

Discovering Functional Communities in Dynamical Networks

Cosma Rohilla Shalizi¹, Marcelo F. Camperi², and Kristina Lisa Klinkner¹

¹ Statistics Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA
cshalizi@cmu.edu, klinkner@cmu.edu

² Physics Department, University of San Francisco, San Francisco, CA 94118 USA
camperi@usfca.edu

Abstract. Many networks are important because they are substrates for dynamical systems, and their pattern of functional connectivity can itself be dynamic — they can functionally reorganize, even if their underlying anatomical structure remains fixed. However, the recent rapid progress in discovering the community structure of networks has overwhelmingly focused on that constant anatomical connectivity. In this paper, we lay out the problem of discovering *functional communities*, and describe an approach to doing so. This method combines recent work on measuring information sharing across stochastic networks with an existing and successful community-discovery algorithm for weighted networks. We illustrate it with an application to a large biophysical model of the transition from beta to gamma rhythms in the hippocampus.

1 Introduction

The *community discovery* problem for networks is that of splitting a graph, representing a group of interacting processes or entities, into sub-graphs (communities) which are somehow modular, so that the nodes belonging to a given sub-graph interact with the other members more strongly than they do with the rest of the network. As the word “community” indicates, the problem has its roots in the study of social structure and cohesion [1, 2, 3], but is related to both general issues of clustering in statistical data mining [4] and to the systems-analysis problem of decomposing large systems into weakly-coupled sub-systems [5, 6, 7].

The work of Newman and Girvan [8] has inspired a great deal of research on statistical-mechanical approaches to community detection in complex networks. (For a recent partial review, see [9].) To date, however, this tradition has implicitly assumed that the network is defined by persistent, if not static, connections between nodes, whether through concrete physical channels (e.g., electrical power grids, nerve fibers in the brain), or through enduring, settled patterns of interaction (e.g., friendship and collaboration networks). However, networks can also be defined through *coordinated behavior*, and the associated sharing of dynamical information; neuroscience distinguishes these as, respectively, “anatomical” and “functional” connectivity [10, 11]. The two sorts of

connectivity do not map neatly onto each other, and it would be odd if functional modules always lined up with anatomical ones. Indeed, the same system could have many different sets of functional communities in different dynamical regimes. For an extreme case, consider globally-coupled map lattices [12], which are important statistical-mechanical models of physical and biological pattern formation [13]. In these systems, the “anatomical” network is fully connected, so there is only a single (trivial) community. Nonetheless, in some dynamical regimes they spontaneously develop many functional communities, i.e., groups of nodes which are internally coherent but with low inter-group coordination [14].³

Coupled map lattices are mathematical models, but the distinction between anatomical and functional communities is not merely a conceptual possibility. Observation of neuronal networks *in vivo* show that it is fairly common for, e.g., central pattern generators to change their functional organization considerably, depending on which pattern they are generating, while maintaining a constant anatomy [15]. Similarly, neuropsychological evidence has long suggested that there is no one-to-one mapping between higher cognitive functions and specialized cortical modules, but rather that the latter participate in multiple functions and vice versa, re-organizing depending on the task situation [16]. Details of this picture, of specialized anatomical regions supporting multiple patterns of functional connectivity, have more recently been filled in by brain imaging studies [11]. Similar principles are thought to govern the immune response, cellular signaling, and other forms of biological information processing [17]. Thus, in analyzing these biological networks, it would be highly desirable to have a way of detecting functional communities, rather than just anatomical ones. Similarly, while much of the work on social network organization concerns itself with the persistent ties which are analogous to anatomy, it seems very likely [18, 19] that these communities cut in complicated ways across the functional ones defined by behavioral coordination [20, 21] or information flow [22]. This is perhaps particularly true of modern societies, which are thought, on several grounds [23, 24, 19] to be more flexibly organized than traditional ones.

In this paper, we propose a two-part method to discover functional communities in network dynamical systems. Section 2.1 describes the first part, which is to calculate, across the whole of the network, an appropriate measure of behavioral coordination or information sharing; we argue that informational coherence, introduced in our prior work [25], provides such a measure. Section 2.2 describes the other half of our method, using our measure of coordination in place of a traditional adjacency matrix in a suitable community-discovery algorithm. Here we employ the Potts model procedure proposed by Reichardt and Bornholdt [26, 27]. Section 2.3 summarizes the method and clarifies the meaning of the functional communities it finds. Sections 3 and 4 apply our method to a detailed biophysical model of collective oscillations in the hippocampus [28], where it allows us to detect the functional re-organization accompanying the transition

³ We are preparing a separate paper on functional communities in coupled map lattices.

from gamma to beta rhythms. Finally, Sect. 5 discusses the limitations of our method and its relations to other approaches (Sect. 5.1) and some issues for future work (Sect. 5.2).

