



Review

Modularity maximization as a flexible and generic framework for brain network exploratory analysis

Farnaz Zamani Esfahlani^{a,1}, Youngheun Jo^{a,1}, Maria Grazia Puxeddu^{a,b,c,1}, Haily Merritt^{d,e,1}, Jacob C. Tanner^{d,e,1}, Sarah Greenwell^a, Riya Patel^a, Joshua Faskowitz^{a,f}, Richard F. Betzel^{a,e,f,g,*}

^a Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, United States

^b Department of Computer, Control and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Rome 00185, Italy

^c IRCCS Fondazione Santa Lucia, Rome 00179, Italy

^d Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47405, United States

^e Cognitive Science Program, Indiana University, Bloomington, IN 47405, United States

^f Program in Neuroscience, Indiana University, Bloomington, IN 47405, United States

^g Network Science Institute, Indiana University, Bloomington, IN 47405, United States

A B S T R A C T

The modular structure of brain networks supports specialized information processing, complex dynamics, and cost-efficient spatial embedding. Inter-individual variation in modular structure has been linked to differences in performance, disease, and development. There exist many data-driven methods for detecting and comparing modular structure, the most popular of which is modularity maximization. Although modularity maximization is a general framework that can be modified and reparameterized to address domain-specific research questions, its application to neuroscientific datasets has, thus far, been narrow. Here, we highlight several strategies in which the "out-of-the-box" version of modularity maximization can be extended to address questions specific to neuroscience. First, we present approaches for detecting "space-independent" modules and for applying modularity maximization to signed matrices. Next, we show that the modularity maximization framework is well-suited for detecting task- and condition-specific modules. Finally, we highlight the role of multi-layer models in detecting and tracking modules across time, tasks, subjects, and modalities. In summary, modularity maximization is a flexible and general framework that can be adapted to detect modular structure resulting from a wide range of hypotheses. This article highlights multiple frontiers for future research and applications.

1. Introduction

The human brain is fundamentally a network and can be described at multiple spatiotemporal scales in terms of nodes – neural elements, e.g. cells, populations, areas – and edges, e.g. synapses, projections, fiber tracts, and correlated activity (Park and Friston, 2013; Sporns et al., 2005). Over the past two decades, network analyses have uncovered key attributes of brain networks that are believed to support distinct modes of brain function. These include small-world architecture to support segregated and specialized information processing and efficient integration of information over long distances (Bassett and Bullmore, 2006; Sporns and Zwi, 2004), influential and central hubs that form integrative rich clubs (Hagmann et al., 2008; Van Den Heuvel and Sporns, 2011), and cost-efficient spatial embeddings (Bassett et al., 2010).

Additionally, brain networks can be decomposed into subgraphs – referred to as modules or communities in the language of network science (Meunier et al., 2009; Sporns and Betzel, 2016). Modules are typically defined as cohesive subgraphs, where neural elements within the same modules tend to form strong connections to one another while those in different modules are weakly connected or not connected at all. This

type of network organization helps support specialized processing, has been used to delineate the brain's functional systems (Gordon et al., 2016; Power et al., 2011), and is thought to support human cognition (Medaglia et al., 2015). Studying the brain's modular structure confers obvious advantages. It helps with dimension reduction (Betzel et al., 2017), facilitates discovery of underlying structure, and allows for classification of nodes' roles (Guimera and Amaral, 2005).

How does one actually go about studying a network's community structure? On the one hand, network neuroscientists can borrow from one of any number of previously discovered community partitions to assign nodes (parcels) to communities or cognitive systems. For instance, Power et al., 2011 describes a division of the brain into thirteen systems based on their network analysis. Other studies have reported similar divisions (Gordon et al., 2016; Schaefer et al., 2018), allowing users to obtain a validated set of assignments for brains in their study.

However, there are many cases where we might expect previously defined communities to be unsuitable, for instance, in case-control studies, where we anticipate ahead of time that a clinical population might exhibit communities dissimilar from those of a neurotypical individual (Alexander-Bloch et al., 2010). More generally, if we want to capture

* Corresponding author.

E-mail address: rbetzel@indiana.edu (R.F. Betzel).

¹ Contributed equally.

individual differences in community structure, then imposing a shared partition onto a cohort may not be especially useful.

In these cases, we can use data-driven methods to algorithmically discover a network's community structure (Fortunato, 2010; Fortunato and Hric, 2016; Newman, 2012). In most cases, these algorithms aim to partition a network's nodes into clusters that maximize some objective function. For instance, Infomap seeks a maximally compressible description of a random walker moving over a network's edges (Rosvall and Bergstrom, 2008; Sanchez-Rodriguez et al., 2021). Stochastic block-models maximize the likelihood that a probabilistic model generated the observed network (Peixoto, 2014). Among this class of data-driven algorithms, however, modularity maximization is among the best known and most widely applied within network neuroscience (Newman and Girvan, 2004). The aim of modularity maximization, in general, is to partition nodes into communities so that the observed density of connections within communities is maximally greater than what we would expect had the network been generated under some null connectivity model. This intuition is captured by the modularity objective function, Q .

Although this definition of communities is general and can accommodate many distinct definitions of the “null connectivity model”, its application in the context of network neuroscience has been narrow. In most applied studies, the null connectivity model generates random networks with a fixed degree and strength distribution. This “enshrined” model, which was promoted in the original studies of modularity, helps correct for communities that are driven by a network's degree distribution, i.e. groups of nodes that due to the fact that they make many connections are therefore more likely to be connected to one another. This model, however, has a number of issues – e.g. it allows for self-connections, which are generally disallowed in the construction of empirical brain networks, and corresponds to a network with prohibitively expensive wiring cost.

The aim of this prospective article is not to provide a definitive roadmap for performing community detection or even how to use modularity maximization in practical contexts. For this, there are many useful reviews (Fortunato and Hric, 2016; Betzel). Rather, the aim of this article is to highlight modifications to a standard community detection algorithm that open up opportunities to address research questions that are central to network neuroscience and would be difficult to address otherwise.

In this study, we will examine several ways in which the modularity function could be extended to be more biologically plausible and to accommodate specific research questions and hypotheses. This work is divided into three sections. In the first section, we cover alternative null connectivity models and discuss their relative strengths and weaknesses, commenting on their suitability for network neuroscience research. In the second section, we show how modularity maximization can be used to discover clusters of brain regions that are dissimilar between conditions or correlated with a continuous variable. In the final section, we will introduce the multi-layer formulation for modularity maximization and highlight strategies for extending this framework to study multi-subject, multi-modal, and time-varying network datasets. Alongside this paper, we provide documented example scripts so that researchers can apply these techniques to their own data.

2. Modularity maximization

2.1. Network construction

The human brain can be divided into functionally, cytoarchitecturally, and connectionally distinct regions or areas (Brodmann, 1909; Desikan et al., 2006; Schaefer et al., 2018). These regions and their pairwise connections can be modeled as the nodes and edges of a graph or network and represented as its connectivity matrix, $A \in \mathbb{R}^{N \times N}$ whose element A_{ij} encodes the existence and weight of the connection between nodes i and j (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010).

This representational framework can be used to model the network architecture of both functional and structural brain networks, although network inference and interpretation differs between modalities.

2.1.1. Structural connectivity

Anatomical or structural connectivity (SC) refers to physical pathways that link neural elements to one another, including synaptic contacts (White et al., 1986) and axonal projections (Markov et al., 2014; Oh et al., 2014). In the case of large-scale brain networks, these pathways represent white-matter fiber tracts which can be inferred non-invasively from diffusion-weighted MRI data using tractography algorithms (Hagmann et al., 2008; Sporns et al., 2005). SC changes over long timescales, e.g. with development Baum et al., 2017, 2020 and aging Betzel et al., 2014, and is relatively invariant over shorter scales (the duration of a typical MRI scan). Although its precise density of connections varies across spatial scales and reconstruction procedure (Ercsey-Ravasz et al., 2013; Horvát et al., 2016; Sarwar et al., 2019), SC is generally sparse, meaning that many pairs of regions are not directly connected. This sparsity, along with the overall configuration of connections, imposes constraints on interregional communication patterns, effectively shaping the spatiotemporal evolution of brain activity (Adachi et al., 2012; Avena-Koenigsberger et al., 2018; Honey et al., 2007).

2.1.2. Functional connectivity

Functional connectivity (FC), on the other hand, is a mathematical construct that measures the statistical relationship between activity recorded at two different points in the brain (Friston, 1994; Horwitz, 2003). Although this definition includes virtually any bivariate (and possibly multivariate) statistical measure, FC is most commonly measured as the zero-lag correlation of fMRI BOLD activity between two regions (and usually extended to all pairs of regions). Whereas SC constrains interregional communication, FC reflects the patterns of correlated activity that emerge as a consequence of those constraints. Unlike SC, which changes across long timescales, FC evolves rapidly and fluctuates over short periods of time (Hutchison et al., 2013; Lurie et al., 2020).

2.2. Community structure of human brain networks

Although SC and FC reflect distinct modes of interregional coupling, they both exhibit modular structure (Sporns and Betzel, 2016). Reported structural modules tend to be compact and contain spatially contiguous regions (Hagmann et al., 2008). The spatial contiguity of structural modules likely reflects evolutionary pressure to reduce the cost of wiring by forming short-range and hence less-costly connections (Stiso and Bassett, 2018). Functional modules, because there is no explicit cost associated with the formation of long-distance correlations, are organized into brain-wide and spatially-distributed clusters (Power et al., 2011; Yeo et al., 2011). Even in the absence of explicit task instructions, functional modules recapitulate well-known patterns of task-evoked activity and previously-delineated functional systems (Smith et al., 2009). These observations give rise to the hypothesis that modularity is a key component of functional specialization (Crossley et al., 2013).

2.3. The modularity or “Q” heuristic

In general, a network's modular structure is unknown ahead of time. Even in the case of SC and FC, whose modular structure has been characterized many times over, we may be interested in how the modular structure of individuals or groups of individuals deviates from composite, population-averaged modules. Assigning a network's nodes or edges to modules by hand or based on visual inspection would be tedious and unfeasible for large networks or even for small networks in large, multi-subject cohorts. Accordingly, the network elements are usually assigned to modules algorithmically through a process known as community detection (Fortunato, 2010).




Fig. 1. Modularity maximization. Modularity maximization works by comparing an observed network (a) with what would be expected were one to simulate a null connectivity model an infinite number of times (b). Specifically, the observed (A) and expected (P) weights of connections are compared to one another via an element-wise subtraction, yielding a modularity matrix B (c). Modularity maximization aims to assign network nodes to communities so that the internal density of connections is maximally greater than chance. The modularity function, Q , is used to assess the quality of modular partitions. In panel d we show examples of a high-quality partition (top; within-community connections tend to be stronger than expected) and a low-quality partition (bottom).

