

## **Time-varying nodal measures with temporal community structure: a cautionary note to avoid misquantification**

Thompson WH<sup>1,2</sup>, Kastrati G<sup>2</sup>, Finc K<sup>3</sup>, Wright J<sup>1,4</sup>, Shine JM<sup>5</sup>, Poldrack RA<sup>1</sup>

1. Department of Psychology, Stanford University, USA
2. Department of Clinical Neuroscience, Karolinska Institutet, Sweden
3. Centre for Modern Interdisciplinary Technologies, Nicolaus Copernicus University in Toruń, Poland
4. Department of Philosophy, Stanford University, USA
5. Brain and Mind Centre, The University of Sydney, NSW, Australia

### **Abstract**

Brain activity can be modelled as a temporal network of interconnected regions. Recently, in network neuroscience, temporal network models have gained popularity and their network properties have been related to cognition and behaviour. Here we demonstrate that calculating nodal properties that are dependent on temporal community structure (such as participation coefficient) in time-varying contexts leads to misleading results due to fluctuations of the community structure over time. Further, we present a temporal extension to the participation coefficient measure (temporal participation coefficient) that circumnavigates this problem by considering all community partitions a node is assigned to through time. Initially, we demonstrate that when controlling for temporal communities, different nodes and time-points are identified as hubs when compared to current approaches. The proposed method allows us to track a node's integration through time while adjusting for the possible changes in community structure of the overall network

Networks of empirical data are a mathematical representation of a recorded phenomenon. Network representations are valuable because there is an array of quantifiable properties of a network that can reveal structure or dynamics. To advance understanding of the modelled system, the properties identified need to be relatable to the underlying phenomena of interest. One example of this is the identification hubs, which are highly interconnected nodes in the network that can be identified using a network-level measure (e.g. participation coefficient). After identifying hubs, these nodes can then be interpreted as having an important role of controlling information within the network. Our knowledge of the phenomenon can be advanced because the network-level measure translates to something about the empirical phenomenon.

There are recent extensions to network theory that allow for the interrogation of multilayer networks (Kivelä et al 2014). One particular type of multi-layer network is temporal network theory, in which a network representation is derived at multiple “snapshots” through time (Holme & Saramäki 2012). This approach has been used to answer questions regarding how nodes, edges, and communities in a network fluctuate in time. To generate knowledge about the underlying network, we also require that the temporal network measures can be mapped back to, or interpreted in terms of, the phenomena they are modelling. Thus, when deriving a time series of network measures (e.g. per node), the values must be comparable across network snapshots. There are many metrics available for quantifying regional topological signatures within temporal networks. Some measures are temporal extensions of static measures (e.g. TempoRank is a temporal extension of PageRank (Rocha & Masudal 2014)) whereas others apply static measures to each time-point (e.g. Bola & Sabel 2015 found changes in rich club coefficients applied to multiple time-points). When applying static measures in a temporal network context, it is important to ensure that the interpretability or clarity of the measure is not changed or distorted when used through time.