2 Discovering Behavioral Communities

There are two parts to our method for finding functional communities. We first calculate a measure of the behavioral coordination between all pairs of nodes in the network: here, the informational coherence introduced in [25]. We then feed the resulting matrix into a suitable community-discovery algorithm, in place of the usual representation of a network by its adjacency matrix. Here, we have used the Reichardt-Bornholdt algorithm [26], owing to its Hamiltonian form, its ability to handle weighted networks, and its close connection to modularity.

2.1 Informational Coherence

We introduced *informational coherence* in [25] to measure the degree to which the behavior of two systems is coordinated, i.e., how much dynamically-relevant information they share. Because of its centrality to our method, we briefly recapitulate the argument of that paper.

The starting point is the strong notion of “state” employed in physics and dynamical systems theory: the state of the system is a variable which determines the distribution of all present and future observables. In inferential terms, the state is a minimal sufficient statistic for predicting future observations [29], and can be formally constructed as measure-valued process giving the distribution of future events conditional on the history of the process. As a consequence, the state always evolves according to a homogeneous Markov process [30, 29].

In a dynamical network, each node i has an associated time-series of observations $X_i(t)$. This is in turn generated by a Markovian state process, $S_i(t)$, which forms its optimal nonlinear predictor. For any two nodes i and j , the informational coherence is

$$IC_{ij} \equiv \frac{I[S_i; S_j]}{\min H[S_i], H[S_j]} \quad (1)$$

where $I[S_i; S_j]$ is the mutual information shared by S_i and S_j , and $H[S_i]$ is the self-information (Shannon entropy) of S_i . Since $I[S_i; S_j] \leq \min H[S_i], H[S_j]$, this is a symmetric quantity, normalized to lie between 0 and 1 inclusive. The construction of the predictive states ensures that $S_i(t)$ encapsulates all information in the past of $X_i(t)$ which is relevant to its future, so a positive value for $I[S_i; S_j]$ means that $S_j(t)$ contains information about the future of $X_i(t)$. That is, a positive value of $I[S_i; S_j]$ is equivalent to the sharing of *dynamically relevant* information between the nodes, manifesting itself as coordinated behavior on the part of nodes i and j .

Clearly, a crucial step in calculating informational coherence is going from the observational time series $X_i(t)$ to the predictive state series $S_i(t)$. In certain

cases with completely specified probability models, this can be done analytically [29, 31]. In general, however, we are forced to reconstruct the appropriate state-space structure from the time series itself. State reconstruction for deterministic systems is based on the Takens embedding theorem, and is now routine [32]. However, biological and social systems are hardly ever deterministic at experimentally-accessible levels of resolution, so we need a stochastic state reconstruction algorithm. Several exist; we use the CSSR algorithm introduced in [33], since, so far as we know, it is currently the only stochastic state reconstruction algorithm which has been proved statistically consistent (for conditionally stationary discrete sequences). We briefly describe CSSR in Appendix A.

Informational coherence is not, of course, the only possible way of measuring behavioral coordination, or functional connectivity. However, it has a number of advantages over rival measures [25]. Unlike measures of strict synchronization, which insist on units doing exactly the same thing at exactly the same time, it accommodates phase lags, phase locking, chaotic synchronization, etc., in a straightforward and uniform manner. Unlike cross-covariance, or the related spectral coherence, it easily handles nonlinear dependencies, and does not require the choice of a particular lag (or frequency, for spectral coherence), because the predictive states summarize the entire relevant portion of the history. Generalized synchrony measures [34] can handle nonlinear relationships among states, but inappropriately assume determinism. Finally, mutual information among the observables, $I[X_i; X_j]$, can handle nonlinear, stochastic dependencies, but suffers, especially in neural systems, because what we really want to detect are coordinated *patterns* of behavior, rather than coordinated instantaneous actions. Because each predictive state corresponds to a unique statistical pattern of behavior, mutual information among these states is the most natural way to capture functional connectivity.

2.2 The Reichardt-Bornholdt Community Discovery Algorithm

The Reichardt-Bornholdt [26, 27] community discovery algorithm finds groups of nodes that are densely coupled to one another, but only weakly coupled to the rest of the network, by establishing a (fictitious) spin system on the network, with a Hamiltonian with precisely the desired properties, and then minimizing the Hamiltonian through simulated annealing. More concretely, every node i is assigned a “spin” σ_i , which is a discrete variable taking an integer value from 1 to a user-defined q . A “community” or “module” will consist of all the nodes with a common spin value. The spin Hamiltonian combines a ferromagnetic term, which favors linked nodes taking the same spin (i.e., being in the same community), and an anti-ferromagnetic term, which favors non-linked nodes taking different spins (i.e., being in different community). Both interactions are of the Potts model type, i.e., they are invariant under permutations of the integers labeling the clusters. After some algebraic manipulation [27], one arrives at the Hamiltonian