There exists a multitude of algorithms and heuristics for detecting a network's modules (Lancichinetti and Fortunato, 2009). Some, like stochastic blockmodels, use statistical inference to group nodes into clusters with similar connectivity profiles. Most methods force nodes to belong to only one module (although others allow for overlap (Ahn et al., 2010; Evans and Lambiotte, 2009; Faskowitz et al., 2020; Palla et al., 2005)). Some methods detect modules based on how network properties shape the evolution of dynamical processes on the network (Pons and Latapy, 2006; Rosvall and Bergstrom, 2008). Others focus on the network's static topology alone.

In this last category is modularity maximization (Newman and Girvan, 2004). As noted in the introduction, modularity maximization operates on an eminently simple principle: compare what we actually see with what we might expect. In the context of networks, this means comparing an observed connectivity matrix, A (Fig. 1a), with another matrix of identical dimensions, $P \in \mathbb{R}^{N \times N}$, whose element P_{ij} encodes the weight of the connection between nodes i and j that we would expect under some null connectivity model (Fig. 1b). The simplest comparison of these matrices is an element-wise subtraction:

$$B = A - P. \quad (1)$$

The resulting matrix, B , has a special name – the *modularity matrix* (Fig. 1c). The element B_{ij} encodes whether the actual connection between i and j is stronger ($B_{ij} > 0$) or weaker ($B_{ij} < 0$) than we would expect. Modularity maximization uses the modularity matrix to evaluate the goodness or quality of a modular partition, i.e. a division of a network's nodes into non-overlapping modules. A partition's quality is rated by the modularity function, Q , which is calculated as:

$$Q = \sum_{ij} B_{ij} \delta(\sigma_i, \sigma_j). \quad (2)$$

Here, $\sigma_i \in \{1, \dots, K\}$ indicates to which of the K modules that node i is assigned and $\delta(\sigma_i, \sigma_j)$ is the Kronecker delta function, whose output is 1

when $\sigma_i = \sigma_j$ and 0 otherwise. Effectively, then, this double summation is over edges (node pairs) that fall within modules. The larger the value of Q , the higher quality the partition (Fig. 1d). That is, Q tends to be large when communities are more internally dense than expected.

Rather than using Q to simply evaluate the quality of a partition, Q can be optimized outright, using a procedure known as modularity maximization. Like many clustering problems, discovering the partition that corresponds to the global maximum Q is computationally intractable, and so many algorithms have been proposed to approximate its value. Among the most popular is the so-called “Louvain algorithm” – a multi-stage, greedy algorithm that, in benchmark testing, is one of the fastest and most accurate (Blondel et al., 2008; Jutla, Jeub, Mucha) (although new versions challenge this (Traag et al., 2019)). The Louvain algorithm is initialized with every node assigned to its own community. In random order, nodes are merged into larger communities if the merger yields an increase in Q . When no single-node moves can increase Q , nodes are aggregated into “meta-nodes” composed of all nodes with the same community label. The previous steps are then repeated until a complete cycle of single-node moves and aggregation yields no improvement.

2.4. Best practices

Modularity maximization is a useful framework for discovering communities in a network. Using it effectively, however, requires that a user navigate several issues including the near-degeneracy of solutions as well as a potential resolution limit. Here, we offer some practical guidance for addressing these issues.

2.4.1. Near-degeneracy of the modularity landscape

Obtaining the optimal partition of a network into communities is impossible in all but the most trivial instances. Stochastic methods like the Louvain algorithm generate estimates of this optimal partition, but

it is understood that no single estimate should be treated as privileged. More generally, and this is especially true as the size of a network grows (Good et al., 2010), there will be many dissimilar partitions of roughly the same quality. The reason for this is simple. Imagine we could obtain the globally optimal partition, i.e. the one corresponding to the largest possible Q . If we were to randomly select a node and move it to a new community, the partition would no longer be optimal, but it would likely still be of high quality. If we did this for every node, for every possible move, and then every pair of nodes, and every triplet, etc., we would obtain an ensemble of near-optimal solutions. In practice, methods like the Louvain algorithm sample from these types of partitions – high quality, but likely sub-optimal. Importantly, there is heterogeneity among these solutions. So how do we solve this?

One possibility is to embrace this heterogeneity and focus on properties of this ensemble rather than those of any single partition (Peixoto, 2012). Another possibility is to condense information from the ensemble of partitions into a co-assignment matrix, whose elements indicate the fraction of partitions in which pairs of nodes were assigned to the same community (Kenett et al., 2020). This results in graded and fuzzy community assignments.

Oftentimes, however, it is useful to obtain a single representative partition to summarize the ensemble of near-optimal solutions. Again, there are multiple strategies for doing so. One possibility is to calculate the pairwise similarity of partitions to one another and select the one that is most similar, on average, to the others (Doron et al., 2012). The most common strategy, however, is to estimate a consensus partition. In general, this involves iteratively reclustering the coassignment matrix (Lancichinetti and Fortunato, 2012), sometimes after applying a threshold to remove elements from node pairs that are infrequently coassigned to the same community. A potentially more useful and parameter-free method is to estimate, along with the empirical coassignment matrix, the expected coassignment matrix, whose elements denote the likelihood that two nodes would be assigned to the same community given a random set of partitions where the number and size of communities is equal to that of the estimated partitions. Note that the elements of the expected co-assignment matrix can be estimated using a permutation model (Betzel, 2020). With the empirical and expected coassignment matrices, we can set up a modularity problem and use the modularity maximization heuristic to discover the consensus clusters.

2.4.2. Resolution limits and multi-scale extensions

If we constructed a network of maximally connected cliques and linked cliques to one another by single edges, we would expect modularity maximization to discover the cliques as communities. In practice, however, for certain sized networks and null connectivity models, the presence of even a single link between communities may appear so “unexpected” that it is advantageous from the perspective of modularity maximization to group cliques together, defying the expectation that each clique correspond to a distinct module. This tendency to merge small, well-defined communities into larger clusters is known as the resolution limit (Fortunato and Barthelemy, 2007). The implication is that, under certain conditions, modularity maximization may be blind to small communities.

To resolve this issue, many authors have incorporated a structural resolution parameter into the modularity equation. In its standard implementation, this parameter scales the relative contribution of the null connectivity model (Arenas et al., 2008; Reichardt and Bornholdt, 2006). That is, $B = A - \gamma P$, where γ is the resolution parameter. In effect, varying the value of γ changes the scale (roughly, the size of communities) that modularity detects. When γ is small, many elements of A are greater than γP and, when grouped into large communities, have the net effect of increasing Q . On the other hand, when γ is larger, a much smaller and exclusive set of connections in A exceed γP , yields much smaller communities. Thus, γ acts as a way to zoom in and out, detecting finer and coarser structure, and helping to circumvent the resolution limit. In some instances, a “sweep” across γ values, in which

modular maximization is performed many times while systematically varying the resolution parameter, can be used to tune the results of modularity maximization to a specific number of communities or to obtain a multiresolution (or hierarchical) organization of the modular structure (Akiki and Abdallah, 2019; Harris et al., 2019; Jeub et al., 2018).

2.5. Modifying modularity

In practice, modularity maximization is used to address the question: “here is a network, what are its modules?” However, even this seemingly straightforward question requires that the user make a series of decisions (Betzel, 2020). Many of these decisions are fundamental and cannot be avoided even if modularity maximization were being used to address a different set of questions. These include how to deal with the degeneracy of the modularity landscape (the number of nearly-optimal partitions grows exponentially with the size of the network) (Good et al., 2010) and how to address issues of scale and the so-called resolution limit (under some conditions, modularity maximization can fail to detect small communities, even if they are unambiguously defined) (Fortunato and Barthelemy, 2007). Other decisions are structural and concern how the modularity maximization problem is formulated. For instance, the user needs to define the expected weights of connections – the variable P . This decision is non-trivial and will impact the character of the detected communities (Expert et al., 2011).

The choice of null model is one decision that a user needs to make and is an example of how the modularity maximization problem can be modified. There are, of course, other ways to alter the way the problem is set up (e.g. by including a resolution parameter that scales the relative contribution of P (Reichardt and Bornholdt, 2006; Schaub et al., 2012; Lambiotte, Delvenne, Barahona)). In fact, any decision that changes the elements of the modularity matrix, B , has the capacity to impact the detected communities. In the following sections, we explore three strategies for modifying the modularity matrix and thereby changing the modularity maximization problem. We show that the modifications help address questions that are at the core of network neuroscience.

3. Single-layer modularity maximization

Most commonly, modularity maximization is used to discover the modular structure of a single network. In neuroscience, this network might represent the structural or functional connectivity from a single subject or it may be a composite (average) of connectivity data from many individuals (Betzel et al., 2019). This is arguably the simplest modularity maximization problem. The aim is to discover modules by maximizing Q , which is a function of two variables: A , the connectivity matrix from the network of interest, and P , the expected weight of connections. All the user needs to do is define P and run the optimization heuristic.

3.1. Null models for brain networks

What is the appropriate choice for P ? Consider, for a moment, two brain regions labeled i and j . Maybe your network of interest is structural, and the voxels along the white-matter tract between i and j have some average fractional anisotropy, which represents the observed weight of the connection. What is the expected FA of that connection? What is the null (connectivity-based) model that gave rise to that expected weight? And further, what network characteristics does the null model consider? Does it take into account the spatial embedding of the connection, and does it avoid rendering self-loops? Perhaps the network of interest is derived from functional imaging data, and the edges represent co-activity patterns between regions. What null model would account for the statistical relationships of this network?

3.1.1. Configuration models

These are questions that the user must confront when performing modularity maximization. However, in most applications the entire en-




Fig. 2. Alternative null models for modularity maximization. Modularity maximization requires that a user specify both the observed and expected weights of connections. (a) An example observed network (in this case an empirical, group-averaged structural connectivity network). (b) Common null models do not preserve spatial information. (c) There exist many null models that preserve spatial information (and in some cases other information about the network, e.g. nodes' degrees). Note that, in general, models in which spatial information is preserved tend to be more similar to the observed network. This is evident even upon visual inspection. (d) Consensus community labels obtained using a null model in which spatial information is absent. (e) Consensus community labels obtained using a null model that preserves spatial information along with nodes' degrees. Panels f and g show module co-assignment matrices for both, revealing that they both generate “internally dense and externally sparse” assortative communities. (h) The element-wise difference between co-assignment probability matrices, however, reveals that the strength-preserving model tends to group nodes in the same hemisphere into the same community, while the spatial model allows nodes in different hemispheres to belong to the same community.

terprise of selecting a null model and defining P is sidestepped. Rather than thinking about what might be the appropriate model, the user lets the software implementation decide for them (if the software even gives the user the option of deciding), usually defining P based on a “configuration model”. The most common model of this class is one that generates null networks whose nodes have the same degree and strength as in the original network, but where connections are otherwise formed at random (Fig. 2b). Under this model, the expected weight of a connection is given by:

$$P_{ij} = \sum_{ij} \frac{k_i k_j}{2m} \quad (3)$$

where $k_i = \sum_j A_{ij}$ is the degree of node i and $2m = \sum_j k_j$ is the total number of edges in the network (or node strength and total weight if the network's edges are weighted).