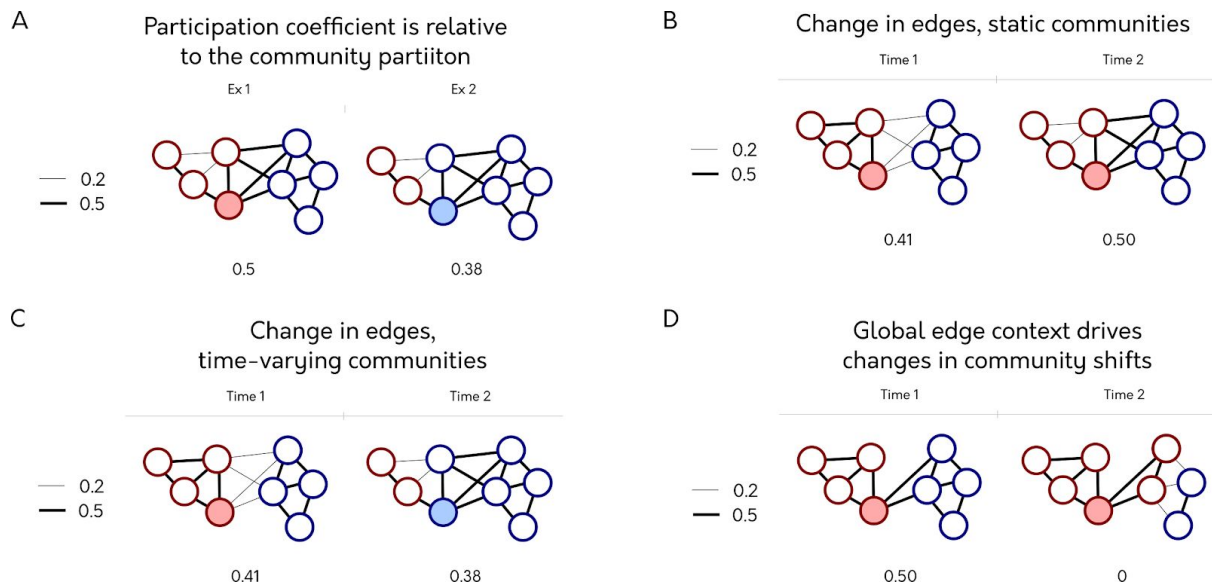
The participation coefficient (PC) is an example of a measure that has been taken from static network theory and applied to networks at multiple time-points in neuroimaging. Briefly, the PC quantifies a node based on the diversity of its connections to other nodes across a community partition (Guimera et al 2005). When the PC has been applied through time, it has often been combined with using a community partition derived for each snapshot (e.g. Betzel et al 2015, Shine et al 2016, Pedersen et al 2017, Tanimizu et al 2017, Xie et al 2018, Fukushima et al 2018; Fukushima & Sporns 2018; Shine et al 2018; Rizkallah et al 2019). Importantly, the PC for any given node is *relative* to the community partition used to calculate it (Figure 1A): if the community partition changes, then the participation coefficient may change. In the two examples in Figure 1A, the shaded node has the same edges but the communities are different, entailing that the participation coefficient changes.

A problem with the interpretation of the PC emerges when it is compared between two (or more) snapshots of a network with different community partitions. The PC is a measure that is relative to the overall community structure of the network. When community boundaries are allowed to fluctuate, as is the case in temporal networks, the participation coefficients calculated at different time points are not based upon the same community-context. This is not an issue in a static network where each node is assigned to one community, nor is it a problem if PC is calculated at multiple time-points in relation to the same static community partition. In temporal networks, communities can merge, split, disappear and reappear through time (Granell et al 2015). In the brain, it is known that

community structure changes in response to task and cognitive demands (Vatansever et al., 2015; Braun et al., 2015; Thompson et al 2019). Since community structure can change between temporal snapshots of the brain, in each temporal snapshot where the PC is calculated has a different community-context and each PC estimate is thus relative to different community-contexts. We argue that calculating the PC per time-point with a temporal community structure is not quantifying the intended property. As a result, the crucial link between the network measure and concrete interpretation breaks down and it becomes unclear what conclusions can be drawn about brain networks on the basis of such comparisons.

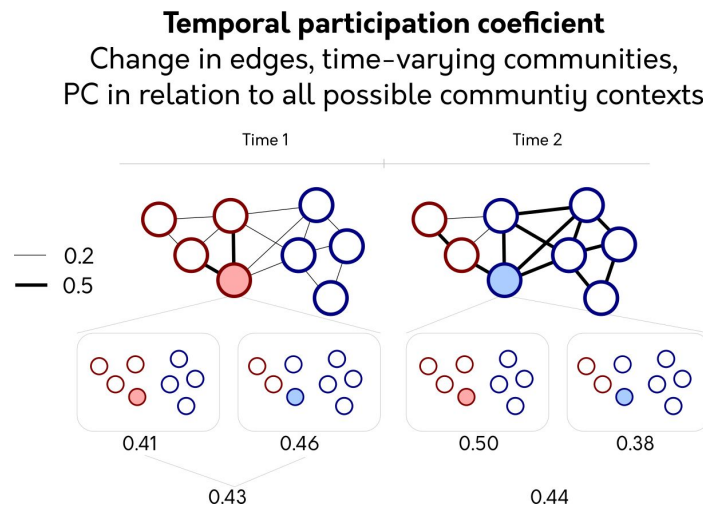
Let us illustrate the problem introduced above by calculating the participation coefficient through time on some toy network examples. Consider a time series of participation coefficients when the community partition is static (Figure 1B). For the two different snapshots in time, there is a change in the edges of the shaded node, which changes the participation coefficient of that node. Specifically, in the second snapshot, the connections of this node have become evenly distributed across nodes in all communities. We can easily relate the two PC values for the two different snapshots to each other, and it makes sense to interpret the increase in the nodes participation coefficient as an increase in the nodes interaction with communities outside of its own.