$$\mathcal{H}(\sigma) = - \sum_{i \neq j} (A_{ij} - \gamma p_{ij}) \delta(\sigma_i, \sigma_j) \quad (2)$$

where A_{ij} is the adjacency matrix, $\delta(\cdot, \cdot)$ is the Kronecker delta function, p_{ij} is a matrix of non-negative constants giving the relative weights of different possible links, and γ gives the relative contribution of link absence to link presence. The choice of p_{ij} is actually fairly unconstrained, but previous experience with community discovery suggests that very good results are obtained by optimizing the Newman modularity Q [35]

$$Q(\sigma) = \frac{1}{2M} \sum_{i,j} \left(A_{ij} - \frac{k_i k_j}{2M} \right) \delta(\sigma_i, \sigma_j) \quad (3)$$

where k_i is the degree of node i , and $2M = \sum_i k_i$ the total number of links. Essentially, Newman’s Q counts the number of edges within communities, minus the number which would be expected in a randomized graph where each node preserved its actual degree [9], and σ_i were IID uniform. Setting $p_{ij} = k_i k_j / 2M$ and $\gamma = 1$, we see that $\mathcal{H}(\sigma)$ and $-Q(\sigma)$ differ only by a term (the diagonal part of the sum for Q) which does not depend on the assignment of nodes to communities. Thus, minimizing $\mathcal{H}(\sigma)$ is the same as maximizing the modularity. Varying γ , in this scheme, effectively controls the trade-off between having many small communities and a few large ones [27], and makes it possible to discover a hierarchical community structure, which will be the subject of future work.

While this procedure was originally developed for the case where A_{ij} is a 0-1 adjacency matrix, it also works perfectly well when links take on (positive) real-valued strengths. In particular, using $A_{ij} = IC_{ij}$, we can still maximize the modularity, taking the “degree” of node i to be $k_i = \sum_j IC_{ij}$ [27]. The interpretation of the modularity is now the difference between the strength of intra-community links, and a randomized model where each node shares its link strength indifferently with members of its own and other communities.

2.3 Summary of the Method

Let us briefly summarize the method for discovering functional communities. We begin with a network, consisting of N nodes. For each node, we have a discrete-value, discrete-time (“symbolic”) time series, $\{x_i(t)\}$, recorded simultaneously over all nodes. The CSSR algorithm is applied to each node’s series separately, producing a set of predictive states for that node, and a time series of those states, $\{s_i(t)\}$. We then calculate the complete set of pairwise informational coherence values, $\{IC_{ij}\}$, using Eq. 1. This matrix is fed into the Reichardt-Bornholdt procedure, with $A_{ij} = IC_{ij}$, which finds an assignment of spins to nodes, $\{\sigma_i\}$, minimizing the Hamiltonian in Eq. 2. The functional communities of the dynamical network consist of groups of nodes with common spin values. Within each community, the average pairwise coherence of the nodes is strictly greater than would be expected from a randomizing null model (as described in the previous paragraph). Furthermore, between any two communities, the average pairwise coherence of their nodes is strictly less than expected from randomization [27].

3 Test on a Model System of Known Structure: Collective Oscillations in the Hippocampus

We use simulated data as a test case, to validate the general idea of our method, because it allows us to work with a substantial network where we nonetheless have a strong idea of what appropriate results should be. Because of our ultimate concern with the functional re-organization of the brain, we employed a large, biophysically-detailed neuronal network model, with over 1000 simulated neurons.

The model, taken from [28], was originally designed to study episodes of gamma (30–80Hz) and beta (12–30Hz) oscillations in the mammalian nervous system, which often occur successively with a spontaneous transition between them. More concretely, the rhythms studied were those displayed by *in vitro* hippocampal (CA1) slice preparations and by *in vivo* neocortical EEGs.

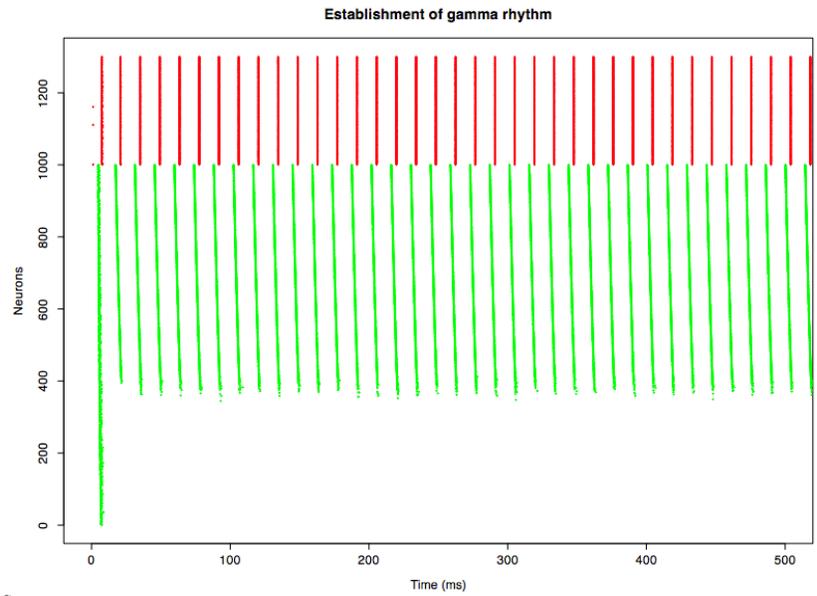
The model contains two neuron populations: excitatory (AMPA) pyramidal neurons and inhibitory (GABA_A) interneurons, defined by conductance-based Hodgkin-Huxley-style equations. Simulations were carried out in a network of 1000 pyramidal cells and 300 interneurons. Each cell was modeled as a one-compartment neuron with all-to-all coupling, endowed with the basic sodium and potassium spiking currents, an external applied current, and some Gaussian input noise. The anatomical, synaptic connections were organized into blocks, as shown in Fig. 2.

