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

Thompson WH[1,2], Kastrati G[2], Finc K[3], Wright J[1,4], Shine JM[5], Poldrack RA[1]

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

Introduction

Network representations are valuable in empirical research because there is an array of quantifiable properties 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 of hubs, which are highly interconnected nodes 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 directing information within the network. Our knowledge of the phenomenon can be advanced because the network-level measure translates to something about the empirical phenomenon.

Quantifying the dynamics of a network often utilizes multilayer networks (Kivelä et al 2014), especially one particular class, temporal networks, which are network representations 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 (Guimerà & Nunes Amaral 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 with temporal communities, 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 a single 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.

Methods

Data used

We used data from the Midnight Scan Club resting-state fMRI (Gordon et al 2017) which is available on openneuro.org (ds000224). The data consists of 10 subjects who were scanned for 10 different resting-state sessions. One subject was excluded due large artifacts. We extracted time series from 200 functionally-defined parcels (using the parcellation of Schaefer et al 2018) from the preprocessed data available on openneuro. See Gordon et al 2017 for more information regarding preprocessing.

Time varying connectivity estimation

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).

Three methods to quantify the participation coefficient

The participation coefficient is defined by Guimerà et al 2015 as:

$$P_i = 1 - \sum_{s}^{N_m} \left(\frac{k_{is}}{k_i} \right)^2 \quad \textit{Static participation coefficient (PC).}$$

where i is a node index; N_M is the number of communities. k_{is} is the within-community degree and k_i is the overall degree of node i.

When calculating the participation coefficient through time, with static communities, the equation is:

$$P_{it} = 1 - \sum_{s}^{N_m} \left(\frac{k_{its}}{k_{it}} \right)^2 \quad \textit{Participation coefficient through time with static communities } (PC_s)$$

Here we see that temporal subscripts have been given to the participation and the degree of the nodes. Note, there is only one static parcellation that every time-point is references to.

The participation coefficient in relation to temporal communities is:

$$P_{it} = 1 - \sum_{s}^{N_{mt}} \left(\frac{k_{its}}{k_{it}} \right)^2 \quad \textit{Participation coefficient through time with temporal communities } (PC_T)$$

Above, we are now summing over the communities N_{mt} which is the number of communities that are found at time-point t. This uses a community partition per time point.

The temporal participation coefficient that we introduce is:

$$P_{it} = 1 - \frac{1}{T}\sum_{u}^{T}\sum_{s}^{N_{mu}}\left(\frac{k_{its}}{k_{it}}\right)^2 \quad \textit{Temporal participation coefficient (TPC)}$$

Here we have added an additional term, where all temporal community partitions are considered for each time-point. See the results section for the motivation behind the TPC. 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).

Quantifying community structure, module degree z-score and flexibility

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 at t that had the largest smallest Jaccard distance with the community at t-1 were assigned the same community index (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.

Aside from the various participation coefficient estimates, we also calculate flexibility (i.e. how many times a node changes community divided by the total number of possible changes, Bassett et al 2011) and within-module degree z-score (MDZ; ie. z-score of node's degree within its community, Guimerà & Nunes Amaral 2005). The MDZ was calculated in two ways, one using static communities (MDZ_S) and temporal communities (MDZ_T) and are used in conjunction with their PC counterparts.

Results

Different underlying reasons for participation coefficient changes through time

Let us illustrate the problem introduced 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 partition varies 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. In the second snapshot, the participation coefficient has changed - it has *decreased* - not because the node decreased its *participation* (i.e. its role in the network) but rather 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 community 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.

