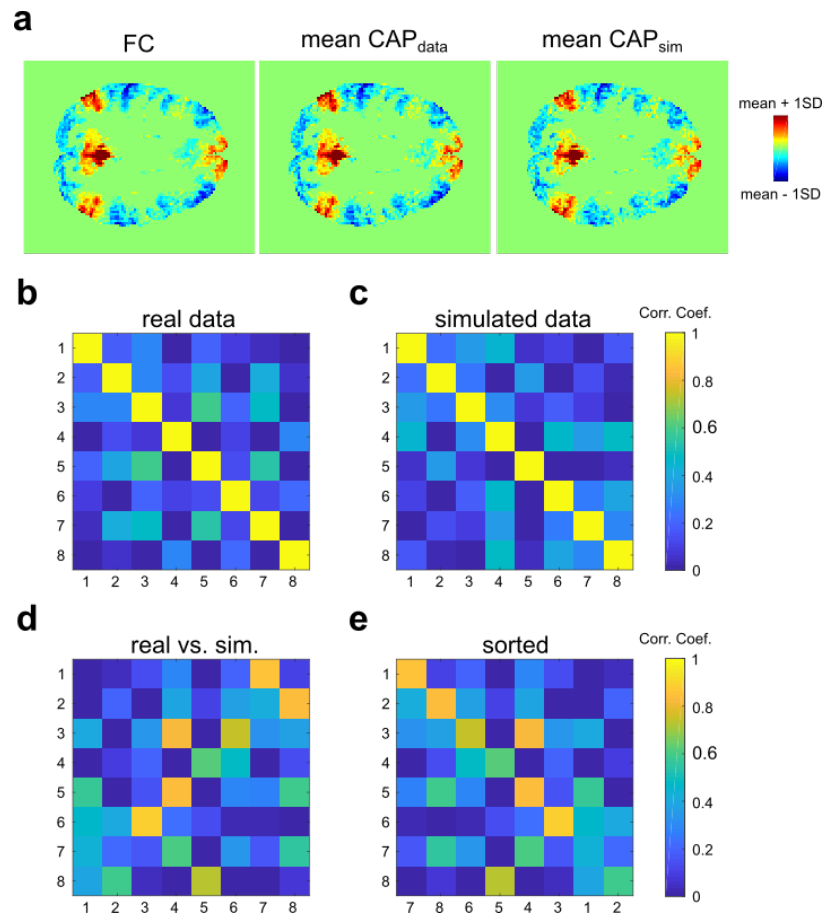


Figure 4 Voxel-based analysis yielded the similar results to the ROI-based analysis.

a. Comparison of the FC map, mean CAP_{data} map, and mean CAP_{sim} map. As in Figure 1, the three maps appear very similar. **b-c.** The similarity between CAP modules within the same data type. Following the same convention as in Figure 2. **d-e.** The similarity between modules across data types. **d.** original ordering as in (b)-(c). **e.** Same as (d) but modules are reordered to maximize matching (mean of diagonal elements = 0.72). See Figure 5 for details of each module.

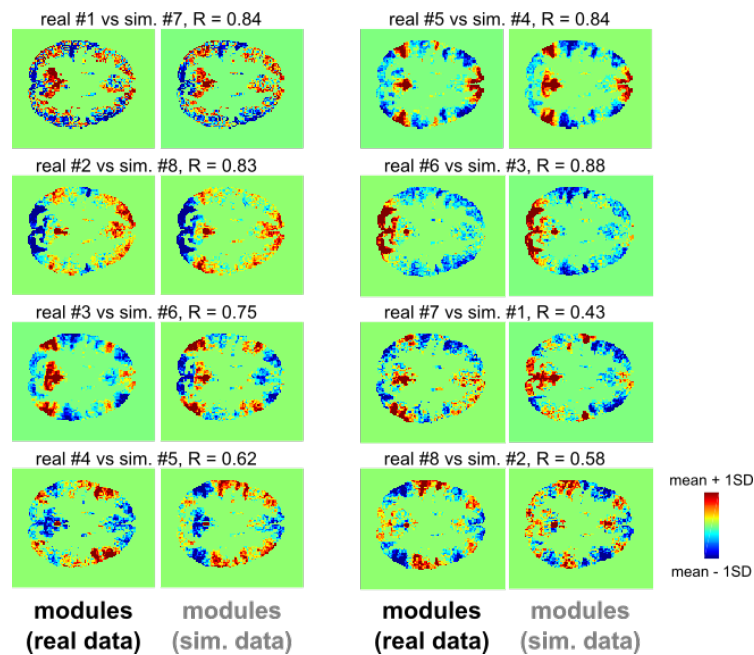


Figure 5 Examples modules derived from the real and simulated data. For the example shown in Figure 4, eight pairs of modules derived from the real data (left) and the simulated data (right) are shown side-by-side. The modules are paired according to Figure 4e.