If instead the community partitions vary over time (Figure 1C), the changes in edges leads to the shaded node being classed as part of the blue community instead of the red community. The node's participation coefficient, in light of this change in community membership, is reduced. However, when viewing the two different values of participation coefficient in Figure 1C, we are unable to say if the difference is because of a switch in the edges or the community assignment. In the second snapshot, the participation coefficient has changed - it has *decreased* - but not because the node decreased its *participation* (i.e. its role in the network) but because the node changed community membership when it increased its connection strengths with the blue community relative to the first time-point. Hence, the interpretation of a temporal series of participation coefficients as reflective of a change in intra-community connections is impeded by the extent with which the modular structure of the network is changing over time. That is, none of the PC values in a time-series can be directly compared to each other nor can we directly translate the abstract measure to an external phenomena.



**Figure 1: Different ways to calculate the participation coefficient through time.** The participation coefficient for the shaded node is shown below each network/snapshot. The border of each node shows the assigned community of the node. There are two different edge weights possible. (A) Two examples of the participation coefficient illustrating how the measure is relative to the community partition. (B) When a static community template is applied across multiple temporal snapshots. (C) When a temporal community parcellation is applied to multiple snapshots. These values can not be directly compared. (D) An example showing that community partitions are driven by changes in edges that do not directly connect to the node of interest. The difference between the two time-points are within the blue community in time-point 1. This changes the community partition and will change the participation coefficient of the shaded node.

A possible objection to this argument is that the temporal communities are calculated on the edges themselves, entailing that there is an interconnection between the community-context and the edge-context of a node. This objection does not fully take into account how communities are calculated. Communities take into account the “global edge context” (i.e. all edges in a network and how they relate to each other) whereas the participation coefficient only considers the “local edge context” (i.e. all edges connected to one node). There is no necessary relationship between these two (exemplified in Figure 1D). A node’s strength can increase with no effect on the community partition. Alternatively, a node can change its community assignment with no change to its own edges.



**Figure 2: The temporal participation coefficient.** Our suggested correction for the problem in *Figure 1CD* is to calculate the participation coefficient of each time-point in relation to all possible community contexts the node can be in. The participation coefficient in relation to each possible temporal community context is shown under each community partition. This fix makes the participation coefficient of the shaded node with temporal communities commensurable.

Temporal communities contain more information than one static community partition and it is understandable there is a desire to calculate the participation coefficient in relation to fluctuating community structure. We present a possible solution to the problem outlined above: *the temporal participation coefficient*. The crux of the problem is that the participation coefficient of a node is relative to the community partition of the network that it is calculated against. If, instead, each participation coefficient estimate is calculated by considering all possible community partitions that the node is known to have been in, then the participation coefficients will be comparable across time points as each estimate is now relative to the same community context (Figure 2). This solution calculates the local edge context at a time-point with all possible community contexts, weighted by their frequency of each community. Then it considers how a node is participating relative to the possible community structure it can have. In figure 2, each TPC estimate is calculated relative to both community contexts. This entails that the shaded node at the second time-point does in fact have the larger participation. As both time-points have their different edge-contexts calculated relative to the same set of community partitions, these values can now be compared.