The first 10 seconds of the simulation correspond to the gamma rhythm, in which only a group of neurons is made to spike via a linearly increasing applied current. The beta rhythm (subsequent 10 seconds) is obtained by activating pyramidal-pyramidal recurrent connections (potentiated by Hebbian pre-processing as a result of synchrony during the gamma rhythm) and a slow outward after-hyper-polarization (AHP) current (the M-current), suppressed during gamma due to the metabotropic activation used in the generation of the rhythm. During the beta rhythm, pyramidal cells, silent during gamma rhythm, fire on a subset of interneurons cycles (Fig. 1).

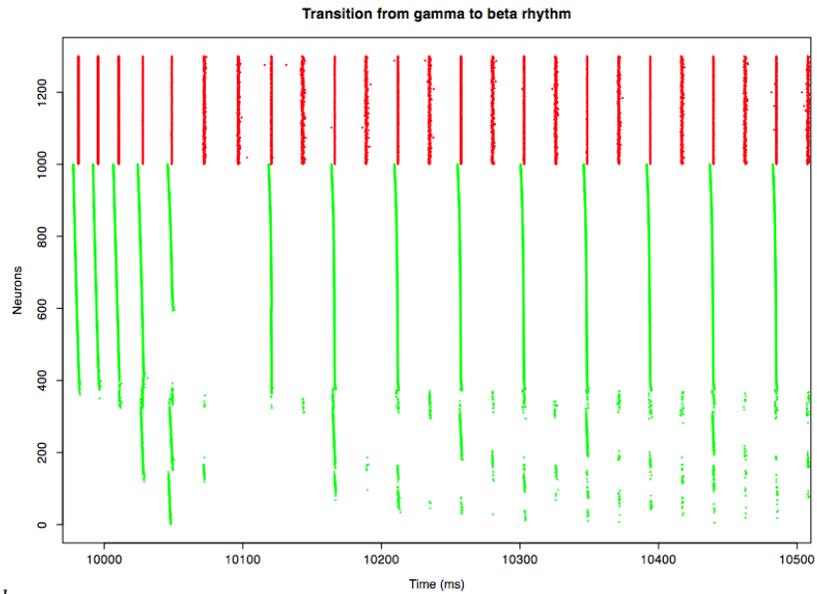
4 Results on the Model

A simple heat-map display of the informational coherence (Fig. 3) shows little structure among the active neurons in either regime. However, visual inspection of the rastergrams (Fig. 1) leads us to suspect the presence of two very large functional communities: one, centered on the inhibitory interneurons and the excitatory pyramidal neurons most tightly coupled to them, and another of the more peripheral excitatory neurons. During the switch from the gamma to the beta rhythm, we expect these groups to re-organize.

These expectations are abundantly fulfilled (Fig. 4). We identified communities by running the Reichardt-Bornholdt algorithm with the maximum number of communities (spin states) set to 25, the modularity Hamiltonian, and $\gamma = 1$. (Results were basically unchanged at 40 or 100 spin values.) In both regimes,



a



b

Fig. 1. Rastergrams of neuronal spike-times in the network. Excitatory, pyramidal neurons (numbers 1 to 1000) are green, inhibitory interneurons (numbers 1001 to 1300) are red. During the first 10 seconds (*a*), the current connections among the pyramidal cells are suppressed and a gamma rhythm emerges (left). At $t = 10$ s, those connections become active, leading to a beta rhythm (*b*, right).