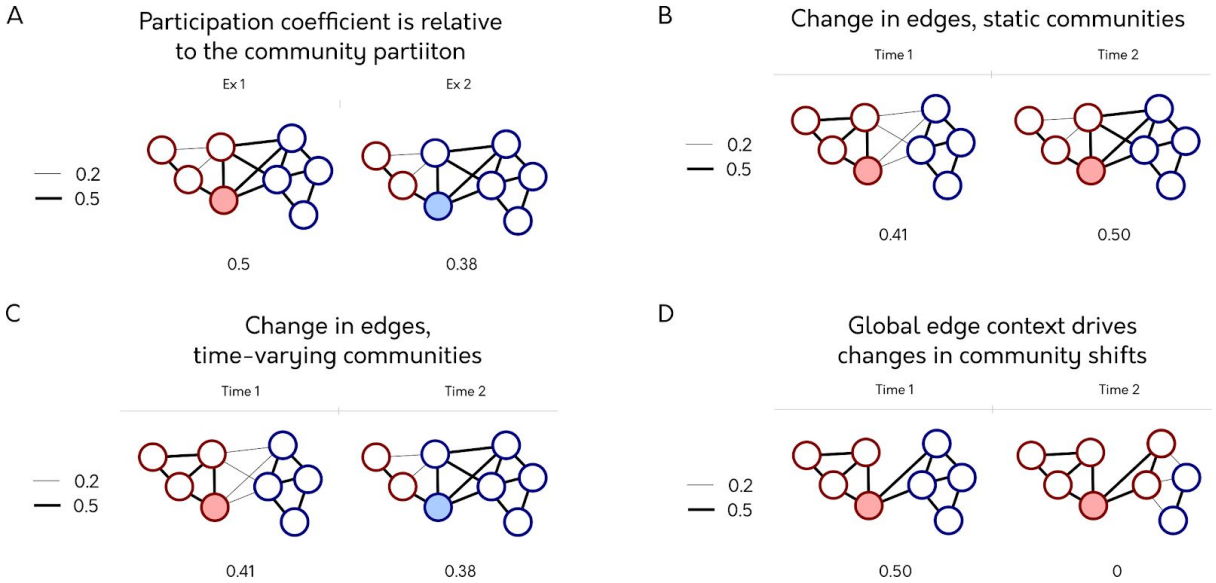

Figure 1: Different ways to calculate the participation coefficient through time. The participation coefficient for the shaded node is shown below each network/snapshot. The border of each node shows the assigned community of the node. There are two different edge weights possible. (A) Two examples of 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.