Intuitively, by selecting this model the user is comparing their observed network against the set of all possible networks with precisely

the same degree/strength sequence. This is an important point; the configuration model preserves a set of features of the observed network (same number/weight of edges, same degree sequence), while randomizing others. By preserving these features, specifically nodes' degrees, the configuration model also guards against the possibility of two nodes being connected to one another and assigned to the same module simply because they form more connections.

While the configuration model is a good general null, is it an appropriate null generator for brain networks? If not, what network features should be preserved in neuroscience-specific null model? How do we efficiently generate randomized networks with those features?

3.1.2. Incorporating spatial relationships

One of the key organizational features of brain networks is the dependence of connections and their weights on nodes' spatial embeddings. In general, the further apart two brain regions are from one another, the less likely they are to be connected and the weaker the weight of

the connection. This is especially true in the case of SC. In other words, connections are not formed independently, but according to the geometry in which the brain is embedded. These dependencies are critical for viability; longer and thicker connections are more costly, requiring more material to form and more energy to maintain and use for signaling. Moreover, they must be compact enough to fit within the skull. Therefore, over time, brains have come to adopt wiring diagrams that favor the formation of short (and less costly) connections. A brain that violates this soft constraint and forms many long connections quickly becomes too materially and metabolically expensive, exceeding the limits of viability.

This is the chief problem with configuration models. In practice, these models assume that connections are formed randomly and independently; the only constraint is that the random networks' nodes have the same degree/strength as the original network. This process, however, incorporates no information about spatial dependencies; networks generated using the configuration model contain many times more long distance connections and their total cost (measured as the summed length of their connections) far exceeds what is observed in any empirical brain network (Gollo et al., 2018). The configuration model, because it fails to preserve this essential feature of brain networks, can be viewed as far too liberal and almost a straw man when it comes to comparisons with empirical data. Note that a second, but possibly less flagrant issue with the configuration model is that it allows for self-connections, i.e. $P_{ii} = \frac{k_i k_i}{2m} \neq 0$. These connections are usually zeroed out in empirical SC and FC data, which means that the null model allows for connections that are explicitly forbidden in the observed data.

Fortunately, generating randomized networks that preserve a network's wiring cost is a simple process. In fact some of these algorithms can simultaneously preserve cost and node degree (although it is much more difficult to preserve node strength as well (Fig. 2c).

One strategy for preserving spatial relationships is to preserve the binary structure of the observed network but to randomly assign weights to the edges (Roberts et al., 2016). This model naturally preserves each node's degree as well as the total cost of wiring. In addition, the model also sets each edge's expected weight to an identical value equal to the mean weight of all connections in the original matrix.

A second possibility is to construct a "minimally-wired" model. In this model, one calculates the mean interregional distance between every pair of nodes and places connections between the K shortest distances, where K is the number of edges in the original network. Weights can then be introduced through rank-preserving procedures. Note that this model does not preserve nodes' degrees.

A third possibility is to parametrically model the effect of distance. In these models, we define a function that monotonically varies as a function of distance (or possibly other variables), e.g. $P_{ij} = \exp(-\gamma D_{ij})$ or $P_{ij} = D_{ij}^{-\eta}$. Here, P_{ij} is the probability that nodes i and j are connected and D_{ij} is the distance between those nodes. As long as $\gamma > 0$ and $\eta > 0$, the probability of forming long connections will be smaller than short connections. Using these models, connections can be populated using two methods – first, one can calculate the probability for every pair of nodes and then, for every $\{i, j\}$, flip a biased coin to determine whether or not a connection should be added. The variables γ and η can be selected so that the expected number of connections generated by the model matches that of the observed network $K = \langle K \rangle$. A second possibility is to change the equalities to relative probabilities. i.e. $P_{ij} = \exp(-\gamma) \rightarrow P_{ij} \propto \exp(-\gamma)$, and iteratively add edges until K edges in total have been placed. This procedure ensures that networks generated by this model have exactly K connections.

So far, each of these models preserve select features of a spatially constrained brain network but never more than one. In general, preserving multiple features is a more challenging task and, for certain sets of features, the space of possible random networks generated by the null model may be small or difficult to access. However, there are several strategies for doing so. One possibility is to use a "space-aware" configuration model. The generic configuration model can be approx-

imated using a rewiring or "edge-swapping" procedure, in which two edges between four distinct nodes are swapped, so that if the initial set of edges were formed between nodes $\{i, j\}$ and $\{u, v\}$, the new set could be formed between $\{i, u\}$ and $\{j, v\}$ or $\{i, v\}$ and $\{j, u\}$. This procedure can be made "space-aware" by forcing the new edges to have similar lengths as the two original. The result is a network whose wiring cost approximates that of the original network and whose degree sequence is identical. This procedure also approximately preserves the relationship between edges' weights and lengths.

The incorporation of spatial information into null models helps to preserve neuroanatomical realism by generating networks whose wiring cost is approximately equal to that of the observed SC matrix. This is accomplished by approximating the observed distribution of connection lengths. However, these models do not preserve other features of the observed network. For instance, a key feature of brain network data is the propensity for homotopic regions to be connected to one another (Deco et al., 2014). Preserving other features is possible but requires that additional constraints be imposed on null models, leaving open opportunities for future work.

The methods described above work well for sparse networks where edges are generated independently. However, for networks whose edges are interrelated to one another, e.g. correlation networks and functional connectivity, this method may be inappropriate (Zalesky et al., 2012). Introducing spatial relationships into null models of functional connectivity is, therefore, more challenging and, in some ways, not as necessary. While the effect of space on SC is well-documented (Honey et al., 2009), functional connections reflect statistical relationships between recorded activity and not physical tracts. Consequently, there is no explicit cost associated with forming a functional connection, whereas there is a cost associated with a structural connection. Indeed, while the correlation of functional connectivity weights and distance is strong for very short connections, beyond ≈ 40 mm, functional connections can take on a wide range of values, including some that are exceptionally strong, e.g. homotopic partners.

Nonetheless, it may be useful to incorporate spatial information into the null model for FC. One possibility for doing so is based on quadratic norms, as in Bellec et al., 2006. Briefly, this procedure entails fitting a parametric model of distance-dependence to data. The model, importantly, preserves transitive properties of correlations, as does other similar models that fit spatial properties of correlation matrices (Burt et al., 2020).

Another strategy is to incorporating spatial relationships into correlation matrices at the level of time series. For instance, one recent paper proposed a method for generating a distance-dependent correlation matrices through a reweighting of uncorrelated regional time series (Esfahlani et al., 2020). For every region, its time series was defined as a weighted sum of all other regions' time series, where the weights were inversely proportionally to the Euclidean distance between regional centroids. Accordingly, nearby regions contributed more than distant regions. As a result, the elements of the resulting correlation matrix decayed monotonically as a function of distance and, because the weights from the summation were parameterized, the rate of decay could be modulated to match the sharp decay in FC over the first 40 mm seen in empirical FC data.

Spatially-informed null models incorporate an element of realism not observed in the configuration model. At the very least, they show that the choice of null model will, in general, impact the character of the detected communities and their subsequent interpretation. This was demonstrated in two recent papers (Betzel et al., 2017; Esfahlani et al., 2020) in which configuration models were explicitly compared with models that encode spatial information. In Betzel et al., 2017, the authors fit a spatial generative model to SC data and used the model to define the expected weights of connections. The resulting modularity matrix emphasized connections whose existence would not be anticipated had the network been generated according to cost-reduction principles alone. The communities detected using the model reflected this,

exhibiting broader spatial extent than those detected with the configuration model, including instances of spatially disjoint communities, a feature that is usually never observed using the standard configuration model.

More recently, Esfahlani et al., 2020 applied a similar model to FC data from the Human Connectome Project (Van Essen et al., 2013). The authors demonstrated that, after correcting for spatial relationships, nodes in sensorimotor systems exhibited increased participation coefficient, suggesting that spatial relationships may obscure the centrality of some regions. Then, using the spatial model as a null model for modularity maximization, the authors detect communities and show that they, in general, disagree with previously delineated communities (Schaefer et al., 2018).

3.1.3. Null models for signed matrices

Another way that the modularity equation can be (and often is) reimagined, is to make it compatible with signed matrices. This is a necessity when dealing with functional connectivity data where weights are estimated using Pearson correlation. The result is a matrix in which every node is connected to every node by a connection of some magnitude, oftentimes of mixed valence (positive and negative connections).

Naively, one might ask why the standard configuration model is inappropriate. Recall that the configuration model calculates expected connection weights based on the (weighted) degrees of nodes. In the case of correlation matrices, every node has identical degree and, because at least some fraction of connections are negative, a simple summation of their weights can return a misleading value for a node's weighted degree (because positive and negative connections can offset one another).

Some of the earliest strategies for dealing with signed matrices is to construct a modularity function that deals separately with the positive and negative elements. That is, given a connectivity matrix, A , to create two distinct matrices, A^+ and A^- , comprised of the positive and negative connections only. Then, to use the configuration model to estimate the expected positive and negative weights of connections as $P_{ij}^{\pm} = \frac{k_i^{\pm} k_j^{\pm}}{2m^{\pm}}$ and, from these elements, two modularity matrices: $B^{\pm} = A^{\pm} - P^{\pm}$. Intuitively, communities should correspond to groups of nodes that are cohesive (mutually positively correlated). To realize this intuition, we can define the following modularity (Gómez et al., 2009):

$$Q^{\text{signed}} = \sum_{ij} [B_{ij}^{+} - B_{ij}^{-}] \delta(\sigma_i \sigma_j). \quad (4)$$

Optimizing this modularity returns groups of nodes that are more correlated with one another than expected by chance and less internally anticorrelated than expected by chance.

Note that in this example, the contributions to Q^{signed} by positive and negative connections are weighted equally. For networks containing approximately the same number of positive/negative connections this may be acceptable. However, for networks with imbalances between positive/negative connection weights, this is more of a problem. Accordingly, recent variations on this general model have aimed to balance these contributions by weighting the positive and negative terms differently (Rubinov and Sporns, 2011).

A second, possibly more important limitation of this approach, is that the configuration model assumes independence between edges. For structural connectivity where there is no mathematical constraint on connection weights, this is not an issue (although as noted previously, the brain's intrinsic geometry imposes a different set of soft constraints). However, for correlation networks like functional connectivity, connection weights are dependent on one another. That is, given the connection between nodes $\{i, j\}$ and $\{j, k\}$, we can place bounds on the weight of the connection between nodes $\{i, k\}$. Accordingly, our null model should also satisfy this property.

There exist several strategies for doing so. One possibility is construct a generative model that operates at the level of time series and create synthetic data, whose correlation matrix can be computed and will satisfy the transitive relationships observed in any correlation matrix. This

model can then be sampled many times, the matrices averaged, and an expected value for each connection estimated. This, however, is computationally expensive, and requires the user to construct the appropriate null model, which may not always be easy to do in the time domain. Another possibility is to use approaches like the Hirschberger-Qi-Steuer (HQS) model (Hirschberger et al., 2007; Zalesky et al., 2012) to generate admissible covariance matrices whose elements follow a specific distribution.