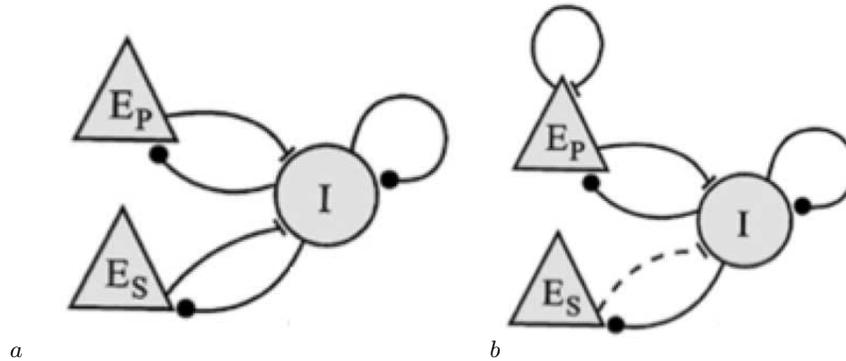


Fig. 2. Schematic depiction of the anatomical network. Here nodes represent populations of cells: excitatory pyramidal neurons (triangles labeled E) or inhibitory interneurons (large circle labeled I). Excitatory connections terminate in bars, inhibitory connections in filled circles. During the gamma rhythm (*a*), the pyramidal neurons are coupled to each other only indirectly, via the interneurons, and dynamical effects separate the pyramidal population into participating (E_P) and suppressed (E_S) subpopulations. During the beta rhythm (*b*), direct connections among the E_P neurons, built up, but not activated, by Hebbian learning under the gamma rhythm are turned on, and the connection from the E_S neurons to the interneurons are weakened by the same Hebbian process (dashed line).

there are two overwhelmingly large communities, containing almost all of the neurons which actually fired, and a handful of single-neuron communities. The significant change, visible in the figure, is in the organization of these communities.

During the gamma rhythm, the 300 interneurons form the core of the larger of these two communities, which also contains 199 pyramidal neurons. Another 430 pyramidal neurons belong to a second community. A final 5 pyramidal cells are in single-neuron communities; the rest do not fire at all. A hierarchical analysis (not shown) has the two large communities merging into a single super-community. The regular alternation of the two communities among the pyramidal neurons, evident in Fig. 4*a*, is due to the fact that the external current driving the pyramidal neurons is not spatially uniform.

With the switch to the beta rhythm, the communities grow and re-organize. The community centered on the interneurons expands, to 733 neurons, largely by incorporating many low-index pyramidal neurons which had formerly been silent, and are now somewhat erratically synchronized, into its periphery. Interestingly, many of the latter are only weakly coherent with any one interneuron (as can be seen by comparing Figs. 3*b* and 4*b*). What is decisive is rather their stronger over-all pattern of coordination with the interneurons, shown by sharing a common (approximate) firing period, which is half that of the high-index pyramidal cells (Fig. 1*b*). Similarly, the other large community, consisting exclusively