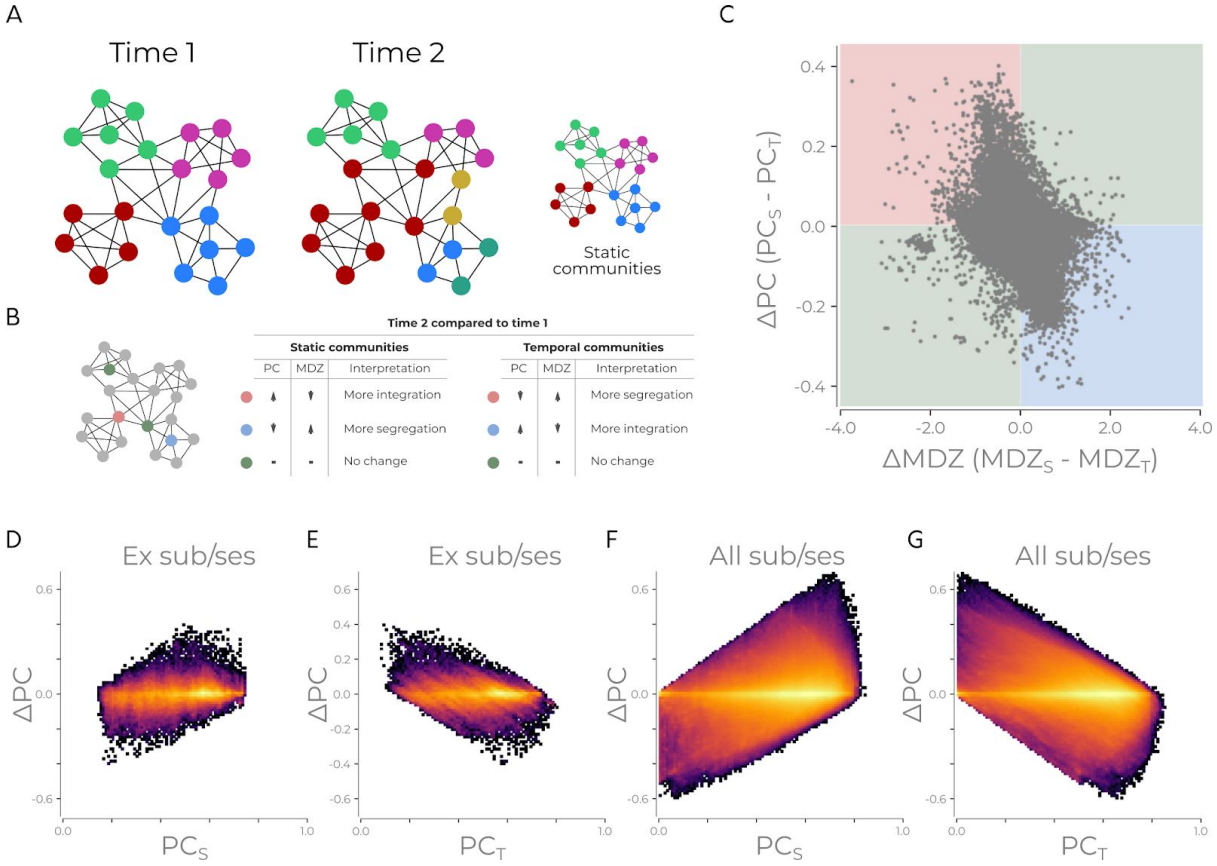

Figure 2. How the different PC methods leads to different interpretations of the network. (A) An example network consisting of four communities over two nodes. (B) Three different types of nodes have been highlighted from A and the accompanying table suggests how these nodes at time-point 2 will be quantified relative to time-point 1 with static or temporal communities for the participation coefficient (PC) and module degree z-score (MDZ). (C) The difference in PC and MDZ for nodes in the static and temporal for all nodes/time-points for one subject/run in the MSC dataset. The red, blue and green quadrant corresponds to the interpretations found in B. (D) Density plot showing the difference in the PC versus PC_S for the example session/subject in C. (E) Same as D but for the difference in PC versus PC_T. (F) Same as D but for all sessions/subject. (G) Same as E but for all sessions/subjects. All density plots have logarithmic color scale.

Misleading network-level interpretations of PC with temporal communities

We have demonstrated that community-context affects the nodal measure when quantified at multiple snapshots. The measures are applied correctly according to their mathematical definitions, which raises the question of why the issue we raise is of any importance. Here we show how the PC_T can cause misinterpretations regarding the property participation is generally trying to quantify.

Consider the toy network demonstrated in Figure 2A. Here we have two different time-points where the community-context changes if considering temporal communities. We have selected a pair of nodes from this network and identified how PC_S and PC_T will assign in relation to increased integration, segregation or no change in relation to previous time-point (Figure 2B). Note how PC_T assigns high PC_T, and thus interpreted as having higher integration, to the node marked in blue in Figure 2B. Here the community that was originally in the first time-point (Time 1) has split into multiple communities in the second time-point (Time 2). It is hard to consider how a decrease in community cohesion should be interpreted as increased participation — but that is what will happen

with PC_T. Likewise, the temporal community that the red node belongs to has extended in the second time-point (Figure 2B, Time 2). This merging of nodes means that the red node is sharing similar information with the more nodes than usual, yet the PC_T method will ascribe a low score entailing that there is less integration. Thus, high PC_T can lead (and has led) to the incorrect interpretation that more integration is occurring in the network, when the opposite is happening. Note, there is nothing mathematically wrong with each PC_T estimate, the problem lies when contrasting PC_T estimates with each other and then interpreting differences between time points in terms of time-varying network structure.

Do situations identified in Figure 2B occur within the contexts that PC_T has been applied (time-varying connectivity with fMRI brain networks)? To test this, we calculated the difference between PC_T and PC_S and between MDZ_T and MDZ_S on an example subject's resting state fMRI data (Figure 2C). In Figure 2C, the coloured quadrants represent situations similar to Figure 2B's node examples. Here the red quadrant indicates that they would be deemed as having "more integration with static communities, more segregation with temporal communities" and vice versa for the blue quadrant. A large portion of the nodes (Example session/subject: 75.59%; All subjects: 66.08%) end up in the red and blue quadrants, entailing that they have alternative interpretations about what is occurring in the network when using the PC_T and PC_S methods.

For completeness, figure 2D-G shows when the absolute difference of PC is high, than one of the methods with have high PC and the other low PC. The difference between PC methods is related to the PC values of both methods for both an example session/subject and all subjects in the dataset. This is expected due to the PC values being bound between 0 and 1, but it is possible that all the differences of e.g. -0.4 occur when the $PC_S = 0$ and $PC_T = 0.4$, which would make our characterization of the quadrants in Figure 2C misleading. However the differences scale with the PC magnitudes.