Why were modules of CAPs in the real data so similar to those in the simulated data? We hypothesized that this observation was because the real and simulated CAPs shared the same underlying statistical distribution. In k-means clustering, data points were grouped into clusters according to the correlation distance between the data points. Evidently, the distribution of the correlation distance was determined by the distribution of the data points. Therefore, if the datasets were derived from the same distribution, provided enough data points were included in each set, the results of k-means clustering should have been similar.

To compare the statistical similarity of the two sets of CAPs, we first compared the covariance matrices. For a representative ROI, we found the covariance matrices of the real and simulated CAPs were very similar ($R = 0.97$; Figure 6a,b). A high degree of similarity was found between covariance matrices of the real and simulated CAPs for all ROIs (mean \pm SD, 0.97 ± 0.0026 , $n = 264$ ROIs; Figure 6c), consistent with the hypothesis that they share the same multivariate distribution. Next, we conducted statistical testing to examine the null hypothesis that the two datasets were drawn from the same multivariate distribution (see Methods). If the real and simulated CAPs were derived from different distributions, the null hypothesis should be rejected. Statistical testing was not

significant for the representative ROI used in Figure 6a ($P > 0.87$). Across all ROIs, p-values were mostly distributed above the typical significance threshold (mean \pm SD, 0.49 ± 0.30 ; Figure 6d; Note that p-values were uncorrected). The null hypothesis was rejected in 20 out of 264 ROIs (7.6% of all ROIs) when the significance threshold was 5%. With a significance threshold of 1%, the null hypothesis was rejected in 7 ROIs (2.7% of all ROIs). Thus, for most ROIs, the statistical distributions of the real and simulated CAPs were similar and hence yielded similar clustering results.

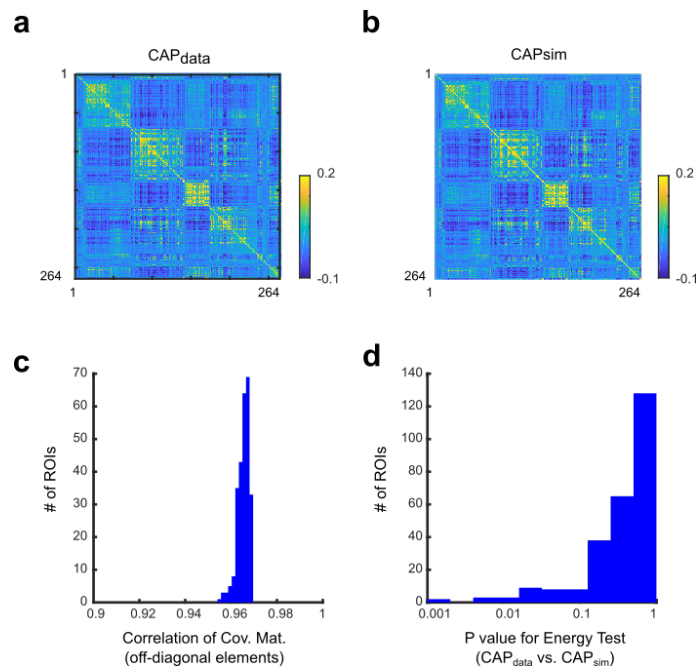


Figure 6 Statistical properties of the real and simulated CAPs were approximately equal. a-b. ROI-by-ROI covariance matrices of the CAP_{data}(a) and CAP_{sim}(b). **c.** The correlation distribution between off-diagonal elements of the covariance matrices of the real and simulated CAPs ($n = 264$ ROIs). **d.** The distribution of p-values from Szekeley & Rizzo's Energy Test, which tests the null hypothesis that CAP_{data} and CAP_{sim} are drawn from the same multivariate distribution (One test per ROI. $n = 264$ ROIs).