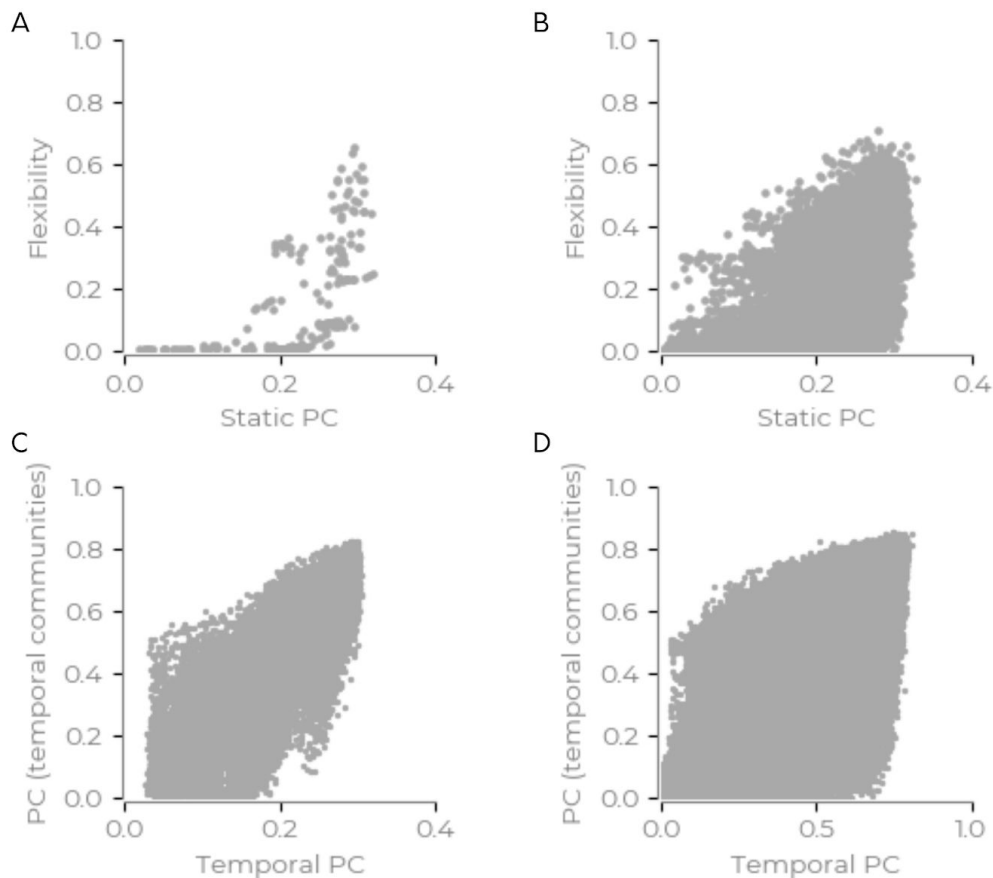
The solution we present does not fit all possible use-cases. One major limitation is that it can only be applied when the network can return to previous states. Some temporal communities may only be possible after certain events have transpired - e.g. during a contagious outbreak, patients could form communities in the hospital. Using our proposed TPC fix on such a data set would entail that post-infection communities could in principle exhibit the pre-infection topographic structure, which would be unrealistic. Thus, the proposed solution only covers networks which can theoretically return to similar states again. This is a reasonable assumption of networks such as the brain. In other cases, quantifying variations in how nodes relate to their community assignments (e.g. Bassett et al 2011) or using time-varying measures in relation to static communities (e.g. Fransson et al 2018) may be more prudent analysis alternatives. The ultimate lesson here is that network measures need to be chosen based on what is known about the system under investigation, and their

sensitivity to relevant and irrelevant properties of the underlying phenomena needs to be validated on a case by case basis.

We have presented a theoretical problem and a potential fix. However, we have yet to show that the problem has any effect on data itself or if our fix corrects it. It *could* be the case that there is no difference when applying the TPC then the PC with temporal communities.

To demonstrate that there is indeed a difference between the different ways to calculate participation, we use data from the Midnight Scan Club resting-state fMRI (Gordon et al 2017) which is available on openneuro.org (ds000224). We extracted time series from 200 functionally-defined parcels (using the parcellation of Schaefer et al 2018) from the preprocessed data available on openneuro. Time-varying connectivity estimation was done by using weighted Pearson correlations with weights based on the Euclidean distance between time-points (Thompson et al 2017; Thompson & Fransson 2018) which is a method that performs well at tracking a fluctuating covariance and also preserves topographical properties of the connectivity matrices (Thompson et al 2018). All negative edges were set to 0 prior to calculating the communities or participation. We calculated the temporal communities using the Louvain algorithm (Blondel et al 2008) with the resolution parameter of 1. Temporal consensus clustering was performed on the temporal communities (Lancichinetti et al 2012). We also calculate static functional connectivity using Pearson correlations and a static community parcellation with the same parameters as the temporal communities. We abbreviate the different participation coefficient methods as follows: static participation coefficient (static PC), participation coefficient per time-point with static communities ( $PC_S$ ), participation coefficient per time-point with temporal communities ( $PC_T$ ), temporal participation coefficient (TPC)

We begin by asking the question: how are nodes with high static PC affected by the temporal communities? This is important as, if nodes with high participation have little change in their community context, this issue we raise may be redundant. To do this, we compare the flexibility (i.e. how many times a node changes community, Bassett et al 2011) with the static participation coefficient (Figure 3AB). Here we see that nodes with high participation increase in their flexibility. If nodes with high participation always remained in the same communities, calculating the PC with temporal communities would be less problematic. Here however we see that these nodes are switching their community allegiance the most.