More recently, another study pointed out that the configuration model has a complicated interpretation in the context of correlation networks (appropriateness aside) (Bazzi et al., 2016; MacMahon, Garlaschelli). As an alternative, the authors suggested that a uniform null model may be appropriate for correlation matrices. The uniform null model is one in which it is assumed that every element in the network is mutually correlated with the same magnitude. That is, $P = \mathbf{1} \cdot \gamma$, where $\mathbf{1} \in \mathbb{R}^{N \times N}$ is a matrix whose elements are all 1 and γ is the magnitude of mutual correlation. This model performed well in benchmark analyses and has the added advantage of not suffering from the resolution limit (Traag et al., 2011). Notably, this model also has a convenient interpretation. Communities discovered at a given γ value correspond to groups of nodes whose mean connectivity to one another exceeds a value of γ on average. This truism facilitates a clearer interpretation of the relationship between communities and the resolution parameter.

In this section, we highlighted several strategies for modifying the modularity equation to make it compatible with correlation matrices. Like strategies for incorporating spatial information, these strategies yield more conservative, but ultimately more realistic, null models that respect the properties of correlation matrices and functional connectivity. Failing to preserve low-level features of a correlation matrix amounts to a straw man argument and increases the risk of detecting spurious or misleading community structure. An important and subtle point, however, concerns the relationships between null models and so-called "generative models" (Betzel et al., 2016; Betzel and Bassett, 2017). In an effort to make null models ever more realistic, one might preserve so many features that networks generated by the null model very closely approximate the observed network data. If this is the case, then the null model has become a generative model of the network – a simple set of rules or parameters that can be used to fully and parsimoniously explain a network's organization (Akarca et al., 2021; Shinn et al., 2021; Vértes et al., 2012). This is not the aim of the null model in modularity maximization. Rather, the goal is to hold a set of low-level features constant and identify those connections that are unanticipated given networks that preserve only these features. While the models described would likely never be considered true generative models, it is critical that, in future explorations, the spirit of null models be retained and that they should avoid becoming fully fledged generative models of the network.

3.2. Condition and group differences

In the previous section, we highlighted potential null models for modularity maximization and noted that each can be viewed as a distinct null hypothesis. From this perspective, modularity maximization can be viewed analogously to the null- or point-hypothesis testing; comparing an observation with some intuition of chance and measuring a summary statistic as output (it is, of course, not truly the same). Whereas the output in traditional null hypothesis testing is a p -value, in the case of modularity maximization, the output is a mapping of nodes to communities and a Q score. Here, the similarity between modularity maximization and hypothesis testing is by analogy only. However, is there a way to use modularity maximization to perform some version of null hypothesis testing, for instance in case-control studies to compare a group of patients with controls or to contrast FC acquired during one task with FC acquired during another? In this section, we explore this question.




Fig. 3. Schematic illustrating how modularity maximization could be used to detect group/condition differences. Modularity maximization works by comparing the connectivity that is observed with what is expected, but is flexible to how those matrices are defined. (a) Here, we construct a modularity matrix by treating the observed network to be FC estimated during movie-watching while the expected network is FC estimated at rest. (b) This modularity matrix can be optimized directly to discover communities of nodes that are more strongly correlated with one another during movie-watching than at rest. (c) By flipping the sign of the matrix, we can discover groups of nodes that are more strongly correlated at rest than during movie-watching. Note that, in principal, this procedure must be accompanied by a statistical model to reduce the possibility of false-positives, as modularity maximization will always find communities in data, even if none exist.

3.2.1. Comparing FC across tasks

One way that modularity maximization could be extended to test distinct null hypotheses is in the comparison of FC acquired during different task states with FC at rest. For example, how does the brain's community structure change when performing a working memory task compared to, say, rest? The conventional response to this question is that the changes are subtle but systematic (Cole et al., 2014). Resting state modular organization is an optimized state, where metabolic demands are minimized, allowing the brain to rapidly reconfigure in service of any forthcoming task-related processing goal (Wig, 2017). Thus, FC reconfigures only slightly from rest-to-task and their connection weights are highly correlated. What if we reframed the question slightly, such that we consider “rest” to be our null or baseline condition – the P in the modularity expression – and working memory FC to be our observed network – the A matrix (Fig. 3)? In this case, we could construct a modularity matrix:

$$B = FC_{task} - FC_{rest}. \quad (5)$$

We could then use this matrix to recover groups of nodes whose mutual connections strengthen during the working memory task compared to rest. Interestingly, the sign of this matrix could also be flipped, so that $B' = FC_{rest} - FC_{task}$. In this case, the detected modules now reflect groups of nodes whose connections are stronger at rest than during the working memory task.

3.2.2. Comparing connectivity between groups

If we are willing to consider the hypothetical scenario described above, then there are other interesting ways that modularity maximization can be modified to accommodate other hypotheses and research questions. For instance, we could use it to discover system-level differences in case-control studies, e.g. where some clinical or treatment population gets compared against a control group (Alexander-Bloch et al., 2010). In this case, we can regard the connectivity of the typical control subject as our expected connectivity pattern; the configuration of connection weights that represent a neuro-typical subject. The clinical group, on the other hand, represents what we actually observe. Using these two components, we can construct a modularity matrix:

$$B = FC_{clinical} - FC_{control}. \quad (6)$$

Optimizing modularity based on this matrix will result in communities whose mutual connections are stronger in the clinical condition than they are in a neuro-typical individual. We note, of course, that we could test the opposite hypotheses (what are the clusters of nodes that are stronger in controls than in the clinical group?) by simply flipping the sign of the modularity matrix or equivalently by subtracting the clinical FC from that of the controls.

This extension of modularity maximization to test for group-level differences in meso-scale structure has yet to be explored. We note, however, that this approach can be viewed as analogous to other methods, most notably the *network-based statistic* (NBS), which uses an element-wise statistical map (e.g. correlations or group differences of edge weights), applies a statistical threshold, and extracts the largest supra-threshold component (Zalesky et al., 2010). With modularity maximization, we instead uncover subsets of nodes with changed connectivity patterns. Importantly, modularity maximization makes no *a priori* assumption about the size of these subsets. NBS, on the other hand, presumes that the group-level effect gets manifested in the largest connected component.

Using modularity maximization to detect clusters of nodes whose connectivity changes with condition requires some additional caveats. For instance, it requires developing methods for controlling false positives. In other words, modularity maximization will *always* return a partition of a network into communities. If there exists element-wise differences in $FC_{clinical}$ and $FC_{control}$ or FC_{rest} and FC_{task} , even in the absence of true clusters, then modularity maximization could spuriously detect modules, i.e. find signal in noise. Controlling for this possibility requires exquisitely accurate estimates of the connectivity matrices (likely not possible) or a carefully constructed null model. In the case of null models, one possibility would be to estimate the group-level matrices using random permutation of group assignments. That is, one could randomly assign some of the clinical data as controls and *vice versa*, optimize the resulting modularity, and compare the detected communities and their quality with those generated using the true condition labels. This very issue, in fact, arises in applications of NBS, and is addressed using the same permutation strategy. As we will also illustrate in the next sections, another solution to mitigate false positives and facilitate the tracking of changes across conditions is adopting a multi-layer framework.

4. Multi-layer modularity maximization

The previous two sections focused on strategies for modifying the modularity quality function by directly changing the expected weights of connections – the P matrix in the equation $B = A - P$. These changes transform modularity maximization from a simple community detection algorithm, into a framework for investigating some of the central questions in neuroscience, including the comparison of groups and experimental conditions.

These analyses, however, also expose a feature (or perhaps limitation) of modularity maximization that is shared by most other community detection algorithms. Namely, these algorithms are designed to estimate communities for one network at a time. If a network's structure varies across time (e.g. time-varying FC [Lurie et al., 2020](#)), if we want to consider multiple subjects in a cohort ([Betzel et al., 2016](#)), the standard modularity function is poorly suited for doing so. To detect modules in datasets made up of multiple connectivity matrices requires that we extend modularity maximization further.

When we consider network datasets comprised of multiple connectivity matrices, there are two general strategies by which they can be analyzed. On one hand, each of the networks can be analyzed independently from all of the others. In the case of a multi-subject cohort, this involves estimating modules for each of the subjects and then comparing them at the end. However, this strategy quickly leads to problems aligning communities from one subject to another. That is, when can we say that communities detected using data from subjects A and B are realizations of the same community? With the exception of trivial cases (where the communities contain identical sets of nodes), this problem becomes challenging and is not easily nor unambiguously resolved, often requiring that the user introduce an additional heuristic.

The second strategy is to simply average the multiple connectivity matrices, forming a single composite matrix. While this strategy allows us to use well-developed community detection methods (or even the variants described in the previous two sections), it also results in a loss of information; by averaging the weights of connections across many matrices, we emphasize their average features, but lose all information about their variability.

Clearly, both strategies present challenges and tradeoffs; the first is capable of capturing inter-subject variation, but has challenges aligning communities from one individual to another. The second approach is more parsimonious and circumvents the alignment problem, at the expense of information about how connectivity patterns (and communities) differ or vary from one matrix to the next. Which should we choose then?

A third strategy for approaching this problem is to take advantage of multi-layer networks ([Kivela et al., 2014](#)) ([Fig. 4a](#)). Whereas an individual connectivity matrix is a two-dimensional object that encodes pairwise connections among a fixed set of elements, multi-layer networks can be conceptualized as a “stack” of these network “slices” or “layers”, forming a three-dimensional multi-layer object. Multi-layer networks are used throughout network science ([De Domenico, 2017](#); [Vaiana and Muldoon, 2018](#)) and recently have begun to be used widely within network neuroscience, where they are typically used to study dynamic FC ([Bassett et al., 2011](#)), where network organization is modeled through a series of time-varying connectivity matrices.

One of the principal advantages of multi-layer networks is that they can aggregate many connectivity matrices into the same (multi-layer) network model, enabling those networks to be analyzed simultaneously without a loss of information. Among the many possible tools for analyzing multi-layer networks is a multi-layer analog of modularity maximization ([Bazzi et al., 2016](#); [Mucha et al., 2010](#)). Although we refer to it as an analog of the traditional single-layer modularity maximization,

it is in fact *identical* to the single-layer method, operating on the exact same principles and formulation.

Multi-layer modularity maximization works by creating a special kind of modularity matrix. Along the diagonal of this matrix are modularity matrices from individual layers; one for each of the K single-layer networks we want to analyze ([Fig. 4b](#)). If we were to submit this matrix to any modularity maximization algorithm nodes from different layers will *always* be assigned to different modules. That is, node i in layer s would never appear in the same community as node i in layer $t \neq s$. This is because there exist no cross-layer groupings of nodes that would lead to an improvement in Q . In other words, there is no practical difference between this approach and the case where we analyze each network independently.