PC in relation to all temporal communities

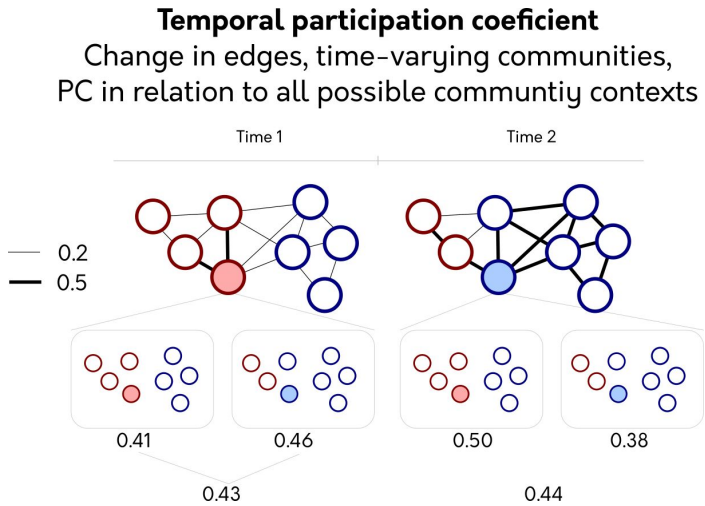

Figure 3: The temporal participation coefficient. Our suggested correction for the problem in *Figure 1CD* is to calculate the participation coefficient of each time-point in relation to all possible community contexts the node can be in. The participation coefficient in relation to each possible temporal community context is shown under each community partition. This fix makes the participation coefficient of the shaded node with temporal communities commensurable.

Given the substantial evidence for temporal changes in community structure, there is an understandable 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 3). 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 3, 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 motivation behind TPC is to have all time points compared to the same community context (like PC_S) but to retain the temporal information inherent in community fluctuations (used in PC_T). Contrasting time-points' participation from different community contexts, as shown above, can lead to misleading interpretations. To rectify this and still utilize the time-varying community structure, we calculate each time-points degree in relation to all possible community structures. If this seems counter intuitive, the logic of using a community vector that does not fully reflect the underlying data is common (e.g. in PC_S where a static community partition is applied across varying time points). The only difference here is that TPC applies multiple community partitions. To successfully do this, two assumptions have to be met. First, a single node's activity does not fully drive the community context (demonstrated in Figure 1D). This entails that the edge context of a node can be considered in relation to multiple instantiable community contexts, not just the one that happens to be the best fit at the specific time-point when considering all edges. Second, the community structure can recur which is a system like the brain, is a reasonable (see below for when this assumption breaks down).

TPC entails that a node's activity should be quantified in relation to all communities it could potentially be in. If one large edge (e.g. 0.9) is responsible for merging two communities and the communities split when the edge is smaller (edge weight 0.1), everything else being equal, TPC will give high participation to the high edge weight because it relates that edge to the other possible situations the network could be in. Since all time-points are quantified in relation to the same community information (all partitions), it is mathematically impossible for any time-point that decreases all of it's edge weights to get higher participation — this guarantee is not possible for the PC_T. This demonstrates that the TPC can utilize all community contexts but avoids the problematic applications shown for PC_T in the previous section.

The solution we present does not fit all possible use-cases. One limitation is that it can only be applied when the network can return to previous states (the recurrence assumption). 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 would influence pre-infection participation, which would be unrealistic. Furthermore, care would also be needed for any of the participation methods if a network bifurcates its community structure between entirely different states that have little or no topographic overlap as combining the network contexts may be hard to interpret. Thus, the proposed solution only covers networks which can theoretically return to similar states again and restraint on the possible temporal community structure exists (e.g. anatomy). This is a reasonable assumption of networks such as the brain. Ine 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. PC_S) 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.

Nodes with high static PC change communities the most.

We have presented a theoretical problem, illustrated how it can lead to differences in interpretation and presented a potential fix. Yet we have to show that the extent of 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 TPC versus PC_T.