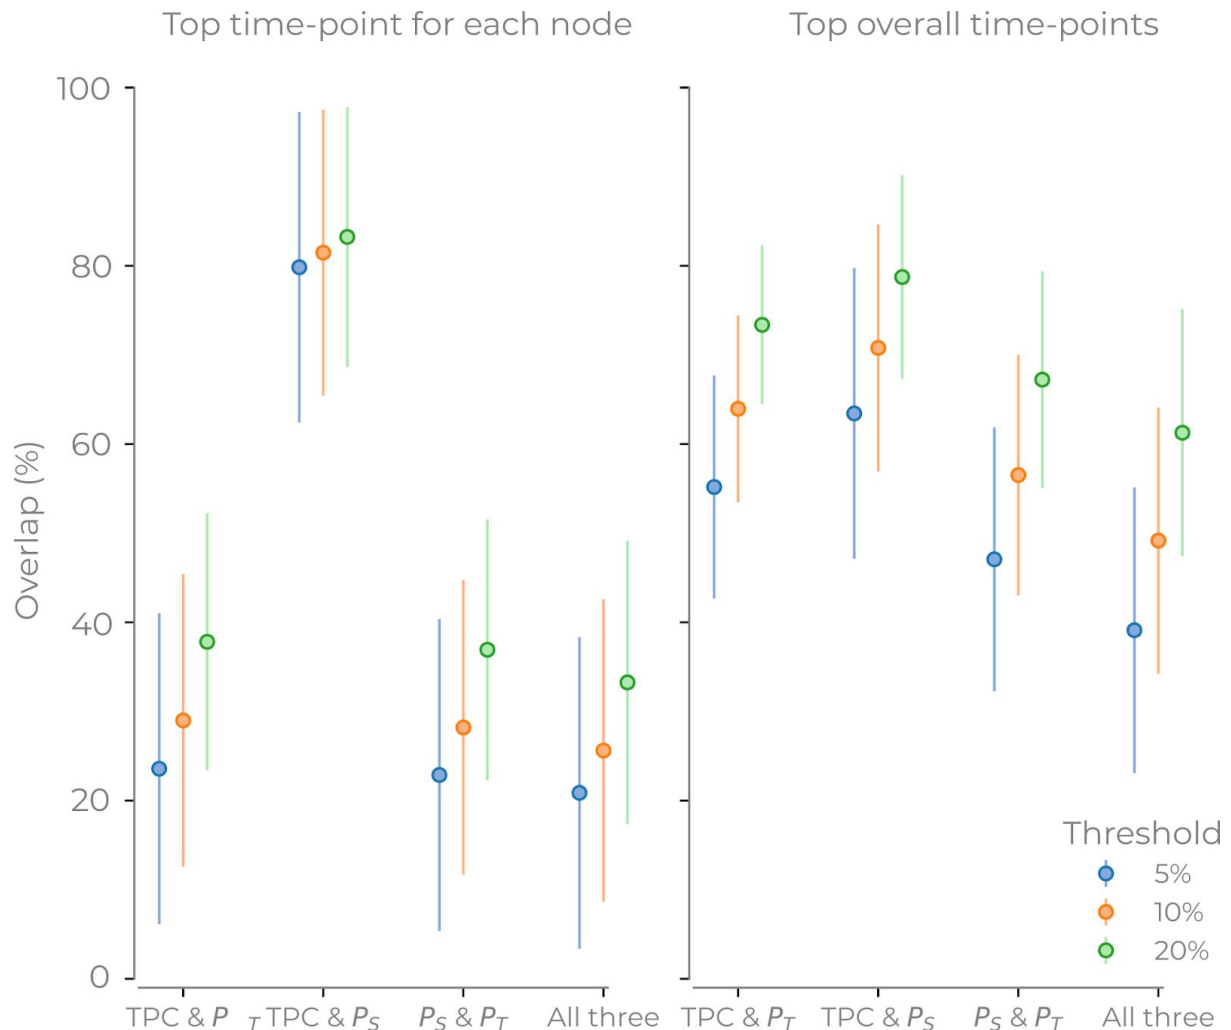


**Figure 3** (A) Static participation coefficient versus the flexibility for one subject/session. (B) Same as A but for all subjects. (C) Temporal participation coefficient versus the participation coefficient per time point with temporal communities for one subject/session. (D) Same as C, but for all subjects.