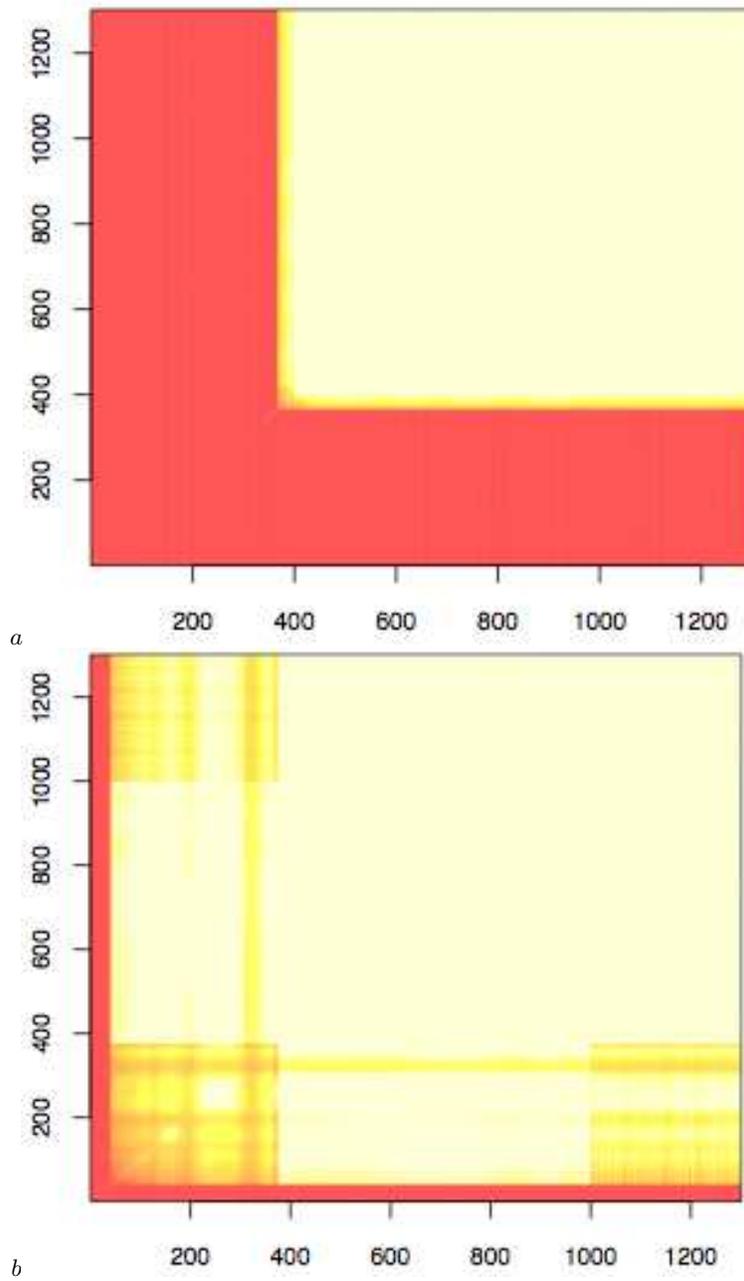


Fig. 3. Heat-maps of coordination across neurons in the network, measured by informational coherence. Colors run from red (no coordination) through yellow to pale cream (maximum).

of pyramidal neurons, also grows (to 518 members), again by expanding into the low-index part of the network; there is also considerable exchange of high-index pyramidal cells between the two communities. Finally, nine low-index neurons, which fire only sporadically, belong in clusters of one or two cells.

5 Discussion and Conclusion

5.1 Limitations and Related Approaches

Our method is distinguished from earlier work on functional connectivity primarily by our strong notion of functional community or module, and secondarily by our measure of functional connectivity. Previous approaches to functional connectivity (reviewed in [10, 11]) either have no concept of functional cluster, or use simple agglomerative clustering [4]; their clusters are just groups of nodes with pairwise-similar behavior. We avoid agglomerative clustering for the same reason it is no longer used to find anatomical communities: it is insensitive to the global pattern of connectivity, and fails to divide the network into coherent components. Recall (Sect. 2.3) that every functional community we find has more intra-cluster information sharing than is expected by chance, and less inter-cluster information sharing. This is a plausible formalization of the intuitive notion of “module”, but agglomeration will not, generally, deliver it.

As for using informational coherence to measure functional connectivity, we discussed its advantages over other measures in Sect. 2.1 above, and at more length in [25]. Previous work on functional connectivity has mostly used surface features to gauge connectivity, such as mutual information between observables. (Some of the literature on clustering general time series, e.g. [36, 37, 38], uses hidden Markov models to extract latent features, but in a mixture-model framework very different from our approach.) The strength of informational coherence is that it is a domain-neutral measure of nonlinear, stochastic coordination; its weakness is that it requires us to know the temporal sequence of predictive states of all nodes in the network.

This need to know the predictive states of each node is the major limitation of our method. For some mathematical models, these states are analytically calculable, but in most cases they must be learned from discrete-value, discrete-time (“symbolic”) time series. Those series must be fairly long; exactly how long is an on-going topic of investigation⁴, but, empirically, good results are rare with less than a few thousand time steps. Similarly, reliable estimates of the mutual information and informational coherence also require long time series.

Predictive states can be mathematically defined for continuous-value, continuous-time systems [30], but all current algorithms for discovering them, not just CSSR, require symbolic time series. (Devising a state-reconstruction procedure for continuous systems is another topic of ongoing research.) Spike trains, like e-mail networks [22], are naturally discrete, so this is not an issue for them, but in most

⁴ CSSR converges on the true predictive states (see the appendix), but the rate of convergence is not yet known.