We begin by asking the question: how are nodes with high static PC affected by the temporal communities? If nodes with high participation have little change in their community context, the problem we raise may be redundant. To clarify this, we compared the flexibility with the static participation coefficient (Figure 4AB). Here we see that nodes with high participation also increase their flexibility. If nodes with high participation always remained in the same communities, calculating the PC with temporal communities would be less problematic. Nodes that switch temporal communities have high static PC, thus it seems concerning for participation calculated through time relative to different community partitions. This does not prove which participation method should be preferred, but it shows that the community contexts are affecting nodes with high participation and will help us understand any differences between participation methods.

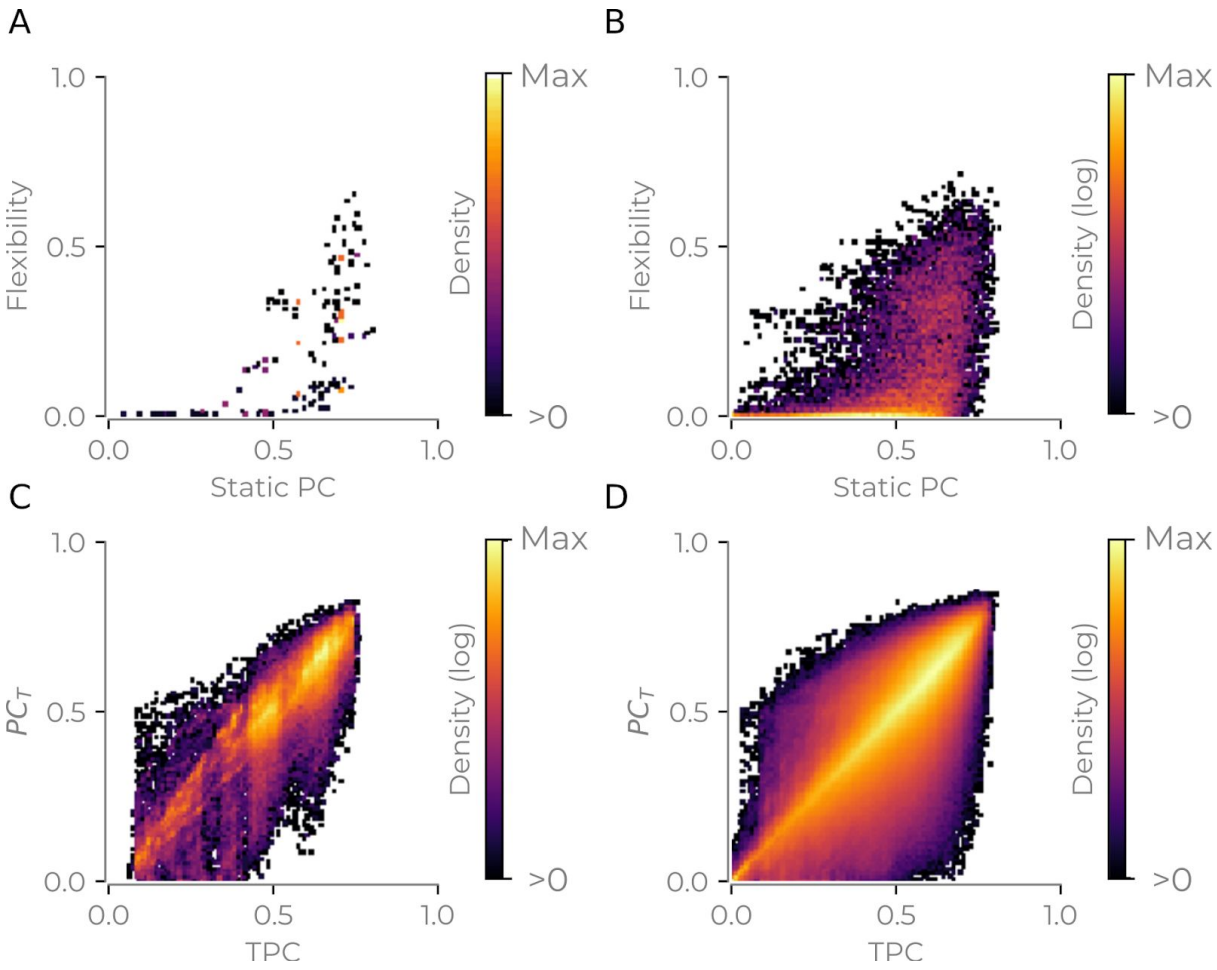

Figure 4 (A) Static participation coefficient versus the flexibility for one subject/session. (B) Same as A but for all subjects. (C) Temporal participation coefficient versus the participation coefficient per time point with temporal communities for one subject/session. (D) Same as C, but for all subjects. All colorbars show density on a 100x100 grid. Panels B-D show logarithmic range between colours.