Nodes that switch temporal communities have high static PC, thus it seems concerning for participation calculated through time relative to different community partitions. Now we contrast the  $PC_T$  with the TPC to see whether they compute similar values or whether they diverge. First we begin by considering all time-points and all nodes together. A heteroskedastic relationship between the two coefficient emerges (Figure 3CD, Bartlett test for heteroskedasticity: all subjects:  $T = 153221.9$ ,  $p < 0.001$ ; example subject:  $T = 1902.3$ ,  $p < 0.001$ ). This heteroskedastic relationship will entail that, while both methods may identify the points that have the highest participation, the relationship quickly then breaks down.

To illustrate the extent that the methods diverge, we compared three different versions of participation coefficient (TPC,  $PC_S$  and  $PC_T$ ). For each method we then identified the highest 5, 10 and 20 percent of values for both for the top time-points for each node and the top overall values. We then quantified the overlap of each of the different methods (Figure 4). If we try and find when each node has its highest participation, we find that the  $PC_T$  has the most unique nodes. When pooling all nodes and time-points together, the overlap of all three methods reached over 60% with large thresholds, but was under 40% for lower thresholds. This shows that the choice of calculating the participation matters. Finally we also observe that, the TPC and the  $PC_S$  overlapped the most of the different versions (reaching 80% of nodes in some instances and always over 60% when combining the paired

and triple intersections). This is reassuring for TPC as we know the  $PC_S$  is a valid method. And the divergence that happens between the TPC and PC with static communities is due to the TPC utilizing the temporal community information.



**Figure 4: intersection of high participation coefficient from different methods.** Here we compare three different participation coefficient methods (temporal participation coefficient, participation coefficient with temporal communities and participation coefficient with static communities). Here we see the intersection of each combination of three methods. The left figure considers, for each subject, the top x% time-point for each node. The right figure, for each subject, the top x% across all nodes and time-points. Error bars show standard deviation.

We have outlined why the  $PC_T$  is problematic. Further we have proposed a solution that fixes the problem. We have also shown the importance of the fix by illustrating how the two different methods will find very different candidate hubs. The extent of the divergence between  $PC_T$ ,  $PC_S$  and TPC will depend on how much the communities fluctuate. This will depend on both the parameters, time-varying connectivity method, community detection algorithm, and the ground truth.

Here the focus has been on temporal communities and its recent application within network neuroscience. However, this can also be a more general warning for such nodal



measures that are relative to the community structure when applied in multilayer cases (e.g. the segregation-integration difference (Fransson et al 2018) could not be extended to temporal communities without the fix we propose here).

Our hope is that this article highlights the problematic nature of quantifying temporal nodal measures relative to a fluctuating temporal community partitions. We have offered one possible solution for this problem which we have shown will reveal candidate “hubs” that are detected compared with current practices.

## Acknowledgements

WHT acknowledges support from the Knut och Alice Wallenbergs Stiftelse (SE) (grant no. 2016.0473, <http://kaw.wallenberg.org>). KF was supported by the Foundation for Polish Science, Poland (START 23.2018).

## References

- Bassett, D. S., Wymbs, N. F., Porter, M. a, Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences of the United States of America*, 108(18), 7641–7646.
- Betzal, R. F., He, Y., Rumschlag, J., & Sporns, O. (2015). Functional brain modules reconfigure at multiple scales across the human lifespan. *ArXiv*, (1510.08045v1).
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10008(10), 6.
- Bola, M., & Sabel, B. A. (2015). Dynamic reorganization of brain functional networks during cognition. *NeuroImage*, 114, 398–413.
- Braun, U., Schäfer, A., Walter, H., Erk, S., Romanczuk-Seiferth, N., Haddad, L., ... Bassett, D. S. (2015). Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 112(37), 11678–11683.
- Fransson, P., Schiffler, B. C., & Thompson, W. H. (2018). Brain network segregation and integration during an epoch-related working memory fMRI experiment . *Neuroimage*, 178, 147–161.
- Fukushima, M., Betzal, R. F., He, Y., & Heuvel, M. P. Van Den. (2018). Structure – function relationships during segregated and integrated network states of human brain functional connectivity, 1091–1106.
- Makoto Fukushima, & Olaf Sporns. (2018). Comparison of fluctuations in global network topology of modeled and empirical brain functional connectivity. *PLoS Computational Biology*, 14(9), 1–27.