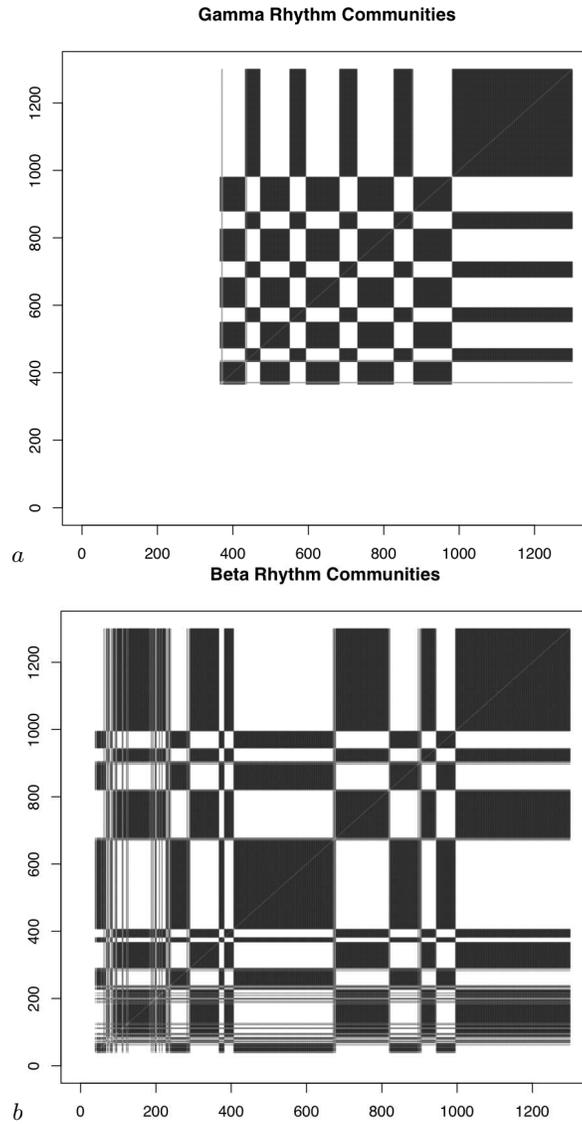


Fig. 4. Division of the network into functional communities. Black points denote pairs of nodes which are both members of a given community. During the gamma rhythm (*a*) the interneurons (numbered 1001 to 1300) form the core of a single community, along with some of the active pyramidal neurons; because of the spatially modulated input received by the latter, however, some of them belong to another community. During beta rhythm (*b*), the communities re-organize, and in particular formerly inactive pyramidal neurons are recruited into the community centered on the interneurons, as suggested by the rastergrams.

other cases we need to find a good symbolic partition first, which is non-trivial [39]. The need for long symbolic time series may be especially difficult to meet with social networks.

5.2 Directions for Future Work

Our results on the model of hippocampal rhythms, described in the previous section, are quite promising: our algorithm discovers functional communities whose organization and properties make sense, given the underlying micro-dynamics of the model. This suggests that it is worthwhile to apply the method to systems where we lack good background knowledge of the functional modules. Without pre-judging the results of those investigations, however, we would like to highlight some issues for future work.

1. Our method needs the full matrix of informational coherences, which is an $O(N^2)$ computation for a network of size N . If we are interested in the organization of only part of the network, can we avoid this by defining a local community structure, as was done for anatomical connectivity by [40]? Alternatively, if we know the anatomical connectivity, can we restrict ourselves to calculating the informational coherence between nodes which are anatomically tied? Doing so with our model system led to basically the same results (not shown), which is promising; but in many real-world systems the anatomical network is itself uncertain.

2. The modularity Hamiltonian of Sect. 2.2 measures how much information each node shares with other members of its community on a pairwise basis. However, some of this information could be redundant across pairs. It might be better, then, to replace the sum over pairs with a higher-order coherence. The necessary higher-order mutual informations are easily defined [10, 41, 42], but the number of measurements needed to estimate them from data grows exponentially with the number of nodes. However, it may be possible to approximate them using the same Chow-Liu bounds employed by [25] to estimate the global coherence.

3. It would be good if our algorithm did not simply report a community structure, but also assessed the likelihood of the same degree of modularity arising through chance, i.e., a significance level. For anatomical communities, Guimera et al. [43] exploit the spin-system analogy to show that random graph processes without community structure will nonetheless often produce networks with non-zero modularity, and (in effect) calculate the sampling distribution of Newman's Q using both Erdős-Rényi and scale-free networks as null models. (See however [44] for corrections to their calculations.) To do something like this with our algorithm, we would need a null model of functional communities. The natural null model of functional connectivity is simply for the dynamics at all nodes to be independent, and (because the states are Markovian) it is easy to simulate from this null model and then bootstrap p -values. We do not yet, however, have a class of dynamical models where there nodes share information, but do so in a completely distributed, a-modular way.

4. A variant of the predictive-state analysis that underlies informational coherence is able to identify coherent structures produced by spatiotemporal dynamics [45]. Moreover, these techniques can be adapted to network dynamics, if the anatomical connections are known. This raises numerous questions. Are functional communities also coherent structures? Are coherent structures in networks [46] necessarily functional communities? Can the higher-order interactions of coherent structures in regular spatial systems be ported to networks, and, if so, could functional re-organization be described as a dynamical process at this level?

5.3 Conclusion

Network dynamical systems have both anatomical connections, due to persistent physical couplings, and functional ones, due to coordinated behavior. These are related, but logically distinct. There are now many methods for using a network’s anatomical connectivity to decompose it into highly modular communities, and some understanding of these methods’ statistical and statistical-mechanical properties. The parallel problem, of using the pattern of functional connectivity to find functional communities, has scarcely been explored. It is in many ways a harder problem, because measuring functional connectivity is harder, and because the community organization is itself variable, and this variation is often more interesting than the value at any one time.

In this paper, we have introduced a method of discovering functional modules in stochastic dynamical networks. We use informational coherence to measure functional connectivity, and combine this with a modification of the Potts-model community-detection procedure. Our method gives good results on a biophysical model of hippocampal rhythms. It divides the network into two functional communities, one of them based on the inhibitory interneurons, the other consisting exclusively of excitatory pyramidal cells. The two communities change in relative size and re-organize during the switch from gamma to beta rhythm, in ways which make sense in light of the underlying model dynamics. While there are theoretical issues to explore, our success on a non-trivial simulated network leads us to hope that we have found a general method for discovering functional communities in dynamic networks.