Methodological divergence of the different methods

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 4CD, 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. This shows that the methods converge.

To quantify the extent that the methods diverge, we considered two different questions: (1) do the time-series of participation coefficients correlate with each other? (2) If selecting the top x% of time-points to be marked as candidate temporal hubs, do the selections intersect? Here, we compared all three of different versions of participation coefficients expressed through time (TPC, PC_S and PC_T).

With regards to how much the time series of participation correlated with each other, the time series of PC_S and PC_T did not correlate highly (Spearman rank (ρ): median: 0.30, σ: 0.21, min: -0.44 , max: 1.0, Figure 5A), especially compared to PC_S and TPC (Spearman rank: median: 0.96, σ: 0.19, min: -0.80 , max: 1.0, Figure 5B). TPC also did not correlate highly with PC_T (median: 0.32, σ: 0.21, min: -0.37 , max: 1.0, Figure 5C). In sum, these correlations of the time series show that TPC and PC_S correspond the most. However, taking the mean PC over time for the various methods, the correlation between the methods is high for all method combinations (PC_S & PC_T: ρ=0.91; PC_S and TPC: 0.91; PC_T & TPC: 1.0). Taken together, this shows that nodes do not dramatically change in identifying which new nodes which have high participation, the methods diverge in when they have high participation. Our challenge to PC_T is that between time-point comparisons are problematic. However, averaging over time-points is not as problematic.

However, correlation of the time series of different participation methods does not mean that the same nodes or time-points will get selected as candidate hubs as the correlation is only quantifying their covariance. 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 (Figure 5D) and the top overall values (Figure 5E). 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 PC method 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.