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Petersen, S. E., Nelson, S. M., Dosenbach, N. U. F., ... Greene, D. J. (2017). Precision Functional Mapping of Individual Human NeuroResource Precision Functional Mapping of Individual Human Brains. *Neuron*, 1–17.

Granell, C., Darst, R. K., Arenas, A., Fortunato, S., & Sergio, G. (2015). Benchmark model to assess community structure in evolving networks. *Physical Review E*, 92(1), 012805.

Guimera, R., Guimerà, R., & Nunes Amaral, L. a. (2005). Functional cartography of complex metabolic networks. *Nature*, 433(7028), 895–900.

Holme, P., & Saramäki, J. (2012). Temporal networks. *Physics Reports*, 519(3), 97–125.

Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, 2(3), 203–271.

Lancichinetti, A., & Fortunato, S. (2012). Consensus clustering in complex networks. *Scientific Reports*, 2.

Pedersen, M., Omidvarnia, A., Jackson, G. D., Zalesky, A., & Walz, J. M. (2017). Spontaneous brain network activity: Analysis of its temporal complexity. *Network Neuroscience*, 1(2), 100–115.

Rizkallah, J., Annen, J., Modolo, J., Gosseries, O., Benquet, P., Mortaheb, S., ... Laureys, S. (2019). NeuroImage : Clinical Decreased integration of EEG source-space networks in disorders of consciousness. *NeuroImage: Clinical*, 23(April), 101841.

Rocha, L. E. C., & Masuda, N. (2014). Random walk centrality for temporal networks. *New Journal of Physics*, 16.

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A., ... Yeo, B. T. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex From Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28, 3095–3114.

Shine, J. M., Bissett, P. G., Bell, P. T., Koyejo, O., Balsters, J. H., Gorgolewski, K. J., ... Poldrack, R. A. (2016). The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. *Neuron*, 92(2), 544–554.

Shine, J. M., Brink, R. L. Van Den, Hernaus, D., Nieuwenhuis, S., & Poldrack, R. A. (2018). Catecholaminergic manipulation alters dynamic network topology across cognitive states. *Network Neuroscience*, 2(3), 381–396.

Tanimizu, T., Kenney, J. W., Okano, E., Kadoma, K., Frankland, P. W., & Kida, X. (2017). Functional Connectivity of Multiple Brain Regions Required for the Consolidation of Social Recognition Memory, 37(15), 4103–4116.

Thompson, W. H., Brantefors, P., & Fransson, P. (2017). From static to temporal network theory: Applications to functional brain connectivity. *Network Neuroscience*, 1(2), 69–99.

Thompson, W. H., & Fransson, P. (2018). A common framework for the problem of deriving estimates of dynamic functional brain connectivity. *NeuroImage*, 172, 896–902.

Thompson, W. H., Richter, C., Plavén-sigraý, P., & Fransson, P. (2018). Simulations to benchmark time-varying connectivity methods for fMRI. *PLoS Computational Biology*, 14(5), e1006196.

Thompson, W. H., Wright, J., Shine, J. M., & Russell, A. (2019). The identification of temporal communities through trajectory clustering correlates with single-trial behavioural fluctuations in neuroimaging data. *bioRxiv* 617027.

Vatansever, D., Menon, D. K., Manktelow, A. E., Sahakian, B. J., & Stamatakis, E. A. (2015). *NeuroImage* Default mode network connectivity during task execution. *NeuroImage*, 122, 96–104.

Xie, H., Gonzalez-castillo, J., Handwerker, D. A., Bandettini, P. A., Calhoun, V. D., Chen, G., ... Mitra, S. (2018). Time-varying whole-brain functional network connectivity coupled to task engagement. *Network Neuroscience*, 3(1), 49–66.