Discussion

In this study, we asked to what extent could the results of CAP analysis be attributed to the non-stationarity of resting brain activity. To this end, we conducted CAP analysis on real fMRI data and simulated data based on a stationary null model with matching covariance structures and spectral contents. Overall, we found that two key observations drawn from the CAP analysis were replicated in the simulated data, which has previously been interpreted as evidences for non-stationary resting brain activity. First, in both the real and simulated data, a small percentage of time points were sufficient to approximate FC calculated using all time points (Figure 1 and Figure 4a). Second, both the real and simulated CAPs were classified into spatially dissimilar modules (Figure 2 and Figure 4b,c). More interestingly, it was possible to find a simulated module closely resembling the spatial pattern of most modules obtained from the real data (Figure 3, Figure 4d,e and Figure 5). The fact that key results of CAP analysis were replicated in simulated stationary data suggests that the results need to be interpreted with care.

The present study adds to a series of previous studies reporting that it is difficult to find signatures of the non-stationarity of resting brain activity in fMRI data (Cifer *et al.*, 2017; Hindriks *et al.*, 2016; Laumann *et al.*, 2016; Novelli and Razi, 2021). The temporal variability in FC observed with sliding-window FC analysis has been attributed to statistical sampling error (Hindriks *et al.*, 2016; Laumann *et al.*, 2016). Notably, Laumann and colleagues developed a stationary null model based on a multivariate Gaussian distribution and matching covariance structures and spectral contents to real data. They used this null model to show that results obtained with sliding window FC analysis are similar for real and simulated data, suggesting that the stationarity of the data cannot be distinguished based on sliding window FC analysis. In relation to CAP analysis, Cifer and colleagues pointed out that long temporal autocorrelation of the (stationary) fMRI signal can explain the finding that a small fraction of time points suffices to approximate static FC (Cifer *et al.*, 2017). More recently, Novelli and Razi asked whether edge-centric FC (Faskowitz *et al.*, 2020; Zamani Esfahlani *et al.*, 2020), a recently developed point-process method similar to CAP analysis performed in the connectivity space, can capture the non-stationarity of FC in resting brain activity. They mathematically showed that the results obtained with the edge FC method, at least in its present form, can be explained by assuming a stationary Gaussian distribution with a covariance matrix matched to real data. In the present study, we extended these previous studies by showing that the results obtained with CAP analysis, another widely used analysis technique for assessing resting brain activity, were similar for real fMRI data and for simulated data generated using the

Laumann null model. This was surprising because the fact that a set of CAPs made from the time course of a single Seed ROI can be clustered into spatially heterogeneous modules appeals to our intuition that these modules represent distinct states of resting brain activity. The present results suggest that this intuition is incorrect. Spatially heterogeneous modules of CAPs were found with simulated data generated by a stationary null model. Moreover, many of the modules found in the simulated data were similar to the modules found in the real data. Thus, the presence of spatially heterogeneous modules of CAPs is insufficient to determine whether the data was generated by a mixture of distributions (*i.e.*, non-stationary activity with multiple states). These studies collectively suggest that extra care needs to be taken when interpreting the results from these analysis techniques designed to extract dynamic structures of resting brain activity.

We would like to emphasize that we are not claiming or trying to prove that resting brain activity is best represented by a stationary Gaussian distribution. In fact, careful statistical analyses suggest that resting brain activity is non-stationary (Liégeois et al., 2017). Similar to a previous study that examined sliding window FC analysis (Laumann *et al.*, 2016), our intention was to make clear what can and cannot be concluded from CAP analysis. We would also like to note that the aim of the present study was not to deny the potential clinical usefulness of CAP analysis. Several studies have applied CAP analysis to clinical data and found valuable features of brain activity that characterize patient groups (Liu *et al.*, 2018; Marshall et al., 2020; Rey et al., 2021; Yang et al., 2021). It is important to emphasize that the usefulness or clinical relevance of these features are not diminished by the present results (Liegeois R., 2021). Nevertheless, the present study sets a limit on how these features might be interpreted. For example, even when spatial patterns of modules are informative for distinguishing between CAPs from patients and healthy controls, it may be incorrect to interpret the result as evidence of distinct “meta-states” between the two groups. We believe that the correct interpretation of CAP analysis and sliding window FC analysis is indispensable for constructing realistic models of resting activity among healthy people and those affected by mental disorders (Wang and Krystal, 2014).