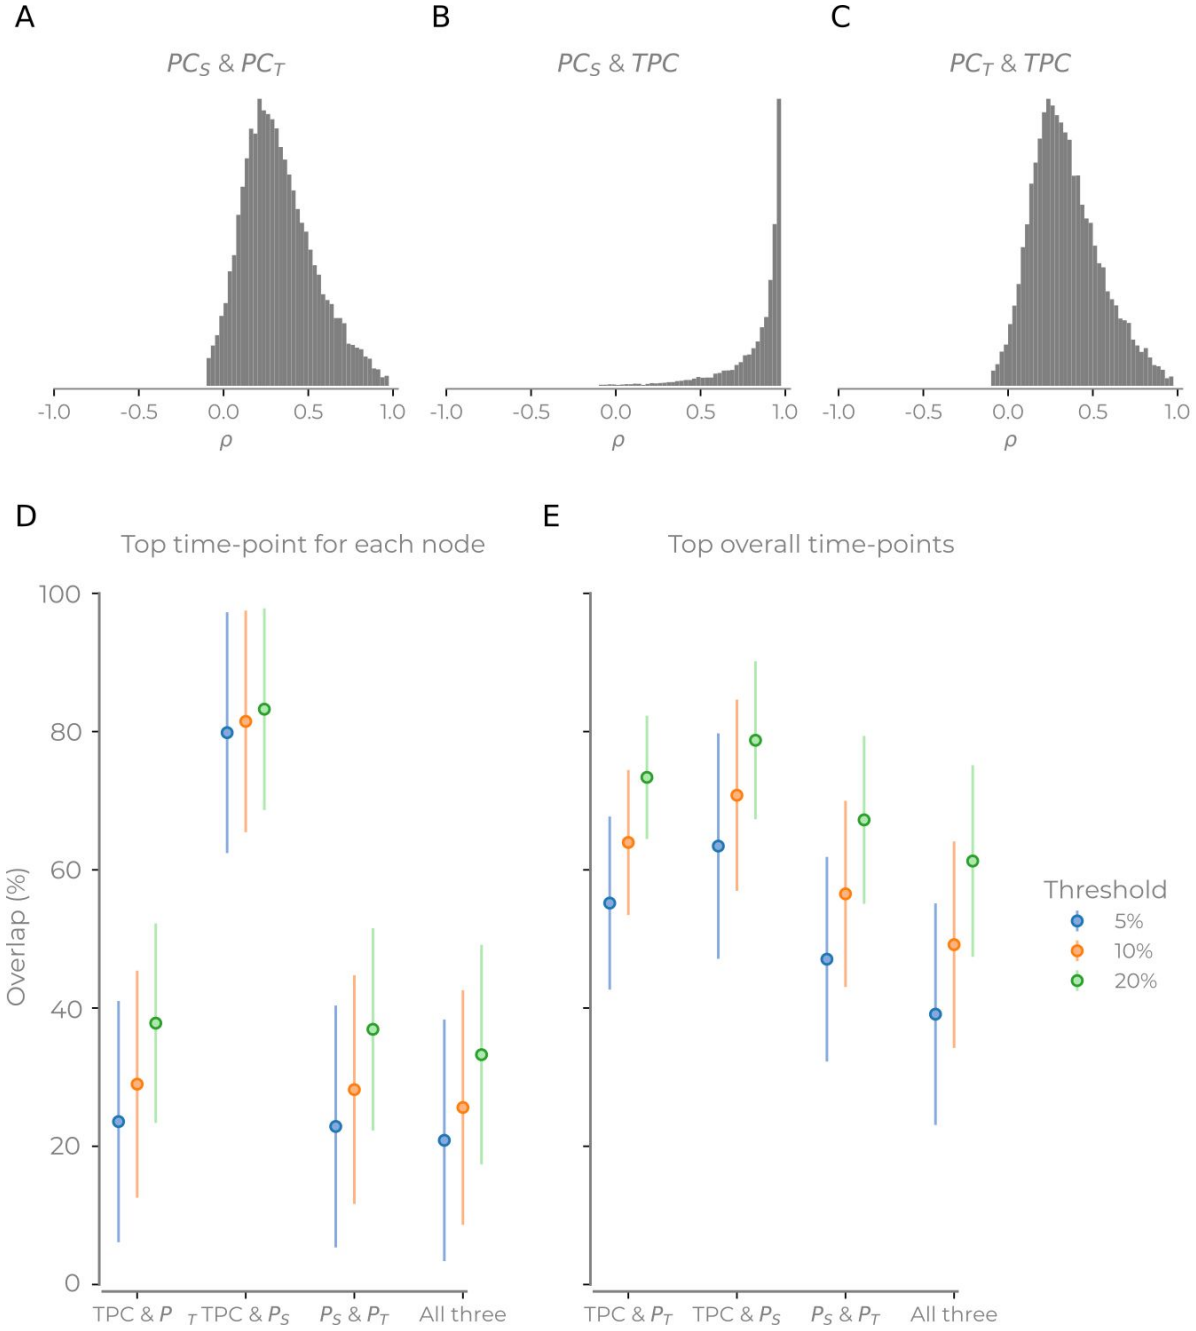

Figure 5: Comparison of the three different PC methods. (A-C) are the histograms of the correlations of the time series for each method (histogram showing all nodes, sessions and subjects). (D-E) Intersection of high participation coefficient from different methods. Here we see the intersection of each combination of three methods. (D) For each subject, the top x% time-point for each node. (E) For each subject, the top x% across all nodes and time-points. Error bars show standard deviation.

Discussion

We have outlined why the PC_T can lead to misleading interpretations when contrasting across different temporal snapshots. Further we have proposed the temporal participation coefficient which allows for participation estimates to be contrasted across time-points without impeding the interpretation. Finally we have also shown that these methods diverge in how much nodal time series correlate and which nodes will be considered 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 but the methods do diverge. We feel that, for any quantification of fluctuations of participation through time will require PC_S or TPC. We are advocating that PC_T can lead to misleading interpretations and should be avoided. Finally, we are not challenging PC_T in all use cases, it is mathematically sound but the interpretation of the values between time-points that is problematic. But there are possible theoretical cases where it can be applied (e.g. a network where communities only merge, then the interpretation is not impeded).

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 that utilizes temporal community information that does not suffer from similar issues regarding interpretation.

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). There are no conflicts of interest to declare.

References

Bassett, D. S., Wymbs, N. F., Porter, M., 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.

Betzel, 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., Betzel, 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.

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-sigray, 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.