However, we can change this fact by, once again, directly modifying the modularity matrix. Specifically, we can add a small amount of weight, ω , linking node i in layer s to itself in other layers². In fact, we can do this for every $i \in [1, \dots, N]$ ([Puxeddu et al., 2019](#)). The effect of the ω parameter is that there now exists scenarios where merging nodes from different layers can lead to improvements in modularity, i.e. $\Delta Q > 0$.

Because community labels are preserved from one layer to the next, we obviate the need to align communities. Comparing communities in different layers is as simple as assessing whether their labels are or are not identical. When averaged across all comparisons, this measure is referred to as *flexibility* ([Bassett et al., 2011, 2013](#), [Braun et al., 2015, 2016](#); [Chai et al., 2017](#); [Gerraty et al., 2018](#)) ([Fig. 4g](#)). This approach requires no averaging procedure and so no information about individual connectivity matrices is lost in the process.

In a way similar to γ , we can tune ω to detect communities more or less consistent across the slices of the multi-layer networks. Suppose we have a multi-subject network. It can be that all the subjects belong to a specific category, or are performing the same task, and we want to identify a common community organization that can be associated with a category or task. In this case, we could use high ω values and encourage the algorithm to find common structures across subjects. On the other hand, if we want to investigate inter-individual variability, or we have a more heterogeneous group of subjects to deal with, we are more likely to choose a low ω , that will highlight features unique to the subjects. Multi-layer modularity optimization with ω -tuning has been shown to be more robust to noise with respect to the single-layer case applied to each slice, and this is true in both conditions where we investigate steady structure within the ensemble of networks, or non-steady ones ([Puxeddu et al., 2021](#)).

Notably, however, multi-layer modularity maximization has other advantages, beyond simply preserving labels across layers. The addition of inter-layer coupling creates dependencies between layers so that, from the perspective of a given layer, global information about connectivity across all/adjacent layers is balanced with local information about a given layer to determine its community structure. The relative balance of global/local information is scaled by the magnitude of the inter-layer coupling parameter. In effect, we can view multi-layer modularity maximization as using group-level information to inform our estimates of communities in each layer. In many cases, this may actually be advantageous – as it helps avoid overfitting partitions to single layers, which may also be prone to noise and inaccuracies.

Although applications of multi-layer community detection are becoming increasingly common within the network neuroscience community, there remain (simple) opportunities for innovation. The following section will continue our discussion of different ways that multi-layer modularity has been used in neuroscience while highlighting some of the possibilities for improvement.

² note that the inter-layer coupling does not have to be in this form – it could, in principle, be added as any pattern of off-diagonal coupling between modularity matrices




Fig. 4. Multi-layer modularity framework. (a) Multi-layer modularity operates on network datasets comprising many distinct connectivity matrices or “layers”. (b) To detect communities, each connectivity matrix (or more accurately, its modularity matrix) is embedded along the block diagonal of a modularity tensor (shown flattened here). Then, a small amount of weight is added between each node and itself across layers. When networks have some temporal ordering, these additional weights may only connect adjacent layers. For networks without temporal order, weights can be added all-to-all. (c) The modularity tensor is then passed to an optimization algorithm to simultaneously partition nodes in each layer into communities. The advantage of this approach is that community labels are preserved across layers, making it possible to directly compare nodes’ community assignments across layers. In previous applications, multi-layer modularity has been applied to time-varying networks (d) or multi-subject datasets (e). In these diagrams, orange and green denote modularity matrices from FC and SC, respectively. However, the framework allows for more complicated configurations, including incorporating SC and FC simultaneously, which allows for the inclusion of a new class of inter-layer coupling (SC to FC; panel f). (g) Multi-layer partitions also make it possible to directly compare communities across layers using a measure of flexibility – how frequently assignments change from layer to layer. This diagram depicts the same communities as in panel c but with rows ordered in terms of flexibility. With different multi-layer tensors, it is possible to calculate different versions of flexibility. In time-varying networks flexibility is calculated by comparing communities only between temporally adjacent layers(h) while in multi-subject datasets, it makes sense to compare communities between all pairs of layers. (j) In multi-subject/multi-modal datasets, we can calculate different categories of flexibility. For instance, in addition to an all-to-all comparison, one could calculate the flexibility between FC layers or SC layers. One could also calculate SC-FC flexibility by comparing SC and FC layers within subjects.

4.1. Applications to time-varying connectivity

Presently, applications of multi-layer modularity maximization to neuroimaging data have been narrow and largely restricted to dynamic or time-varying networks. In these studies, each layer represents a functional network that describes connectivity at a given point in time, e.g. a narrow window in time during a scan session (Bassett et al., 2011; Telesford et al., 2016; Yang et al., 2021) or a particular developmental stage (Puxeddu et al., 2020; Yin et al., 2020).

In these examples, multi-layer modularity maximization makes it possible to track the formation, evolution, and dissolution of community structure across time. This is accomplished by adding weak connections between a node and itself in successive layers and using algorithms like the Louvain algorithm to optimize this multi-layer version of Q . The outcome, as noted earlier, is a series of community labels that are propagated across layers. Because these labels are preserved from one layer to another, comparing communities across layers is trivial.

Again, we want to emphasize that this feature – the correspondence of community labels across layers – is one of the principal advantages of multi-layer modularity maximization. Matching communities from one partition to another is, in general, non-trivial for single-layer networks and in some cases is ill-posed. Consider, as an example, two networks, made up of four nodes. In network 1, we find those nodes all assigned to the same community $\sigma_x = \{A, B, C, D\}$. In network 2, we find them split evenly into two communities $\sigma_y = \{A, B\}$ and $\sigma_z = \{C, D\}$. How is σ_x related to σ_y and σ_z ? Does σ_x fracture into two novel communities or does it continue on through one of either σ_y or σ_z ? If so, which of the two represents the continuation and which represents a novel community? Or are they both novel?

Addressing these types of questions can be challenging when communities are detected using single-layer modularity maximization. In all but the most trivial cases, matching communities across partitions requires the inclusion of some additional heuristic, e.g. to “break ties” like those described above. Multi-layer modularity maximization, on the other hand, incurs a greater up-front computational cost (increased memory to store multilayer networks, greater runtime to estimate communities), but automatically tracks communities across layers, effectively solving the matching problem. In doing so, it facilitates the calculation of a large suite of metrics, including *flexibility* (Fig. 4g), which measures how often a node changes communities from one layer to the next. Flexibility, in particular, has been used in many contexts, and has been linked to learning rate (Bassett et al., 2011), mood and affect (Betzel et al., 2017), clinical status (Braun et al., 2016), and executive function (Braun et al., 2015), among others. Note that multi-layer and single-layer modularity maximization are mathematically identical operations, and that the increased computational burden is a consequence of the size of the multi-layer modularity matrix, e.g. the need for additional random access memory, and, in the case of the Louvain algorithm, not the optimization heuristic itself (Lancichinetti and Fortunato, 2009).

4.2. Applications to multi-subject, multi-task, and dense sampling datasets

While most applications of multi-layer modularity maximization and flexibility analysis have been to time-varying FC, they have also been applied in other contexts. For instance, to investigate differences in modular structure across task state and across subjects (Betzel et al., 2019). The community structure of functional brain networks exhibits scale-specific patterns of inter- and intra-subject variability (Betzel et al., 2019). Just like with time-varying connectivity, multi-layer modularity can be used to detect modules in networks where layers represent connectivity estimated under different conditions (e.g. different tasks in the scanner) or even different individuals (Fig. 4e). As with applications to time-varying connectivity, flexibility analyses can be used to characterize changes in community structure across layers. In this case, however, the interpretation changes – flexibility becomes a measure task-induced reconfiguration or inter-subject stability of communities. Note that one could also com-

bine these two approaches and use multi-layer modularity to investigate changes in modular structure in multiple brains across time. This approach might be especially useful in studies using naturalistic stimuli (Betzel et al., 2020; Finn and Bandettini, 2021; Sonkusare et al., 2019), or hyperscanning experiments, in which brain activity is recorded simultaneously from two (or more) individuals while interacting with one another (Babiloni and Astolfi, 2014; Dumas et al., 2011).

Interestingly, flexibility can be broken down into components by comparing specific pairs of layers. Suppose, for instance, that a dataset was comprised of a control and clinical population. You could measure the flexibility (community dissimilarity) separately between all pairs of control and clinical subjects and then within both populations. These values could then be evaluated further using basic statistics to compare the variability of communities within and between populations. This approach could also be especially powerful in dense-sampling studies, where subjects are scanned many times in order to better characterize their subject-specific organization. In these types of studies, one could treat connectivity derived from repeated scans of many subjects and conditions and subsequently calculate different versions of flexibility. For instance, one could calculate flexibility only comparing communities within subjects to estimate the within-subject variability of communities. This would allow a user to trivially generate spatial maps of regions and the baseline variability of their community assignments within subjects. The same could be done between subjects, between conditions, and for different combinations. In spirit, this flexibility analysis is similar to other studies that have examined variability of FC patterns to ascertain what factors account for their variability (Gratton et al., 2018).

4.3. Multi-layer community detection for multi-modal datasets

There are yet other ways that multilayer modularity maximization could be further extended. For instance, one interesting application involves treating different connectivity modalities as layers (Bentley et al., 2016; Betzel et al., 2019). For instance, one could assemble a multi-layer network from structural and functional connectivity, along with morphometric and transcriptomic similarity matrices and cluster them simultaneously. Importantly, consideration should be taken that the weights of the matrices are comparable. If one layer has much stronger weights than the others, it will tend to contribute disproportionately to the multi-layer modularity and impact estimates of modular structure from other layers.

Looking even further into the future, one could also imagine jointly detecting communities using multi-modal data from many different subjects (Fig. 4f). A clear example is the case where one has a dataset of SC and FC matrices from many different individuals. As before, ensuring that the total weight of each layer and modality are comparable to one another is important. For instance, while SC matrices are usually sparse and made of positive weights, FC matrices are fully connected and made of positive and negative weights. This difference can be overcome by transforming the SC matrix into a structural correlation matrix by computing the Pearson's correlation between each pair of rows. In this way, the weights of both SC and FC would fall within the same range (Amico and Goñi, 2018). The user also has the opportunity to be creative in terms of how the inter-layer coupling parameter, ω , is introduced. For instance, one might imagine only coupling structural matrices to structural matrices across individuals but, within subjects, introducing a new parameter that governs the subject-specific coupling of structural and functional connectivity to one another.

When the inter-subject coupling of FC-to-FC or SC-to-SC is strong and the SC-to-FC coupling weak, optimizing the multi-layer modularity will return group-representative partitions for both SC and FC. However, as we increase the coupling between SC and FC at the subject level we also force the modularity maximization algorithm to detect modular structures that are shared between the two modalities. Again, this procedure allows for the calculating various flexibilities, including the flexibility between SC and FC, both globally (considering all sub-

jects) and at the micro-level of individual subjects. These subject-level measures of structure-function coupling could be used as phenotypical measures of an individual and their variation tracked over the course of development, across tasks, or clinical assessments.