References

1. Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., and Calhoun, V.D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cereb Cortex* 24, 663-676. 10.1093/cercor/bhs352.
2. Cifer, L., Zarepour, M., Horovitz, S.G., Cannas, S., and Chialvo, D.R. (2017). On Why a Few Points Suffice to Describe Spatiotemporal Large-scale Brain Dynamics. arXiv, arxiv:1707.00759.
3. Faskowitz, J., Esfahlani, F.Z., Jo, Y., Sporns, O., and Betzel, R.F. (2020). Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nat Neurosci* 23, 1644-1654. 10.1038/s41593-020-00719-y.
4. Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., and Deco, G. (2016). Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *Neuroimage* 127, 242-256. 10.1016/j.neuroimage.2015.11.055.
5. Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., et al. (2013). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* 80, 360-378. 10.1016/j.neuroimage.2013.05.079.
6. Laumann, T.O., Snyder, A.Z., Mitra, A., Gordon, E.M., Gratton, C., Adeyemo, B., Gilmore, A.W., Nelson, S.M., Berg, J.J., Greene, D.J., et al. (2016). On the Stability of BOLD fMRI Correlations. *Cereb Cortex*. 10.1093/cercor/bhw265.
7. Liegeois R., Y.B., Van De Ville D. (2021). Interpreting Null Models of Resting-State Functional MRI. bioRxiv, <https://doi.org/10.1101/2021.1103.1130.437514>.
8. Liu, X., Chang, C., and Duyn, J.H. (2013). Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Front Syst Neurosci* 7, 101. 10.3389/fnsys.2013.00101.
9. Liu, X., and Duyn, J.H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc Natl Acad Sci U S A* 110, 4392-4397. 10.1073/pnas.1216856110.
10. Liu, X., Zhang, N., Chang, C., and Duyn, J.H. (2018). Co-activation patterns in resting-state fMRI signals. *Neuroimage* 180, 485-494. 10.1016/j.neuroimage.2018.01.041.
11. Liégeois, R., Laumann, T.O., Snyder, A.Z., Zhou, J., and Yeo, B.T.T. (2017). Interpreting temporal fluctuations in resting-state functional connectivity MRI.

- Neuroimage 163, 437-455. 10.1016/j.neuroimage.2017.09.012.
12. Marshall, E., Nomi, J.S., Dirks, B., Romero, C., Kupis, L., Chang, C., and Uddin, L.Q. (2020). Coactivation pattern analysis reveals altered salience network dynamics in children with autism spectrum disorder. *Netw Neurosci* 4, 1219-1234. 10.1162/netn_a_00163.
 13. Novelli, L., and Razi, A. (2021). A mathematical perspective on edge-centric functional connectivity. arXiv, arxiv:2106.10631.
 14. Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., and Petersen, S.E. (2011). Functional network organization of the human brain. *Neuron* 72, 665-678. 10.1016/j.neuron.2011.09.006.
 15. Preti, M.G., Bolton, T.A., and Van De Ville, D. (2016). The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage*. 10.1016/j.neuroimage.2016.12.061.
 16. Rey, G., Bolton, T.A.W., Gaviria, J., Piguet, C., Preti, M.G., Favre, S., Aubry, J.M., Van De Ville, D., and Vuilleumier, P. (2021). Dynamics of amygdala connectivity in bipolar disorders: a longitudinal study across mood states. *Neuropsychopharmacology* 46, 1693-1701. 10.1038/s41386-021-01038-x.
 17. Szekely, G., and Rizzo, M. (2013). Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference* 143, 1249-1272. 10.1016/j.jspi.2013.03.018.
 18. Wang, X.J., and Krystal, J.H. (2014). Computational psychiatry. *Neuron* 84, 638-654. 10.1016/j.neuron.2014.10.018.
 19. Yang, H., Zhang, H., Di, X., Wang, S., Meng, C., Tian, L., and Biswal, B. (2021). Reproducible coactivation patterns of functional brain networks reveal the aberrant dynamic state transition in schizophrenia. *Neuroimage* 237, 118193. 10.1016/j.neuroimage.2021.118193.
 20. Zamani Esfahlani, F., Jo, Y., Faskowitz, J., Byrge, L., Kennedy, D.P., Sporns, O., and Betzel, R.F. (2020). High-amplitude cofluctuations in cortical activity drive functional connectivity. *Proc Natl Acad Sci U S A* 117, 28393-28401. 10.1073/pnas.2005531117.