5. Outlook and concluding remarks

This prospective article deals with brain network communities and focuses specifically on the modularity maximization heuristic. We argue that, despite its widespread application and success, modularity maximization has, so far, been applied narrowly within neuroscience and often without taking into account domain-specific knowledge. We highlight several simple strategies for broadening and extending the modularity maximization framework by directly modifying the modularity matrix to incorporate realistic features of brain networks, including spatial dependences between connections and the preservation of statistical relationships among connections. The primary aim of this work is not to present new findings or to overturn previous observations, but to draw attention to these modifications with the hope that they encourage users to creatively explore modifications of modularity maximization, leading to new neuroscientific insight.

Community structure is omnipresent across networks (Newman, 2012). Detecting this structure and designing new algorithms to do so is of great interest to practitioners and theoreticians, alike (Fortunato and Hric, 2016; Leskovec et al., 2009). Like other clustering algorithms, community detection reduces the dimensionality of large complex systems, identifies patterns, uncovers functionally-related groups of nodes, diagnoses node roles, and broadly generates new insight into system organization and function, prompting new scientific hypotheses.

Community detection is also important for neuroscience (Meunier et al., 2009; Sporns and Betzel, 2016). Shifts in data-sharing practices have generated massive, publicly-available datasets that include imaging data from thousands of individuals (Horien et al., 2021; Van Essen et al., 2013). In parallel, advances in cellular-level recordings have made it possible to record from up to millions of neurons simultaneously (Demas et al., 2021). Making sense of these data requires flexible tools for dimension reduction and for extracting neurobiologically meaningful features that can be propagated to secondary analyses.

Indeed, community detection, in general, and modularity maximization, specifically, have been applied to brain data at virtually all spatiotemporal scales, offering a convenient method for partitioning nervous systems into communities. However, in almost every instance, these applications have leveraged an “out-of-the-box” version of modularity maximization. On one hand, the fact that neuroscientific insight can be gleaned in this way is a testament to the universality of network science and the tools used to interrogate networks. Modularity maximization was designed to be generic and not necessarily with brain networks in mind. On the other hand, this leaves lots of room for improvement. As noted earlier, the generic version of modularity maximization features a null model that is poorly-suited for both SC and FC, where it fails to preserve basic features of brain networks or outright violates statistical truisms.

The aim of this article, however, is not to wag a finger at past applications of modularity maximization. Rather, we endeavor to highlight several strategies for extending modularity maximization so that future applications can make stronger and more nuanced scientific claims and target increasingly specific research questions. Additionally, we want to highlight the components of modularity maximization – null models, multi-layer formulations, parameterizations – that are fundamentally left up to the user and to encourage exploration of these components. It is our view that the utility of “out-of-the-box” and generic methods, although they have been used to make important contributions up to this point, is limited and can only carry the field so far. Future work must be directed to link data-driven and network science methods with domain-specific knowledge from neuroscience. In the following paragraphs, we

cover several topics that are within the scope of this prospective article but were not discussed in depth in the main text.

Occasionally, modularity maximization may place some nodes into isolated communities referred to as “singletons.” In general, there are several reasons why singletons might appear. If a network has been sparsified by thresholding weak/inconsistent connections, it is possible that the resulting network fragments into disconnected components, some of which may comprise only a single node. In this case, a simple solution is to construct a maximum spanning tree (ensuring connectedness) and adding to this network the strongest and/or most consistent connections to reach a desired threshold. Another reason why a node might be partitioned into a singleton community is that all its connections to the rest of the network are weaker than would be expected under the null connectivity model. This type of singleton can often be reincorporated into a larger community by changing the value of the resolution parameter, γ . Alternatively, the multi-layer approach can also help avoid this issue. In the multi-layer framework, information about other layers is used to determine the community structure for any single layer. Even if a node is disconnected, makes weak connections to the rest of the network, or the inter-layer coupling parameter is sufficiently strong, it may be advantageous from the perspective of the multi-layer modularity function to incorporate that node into a larger community. Of course, a final possibility is that a singleton community reflects a true architectural feature of a network. It could represent a brain area that, under a specific set of conditions (task, disease, etc.) dissociates from the rest of the network.

The focus of this article has been on modularity maximization, and for good reason. It operates according to an eminently simple principle, seeking communities of nodes whose internal density of connections is maximally greater than what would be expected under a chance model. Modularity maximization is a relatively old method (Newman and Girvan, 2004) and has been applied in virtually every scientific domain in which networks have made ingress. Optimizing modularity yields communities that are strictly assortative – nodes in the same community are more likely to connect to other nodes in the same community than to nodes in other communities. This type of organization, however, is one of many possible types of community configurations. Among these are disassortative – where nodes in the same community preferentially avoid making connections to one another – and core-periphery structure, where a densely connected core projects to a periphery and the peripheral nodes avoid making connections to other peripheral nodes. Modularity maximization is incapable of resolving these types of organization.

However, there exist many other methods that are capable of discovering non-assortative communities. Among the most popular methods is the stochastic blockmodel, which seeks to discover groups of nodes whose connectivity patterns are similar to one another. This allows for the possibility that groups of nodes are not connected to one another but nonetheless exhibit similar connectivity profiles, e.g. core-periphery or disassortative communities (Betzel et al., 2018; Faskowitz et al., 2018; Moyer et al., 2015).

Unlike modularity maximization (and Infomap), blockmodels can flexibly detect different categories of community interactions. This is an important point; modularity and Infomap are capable of detecting only assortative communities. If the aim is to obtain an estimate of the brain's true community structure, we might expect these methods to perform reasonably well, provided those ground-truth communities are assortative. However, if the brain deviates from this type of organization, then modularity and Infomap may perform poorly and misclassify nodes' communities. Indeed, these misclassifications can further impact our understanding of the brain's organization and function. Misclassified nodes may make connections outside of the assortative community to which they are assigned, giving the impression that they are integrative hubs that span modules. A blockmodel, however, might discover a more parsimonious description of communities where these putative “hubs” are grouped together with other hubs and form far fewer cross-module links (Pavlovic et al., 2014). Additionally, blockmodels are, effectively,

generative models of a network and serve as a powerful way of performing posterior predictive checks – i.e. simulating data from the fit model – which can be used to assess how well the model recapitulates the observed network and to identify features that are not accounted for by the current model (Gelman et al., 2013).

In general, no community detection algorithm is privileged over the others, including modularity maximization. Community detection can be viewed as an inversion of some generative function for networks that depends on nodes' latent community assignments. Different algorithms and methods effectively “guess” at what that function is to produce estimates of communities. However, there is almost always a many-to-one mapping of a fixed set of communities to a given network and, similarly, there may be different sets of communities that, through different generative functions, explain the observed network equally well (Peel et al., 2017). Unless one knows the true generative function, then it is impossible to determine with certainty which method and its inferred communities is optimal. One might imagine that this issue could be addressed by cross-validating with respect to some additional data, e.g. verifying that detected structural communities are enriched for functional connectivity or gene coexpression patterns. However, this strategy suffers from the same issue, in that there may exist multiple mappings between these additional data and the detected communities. In short, these observations motivate exploring multiple community detection algorithms and not weighting the results of any single algorithm heavily, but focusing instead on communities that are shared and robust, but also points of frustration, where different algorithms fail to arrive at consensus.

Since the recent explosion of interest in networks, neuroscience has repeatedly borrowed methodology from network science proper. Small-worlds, hubs, rich-clubs, and modules are concepts that make sense for brain network data, but the tools used to detect these types of structures originated in network science, often with applications to other systems in mind, e.g. social networks. While we focus here on extending one of these tools, this represents only the first steps. A key challenge for the future is to reverse this directionality and to develop analysis tools that are specific to nervous systems, taking into account their generative mechanisms, peculiarities, and domain-specific knowledge. Doing so creates bridges with other sub-disciplines in the computational neurosciences, opening up opportunities for increasingly sophisticated exploration of brain network organization. Specifically, modularity maximization presents a flexible framework for exploring community structure in brain networks. It allows for the user to test and compare communities detected under distinct null hypotheses, identifying commonalities and idiosyncrasies. While these types of approaches can sometimes lead to confusion – comparing modules detected using different null models may present issues with interpretation – they also help facilitate neuroscientific discovery, helping us to identify what is being lost/gained by choosing increasingly domain-specific null models, while offering new perspectives and insight on the organization and function of brain networks.

5.1. Software

Matlab code for implementing the models discussed in the paper can be found via the following link: https://github.com/brain-networks/nonstandard_modularity_maximization.

Acknowledgment

RFB wrote initial draft of manuscript. All authors edited and revised manuscript and contributed and tested code.

References

Adachi, Y., Osada, T., Sporns, O., Watanabe, T., Matsui, T., Miyamoto, K., Miyashita, Y., 2012. Functional connectivity between anatomically unconnected areas is shaped by collective network-level effects in the macaque cortex. *Cereb. Cortex* 22 (7), 1586–1592.

Ahn, Y.-Y., Bagrow, J.P., Lehmann, S., 2010. Link communities reveal multiscale complexity in networks. *Nature* 466 (7307), 761–764.

Akarc, D., Vértés, P.E., Bullmore, E.T., Astle, D.E., 2021. A generative network model of neurodevelopmental diversity in structural brain organization. *Nat. Commun.* 12 (1), 1–18.

Akiki, T.J., Abdallah, C.G., 2019. Determining the hierarchical architecture of the human brain using subject-level clustering of functional networks. *Sci. Rep.* 9 (1), 1–15.

Alexander-Bloch, A.F., Gogtay, N., Meunier, D., Birn, R., Clasen, L., Lalonde, F., Lenroot, R., Giedd, J., Bullmore, E.T., 2010. Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia. *Front. Syst. Neurosci.* 4, 147.

Amico, E., Goñi, J., 2018. Mapping hybrid functional-structural connectivity traits in the human connectome. *Netw. Neurosci.* 2 (3), 306–322.

Arenas, A., Fernandez, A., Gomez, S., 2008. Analysis of the structure of complex networks at different resolution levels. *New J. Phys.* 10 (5), 053039.

Avena-Koenigsberger, A., Misic, B., Sporns, O., 2018. Communication dynamics in complex brain networks. *Nat. Rev. Neurosci.* 19 (1), 17.

Babiloni, F., Astolfi, L., 2014. Social neuroscience and hyperscanning techniques: past, present and future. *Neurosci. Biobehav. Rev.* 44, 76–93.

Bassett, D.S., Bullmore, E., 2006. Small-world brain networks. *Neuroscientist* 12 (6), 512–523.

Bassett, D.S., Greenfield, D.L., Meyer-Lindenberg, A., Weinberger, D.R., Moore, S.W., Bullmore, E.T., 2010. Efficient physical embedding of topologically complex information processing networks in brains and computer circuits. *PLoS Comput. Biol.* 6 (4), e1000748.

Bassett, D.S., Porter, M.A., Wymbs, N.F., Grafton, S.T., Carlson, J.M., Mucha, P.J., 2013. Robust detection of dynamic community structure in networks. *Chaos* 23 (1), 013142.

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. *Proc. Natl. Acad. Sci.* 108 (18), 7641–7646.

Baum, G.L., Ciric, R., Roalf, D.R., Betzel, R.F., Moore, T.M., Shinohara, R.T., Kahn, A.E., Vandekar, S.N., Rupert, P.E., Quarmley, M., et al., 2017. Modular segregation of structural brain networks supports the development of executive function in youth. *Curr. Biol.* 27 (11), 1561–1572.

Baum, G.L., Cui, Z., Roalf, D.R., Ciric, R., Betzel, R.F., Larsen, B., Cieslak, M., Cook, P.A., Xia, C.H., Moore, T.M., et al., 2020. Development of structure–function coupling in human brain networks during youth. *Proc. Natl. Acad. Sci.* 117 (1), 771–778.

Bazzi, M., Porter, M.A., Williams, S., McDonald, M., Penn, D.J., Howison, S.D., 2016. Community detection in temporal multilayer networks, with an application to correlation networks. *Multiscale Model. Simul.* 14 (1), 1–41.

Bellec, P., Perlbarg, V., Jbabdi, S., Pélégriani-Issac, M., Anton, J.-L., Doyon, J., Benali, H., 2006. Identification of large-scale networks in the brain using fmri. *NeuroImage* 29 (4), 1231–1243.

Zalesky, A., Fornito, A., Bullmore, E.T., 2010. Network-based statistic: identifying differences in brain networks. *NeuroImage* 53 (4), 1197–1207.

Zalesky, A., Fornito, A., Bullmore, E., 2012. On the use of correlation as a measure of network connectivity. *NeuroImage* 60 (4), 2096–2106.

Betzel, R. F., 2020arXiv preprint arXiv:2011.06723

Bentley, B., Brancic, R., Barnes, C.L., Chew, Y.L., Yemini, E., Bullmore, E.T., Vértés, P.E., Schäfer, W.R., 2016. The multilayer connectome of *Caenorhabditis elegans*. *PLoS Comput. Biol.* 12 (12), e1005283.

Betzel, R.F., Avena-Koenigsberger, A., Goñi, J., He, Y., De Reus, M.A., Griffa, A., Vértés, P.E., Mišić, B., Thiran, J.-P., Hagmann, P., et al., 2016. Generative models of the human connectome. *NeuroImage* 124, 1054–1064.

Betzel, R.F., Bassett, D.S., 2017. Generative models for network neuroscience: prospects and promise. *J. R. Soc. Interface* 14 (136), 20170623.

Betzel, R.F., Byrge, L., Esfahlani, F.Z., Kennedy, D.P., 2020. Temporal fluctuations in the brain's modular architecture during movie-watching. *NeuroImage* 116687.

Betzel, R.F., Byrge, L., He, Y., Goñi, J., Zuo, X.-N., Sporns, O., 2014. Changes in structural and functional connectivity among resting-state networks across the human lifespan. *NeuroImage* 102, 345–357.

Betzel, R.F., Griffa, A., Hagmann, P., Mišić, B., 2019. Distance-dependent consensus thresholds for generating group-representative structural brain networks. *Netw. Neurosci.* 3 (2), 475–496.

Betzel, R.F., Medaglia, J.D., Bassett, D.S., 2018. Diversity of meso-scale architecture in human and non-human connectomes. *Nat. Commun.* 9 (1), 1–14.

Betzel, R.F., Medaglia, J.D., Kahn, A.E., Soffer, J., Schonhaut, D.R., Bassett, D.S., 2019. Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography. *Nat. Biomed. Eng.* 3 (11), 902–916.

Betzel, R.F., Medaglia, J.D., Papadopoulos, L., Baum, G.L., Gur, R., Gur, R., Roalf, D., Satterthwaite, T.D., Bassett, D.S., 2017. The modular organization of human anatomical brain networks: Accounting for the cost of wiring. *Netw. Neurosci.* 1 (1), 42–68.

Betzel, R.F., Satterthwaite, T.D., Gold, J.I., Bassett, D.S., 2017. Positive affect, surprise, and fatigue are correlates of network flexibility. *Sci. Rep.* 7 (1), 1–10.

BetzelThe community structure of functional brain networks exhibits scale-specific patterns of inter- and intra-subject variability. R.F., Bertolero, M.A., Gordon, E.M., Gratton, C., Dosenbach, N.U., Bassett, D.S., 2019. *NeuroImage* 202, 115990.

Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E., 2008. Fast unfolding of communities in large networks. *J. Stat. Mech.* 2008 (10), P10008.

Braun, U., Schäfer, A., Bassett, D.S., Rausch, F., Schweiger, J.I., Bilek, E., Erk, S., Romanczuk-Seiferth, N., Grimm, O., Geiger, L.S., et al., 2016. Dynamic brain network reconfiguration as a potential schizophrenia genetic risk mechanism modulated by nmda receptor function. *Proc. Natl. Acad. Sci.* 113 (44), 12568–12573.

Braun, U., Schäfer, A., Walter, H., Erk, S., Romanczuk-Seiferth, N., Haddad, L., Schweiger, J.I., Grimm, O., Heinz, A., Tost, H., et al., 2015. Dynamic reconfiguration

- of frontal brain networks during executive cognition in humans. *Proc. Natl. Acad. Sci.* 112 (37), 11678–11683.
- Brodman, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Barth.
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10 (3), 186–198.
- Burt, J.B., Helmer, M., Shinn, M., Anticevic, A., Murray, J.D., 2020. Generative modeling of brain maps with spatial autocorrelation. *NeuroImage* 220, 117038.
- Chai, L.R., Khambhati, A.N., Ciric, R., Moore, T.M., Gur, R.C., Gur, R.E., Satterthwaite, T.D., Bassett, D.S., 2017. Evolution of brain network dynamics in neurodevelopment. *Netw. Neurosci.* 1 (1), 14–30.
- Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83 (1), 238–251.
- Crossley, N.A., Mechelli, A., Vertes, P.E., Winton-Brown, T.T., Patel, A.X., Ginetet, C.E., McGuire, P., Bullmore, E.T., 2013. Cognitive relevance of the community structure of the human brain functional coactivation network. *Proc. Natl. Acad. Sci.* 110 (28), 11583–11588.
- De Domenico, M., 2017. Multilayer modeling and analysis of human brain networks. *Giga Sci.* 6 (5), gix004.
- Deco, G., McIntosh, A.R., Shen, K., Hutchison, R.M., Menon, R.S., Everling, S., Hagmann, P., Jirsa, V.K., 2014. Identification of optimal structural connectivity using functional connectivity and neural modeling. *J. Neurosci.* 34 (23), 7910–7916.
- Demas, J., Manley, J., Tejera, F., Kim, H., Traub, F.M., Chen, B., Vaziri, A., 2021. High-speed, cortex-wide volumetric recording of neuroactivity at cellular resolution using light beads microscopy. *bioRxiv*.
- Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., et al., 2006. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *NeuroImage* 31 (3), 968–980.
- Doron, K.W., Bassett, D.S., Gazzaniga, M.S., 2012. Dynamic network structure of inter-hemispheric coordination. *Proc. Natl. Acad. Sci.* 109 (46), 18661–18668.
- Dumas, P., Lachat, F., Martinerie, J., Nadel, J., George, N., 2011. From social behaviour to brain synchronization: review and perspectives in hyperscanning. *Irbm* 32 (1), 48–53.
- Ercsey-Ravasz, M., Markov, N.T., Lamy, C., Van Essen, D.C., Knoblauch, K., Toroczkai, Z., Kennedy, H., 2013. A predictive network model of cerebral cortical connectivity based on a distance rule. *Neuron* 80 (1), 184–197.
- Esfahlani, F.Z., Bertolero, M.A., Bassett, D.S., Betzel, R.F., 2020. Space-independent community and hub structure of functional brain networks. *NeuroImage* 211, 116612.
- Evans, T., Lambiotte, R., 2009. Line graphs, link partitions, and overlapping communities. *Phys. Rev. E* 80 (1), 016105.
- Expert, G., Evans, T.S., Blondel, V.D., Lambiotte, R., 2011. Uncovering space-independent communities in spatial networks. *Proc. Natl. Acad. Sci.* 108 (19), 7663–7668.
- Faskowitz, J., Esfahlani, F.Z., Jo, Y., Sporns, O., Betzel, R.F., 2020. Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nat. Neurosci.* 23 (12), 1644–1654.
- Faskowitz, J., Yan, X., Zuo, X.-N., Sporns, O., 2018. Weighted stochastic block models of the human connectome across the life span. *Sci. Rep.* 8 (1), 1–16.
- Finn, E.S., Bandettini, P.A., 2021. Movie-watching outperforms rest for functional connectivity-based prediction of behavior. *NeuroImage* 235, 117963.
- Fortunato, S., 2010. Community detection in graphs. *Phys. Rep.* 486 (3–5), 75–174.
- Fortunato, S., Barthelemy, M., 2007. Resolution limit in community detection. *Proc. Natl. Acad. Sci.* 104 (1), 36–41.
- Fortunato, S., Hric, D., 2016. Community detection in networks: A user guide. *Phys. Rep.* 659, 1–44.
- Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* 2 (1–2), 56–78.
- Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. *Bayesian Data Analysis*. CRC press.
- Gerraty, R.T., Davidow, J.Y., Foerde, K., Galvan, A., Bassett, D.S., Shohamy, D., 2018. Dynamic flexibility in striatal-cortical circuits supports reinforcement learning. *J. Neurosci.* 38 (10), 2442–2453.
- Gollo, L.L., Roberts, J.A., Cropley, V.L., Di Biase, M.A., Pantelis, C., Zalesky, A., Breakpear, M., 2018. Fragility and volatility of structural hubs in the human connectome. *Nat. Neurosci.* 21 (8), 1107–1116.
- Gómez, S., Jensen, P., Arenas, A., 2009. Analysis of community structure in networks of correlated data. *Phys. Rev. E* 80 (1), 016114.
- Good, B.H., De Montjoye, Y.-A., Clauset, A., 2010. Performance of modularity maximization in practical contexts. *Phys. Rev. E* 81 (4), 046106.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2016. Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cereb. Cortex* 26 (1), 288–303.
- Gratton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., et al., 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or metabolic variation. *Neuron* 98 (2), 439–452.
- Guimera, R., Amaral, L.A.N., 2005. Functional cartography of complex metabolic networks. *Nature* 433 (7028), 895–900.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C.J., Wedeen, V.J., Sporns, O., 2008. Mapping the structural core of human cerebral cortex. *PLoS Biol.* 6 (7), e159.
- Harris, J.A., Mihalas, S., Hirokawa, K.E., Whitesell, J.D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., et al., 2019. Hierarchical organization of cortical and thalamic connectivity. *Nature* 575 (7781), 195–202.
- Hirschberger, M., Qi, Y., Steuer, R.E., 2007. Randomly generating portfolio-selection covariance matrices with specified distributional characteristics. *Eur. J. Oper. Res.* 177 (3), 1610–1625.
- Honey, C.J., Kötter, R., Breakpear, M., Sporns, O., 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc. Natl. Acad. Sci.* 104 (24), 10240–10245.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.-P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci.* 106 (6), 2035–2040.
- Horien, C., Noble, S., Greene, A.S., Lee, K., Barron, D.S., Gao, S., O'Connor, D., Salehi, M., Dadashkarimi, J., Shen, X., et al., 2021. A hitchhiker's guide to working with large, open-source neuroimaging datasets. *Nat. Hum. Behav.* 5 (2), 185–193.
- Horvát, S., Gámánut, R., Ercsey-Ravasz, M., Magrou, L., Gámánut, B., Van Essen, D.C., Burkhalter, A., Knoblauch, K., Toroczkai, Z., Kennedy, H., 2016. Spatial embedding and wiring cost constrain the functional layout of the cortical network of rodents and primates. *PLoS Biol.* 14 (7), e1002512.
- Horwitz, B., 2003. The elusive concept of brain connectivity. *NeuroImage* 19 (2), 466–470.
- 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.
- Jutla, I. S., Jeub, L. G., Mucha, P. J., 2011. A generalized Louvain method for community detection implemented in MATLAB. URL <http://netwiki.amath.unc.edu/GenLouvain>.
- Jeub, L.G., Sporns, O., Fortunato, S., 2018. Multiresolution consensus clustering in networks. *Sci. Rep.* 8 (1), 1–16.
- Kenett, Y.N., Betzel, R.F., Beaty, R.E., 2020. Community structure of the creative brain at rest. *NeuroImage* 210, 116578.
- Lambiotte, R., Delvenne, J.-C., Barahona, M., 2008. arXiv preprint arXiv:0812.1770.
- Kivela, M., Arenas, A., Barthelemy, M., Gleeson, J.P., Moreno, Y., Porter, M.A., 2014. Multilayer networks. *J. Complex Netw.* 2 (3), 203–271.
- Lancichinetti, A., Fortunato, S., 2009. Community detection algorithms: a comparative analysis. *Phys. Rev. E* 80 (5), 056117.
- Lancichinetti, A., Fortunato, S., 2012. Consensus clustering in complex networks. *Sci. Rep.* 2 (1), 1–7.
- Leskovec, J., Lang, K.J., Dasgupta, A., Mahoney, M.W., 2009. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Math.* 6 (1), 29–123.
- MacMahon, M., Garlaschelli, D., 2013. arXiv preprint arXiv:1311.1924.
- Lurie, D.J., Kessler, D., Bassett, D.S., Betzel, R.F., Breakpear, M., Kheilholtz, S., Kucyi, A., Liégeois, R., Lindquist, M.A., McIntosh, A.R., et al., 2020. Questions and controversies in the study of time-varying functional connectivity in resting fmri. *Netw. Neurosci.* 4 (1), 30–69.
- Markov, N.T., Ercsey-Ravasz, M., Ribeiro Gomes, A., Lamy, C., Magrou, L., Vezoli, J., Misery, P., Falchier, A., Quilodran, R., Gariel, M., et al., 2014. A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cereb. Cortex* 24 (1), 17–36.
- Medaglia, J.D., Lynall, M.-E., Bassett, D.S., 2015. Cognitive network neuroscience. *J. Cogn. Neurosci.* 27 (8), 1471–1491.
- Meunier, D., Lambiotte, R., Fornito, A., Ersche, K., Bullmore, E.T., 2009. Hierarchical modularity in human brain functional networks. *Front. Neuroinform.* 3, 37.
- Moyer, D., Gutman, B., Prasad, G., Faskowitz, J., Ver Steeg, G., Thompson, P., 2015. Blockmodels for connectome analysis. In: 11th International Symposium on Medical Information Processing and Analysis, vol. 9681. International Society for Optics and Photonics, p. 96810A.
- Mucha, P.J., Richardson, T., Macon, K., Porter, M.A., Onnela, J.-P., 2010. Community structure in time-dependent, multiscale, and multiplex networks. *Science* 328 (5980), 876–878.
- Newman, M.E., 2012. Communities, modules and large-scale structure in networks. *Nat. Phys.* 8 (1), 25–31.
- Newman, M.E., Girvan, M., 2004. Finding and evaluating community structure in networks. *Phys. Rev. E* 69 (2), 026113.
- Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M., et al., 2014. A mesoscale connectome of the mouse brain. *Nature* 508 (7495), 207–214.
- Palla, G., Derényi, I., Farkas, I., Vicsek, T., 2005. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* 435 (7043), 814–818.
- Park, H.-J., Friston, K., 2013. Structural and functional brain networks: from connections to cognition. *Science* 342 (6158).
- Pavlovic, D.M., Vertes, P.E., Bullmore, E.T., Schafer, W.R., Nichols, T.E., 2014. Stochastic blockmodeling of the modules and core of the caenorhabditis elegans connectome. *PLoS ONE* 9 (7), e97584.
- Peel, L., Larremore, D.B., Clauset, A., 2017. The ground truth about metadata and community detection in networks. *Sci. Adv.* 3 (5), e1602548.
- Peixoto, T.P., 2012. Entropy of stochastic blockmodel ensembles. *Phys. Rev. E* 85 (5), 056122.
- Peixoto, T.P., 2014. Hierarchical block structures and high-resolution model selection in large networks. *Phys. Rev. X* 4 (1), 011047.
- Pons, P., Latapy, M., 2006. Computing communities in large networks using random walks. *J. Graph Algorithms Appl. CiteSeer*.
- Puxeddu, M.G., Petti, M., Astolfi, L., 2021. A comprehensive analysis of multilayer community detection algorithms for application to eeg-based brain networks. *Front. Syst. Neurosci.* 15.
- 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., et al., 2011. Functional network organization of the human brain. *Neuron* 72 (4), 665–678.
- Puxeddu, M.G., Faskowitz, J., Betzel, R.F., Petti, M., Astolfi, L., Sporns, O., 2020. The modular organization of brain cortical connectivity across the human lifespan. *NeuroImage* 218, 116974.

- Puxeddu, M. G., Petti, M., Mattia, D., Astolfi, L., 2019. The optimal setting for multilayer modularity optimization in multilayer brain networks. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, pp. 624–627.
- Reichardt, J., Bornholdt, S., 2006. Statistical mechanics of community detection. *Phys. Rev. E* 74 (1), 016110.
- Roberts, J.A., Perry, A., Lord, A.R., Roberts, G., Mitchell, P.B., Smith, R.E., Calamante, F., Breakspear, M., 2016. The contribution of geometry to the human connectome. *NeuroImage* 124, 379–393.
- Rosvall, M., Bergstrom, C.T., 2008. Maps of random walks on complex networks reveal community structure. *Proc. Natl. Acad. Sci.* 105 (4), 1118–1123.
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. *NeuroImage* 52 (3), 1059–1069.
- Rubinov, M., Sporns, O., 2011. Weight-conserving characterization of complex functional brain networks. *NeuroImage* 56 (4), 2068–2079.
- Sanchez-Rodriguez, L.M., Iturria-Medina, Y., Mouches, P., Sotero, R.C., 2021. Detecting brain network communities: considering the role of information flow and its different temporal scales. *NeuroImage* 225, 117431.
- Sarwar, T., Ramamohanarao, K., Zalesky, A., 2019. Detecting brain network communities: considering the role of information flow and its different temporal scales. *Magn. Reson. Med.* 81 (2), 1368–1384.
- Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T., 2018. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cereb. Cortex* 28 (9), 3095–3114.
- Schaub, M.T., Delvenne, J.-C., Yaliraki, S.N., Barahona, M., 2012. Markov dynamics as a zooming lens for multiscale community detection: non clique-like communities and the field-of-view limit. *PLoS ONE* 7 (2), e32210.
- Shinn, M., Hu, A., Turner, L., Noble, S., Achard, S., Anticevic, A., Scheinost, D., Constable, R.T., Lee, D., Bullmore, E.T., et al., 2021. Spatial and temporal autocorrelation weave human brain networks. *bioRxiv*.
- Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., et al., 2009. Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci.* 106 (31), 13040–13045.
- Sonkusare, S., Breakspear, M., Guo, C., 2019. Naturalistic stimuli in neuroscience: critically acclaimed. *Trends Cogn. Sci.* 23 (8), 699–714.
- Sporns, O., Betzel, R.F., 2016. Modular brain networks. *Annu. Rev. Psychol.* 67, 613–640.
- Sporns, O., Tononi, G., Kötter, R., 2005. The human connectome: a structural description of the human brain. *PLoS Comput. Biol.* 1 (4), e42.
- Sporns, O., Zwi, J.D., 2004. The small world of the cerebral cortex. *Neuroinformatics* 2 (2), 145–162.
- Stiso, J., Bassett, D.S., 2018. Spatial embedding imposes constraints on neuronal network architectures. *Trends Cogn. Sci.* 22 (12), 1127–1142.
- Telesford, Q.K., Lynall, M.-E., Vettel, J., Miller, M.B., Grafton, S.T., Bassett, D.S., 2016. Detection of functional brain network reconfiguration during task-driven cognitive states. *NeuroImage* 142, 198–210.
- Traag, V.A., Van Dooren, P., Nesterov, Y., 2011. Narrow scope for resolution-limit-free community detection. *Phys. Rev. E* 84 (1), 016114.
- Traag, V.A., Waltman, L., Van Eck, N.J., 2019. From louvain to leiden: guaranteeing well-connected communities. *Sci. Rep.* 9 (1), 1–12.
- Vaiana, M., Muldoon, S.F., 2018. Multilayer brain networks. *J. Nonlinear Sci.* 1–23.
- Van Den Heuvel, M.P., Sporns, O., 2011. Rich-club organization of the human connectome. *J. Neurosci.* 31 (44), 15775–15786.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Wu-Minn HCP Consortium, et al., 2013. The wu-minn human connectome project: an overview. *NeuroImage* 80, 62–79.
- Vértes, P.E., Alexander-Bloch, A.F., Gogtay, N., Giedd, J.N., Rapoport, J.L., Bullmore, E.T., 2012. Simple models of human brain functional networks. *Proc. Natl. Acad. Sci.* 109 (15), 5868–5873.
- White, J.G., Southgate, E., Thomson, J.N., Brenner, S., 1986. The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 314 (1165), 1–340.
- Wig, G.S., 2017. Segregated systems of human brain networks. *Trends Cogn. Sci.* 21 (12), 981–996.
- Yang, Z., Telesford, Q.K., Franco, A.R., Lim, R., Gu, S., Xu, T., Ai, L., Castellanos, F.X., Yan, C.-G., Colcombe, S., et al., 2021. Measurement reliability for individual differences in multilayer network dynamics: Cautions and considerations. *NeuroImage* 225, 117489.
- Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.*
- Yin, W., Li, T., Hung, S.-C., Zhang, H., Wang, L., Shen, D., Zhu, H., Mucha, P.J., Cohen, J.R., Lin, W., 2020. The emergence of a functionally flexible brain during early infancy. *Proc. Natl. Acad. Sci.* 117 (38), 23904–23913.