Acknowledgments

Thanks to L. A. N. Amaral, S. Bornholdt, A. Clauset, R. Haslinger, C. Moore and M. E. J. Newman for discussing community discovery and/or network dynamics; to J. Reichardt for valuable assistance with the implementation of his algorithm and general suggestions; and to the editors for their patience.

A The CSSR Algorithm

This appendix briefly describes the CSSR algorithm we use to reconstruct the effective internal states of each node in the network. For details, see [33]; for an

open-source C++ implementation, see <http://bactra.org/CSSR/>. For recent applications of the algorithm to problems in crystallography, anomaly detection and natural language processing, see [47, 48, 49, 50, 51].

We wish to predict a dynamical system or stochastic process $\{X_t\}$. By X_s^t we will denote the whole trajectory of the process from time s to time t , inclusive, by X_t^+ the whole “past” or “history” of the process through time t , and by X_t^- its “future”, its trajectory at times strictly greater than t . The “state” of $\{X_t\}$ at time t is a variable, S_t , which fixes the distribution of all present or future observations, i.e., the distribution of $X^+(t)$ [29, 31]. As such, the state is a minimal sufficient statistic for predicting the future of the process. Sufficiency is equivalent to the requirement that $I[X_t^+; X_t^-] = I[X_t^+; S_t]$, where $I[\cdot; \cdot]$ is the mutual information [52]. In general, $S_t = \epsilon(X_t^-)$, for some measurable functional $\epsilon(\cdot)$ of the whole past history of the process up to and including time t . If $\{X_t\}$ is Markovian, then ϵ is a function only of X_t , but in general the state will incorporate some history or memory effects. Each state, i.e., possible value of ϵ , corresponds to a predictive distribution over future events, and equally to an equivalence class of histories, all of which lead to that conditional distribution over future events. State-reconstruction algorithms use sample paths of the process to find approximations $\hat{\epsilon}$ to the true minimal sufficient statistic ϵ , and ideally the approximations converge, at least in probability. The CSSR algorithm [33] does so, for discrete-valued, discrete-time, conditionally-stationary processes.

CSSR is based on the following result about predictive sufficiency [29, pp. 842–843]. Suppose that ϵ is next-step sufficient, i.e., $I[X_{t+1}; X_t^-] = I[X_{t+1}; \epsilon(X_t^-)]$, and that it can be updated recursively: for some measurable function T , $\epsilon(X_{t+1}^-) = T(\epsilon(X_t^-), X_{t+1})$. Then ϵ is predictively sufficient for the whole future of the process — intuitively, the recursive updating lets us chain together accurate next-step predictions to go as far into the future as we like. CSSR approximates ϵ by treating it as a partition, or set of equivalence classes, over histories, and finding the coarsest partition which meets both of the conditions of this theorem. Computationally, CSSR represents states as sets of suffixes, so a history belongs to a state (equivalence class) if it terminates in one of the suffixes in that state’s representation. That is, a history, x_t^- , will belong to the class C , $x_t^- \in C$, if $x_{t-|c|+1}^t = c$, for some suffix c assigned to C , where $|c|$ is the length of the suffix.⁵

In the first stage, CSSR tries to find a partition of histories which is sufficient for next-step prediction. It begins with the trivial partition, in which all histories belong to the same equivalence class, defined by the null suffix (corresponding to an IID process), and then successively tests whether longer and longer suffixes give rise to the same conditional distribution for the next observation which differ significantly from the class to which they currently belong. That is, for each class C , suffix c in that class, and possible observable value a , it tests whether $\Pr(X_{t+1}|X_t^- \in C)$ differs from $\Pr(X_{t+1}|X_{t-|c|+1}^t = c, X_{t-|c|} = a)$. (We use stan-

⁵ The algorithm ensures that there are never overlapping suffixes in distinct states.

dard tests for discrepancy between sampled distributions.) If an extended, child suffix (ac) does not match its current classes, the parent suffix (c) is deleted from its class (C), and CSSR checks whether the child matches any existing class; if so it is re-assigned to the closest one, and the partition is modified accordingly. Only if a suffix’s conditional distribution ($\Pr(X_{t+1}|X_{t-L}^t = ac)$) differs significantly from all existing classes does it get its own new cell in the partition.

The result of this stage is a partition of histories (i.e., a statistic) which is close to being next-step sufficient, the sense of “close” depending on the significance test. In the second stage, CSSR iteratively refines this partition until it can be recursively updated. This can always be done, though it is potentially the most time-consuming part of the algorithm⁶. The output of CSSR, then, is a set of states which make good next-step predictions and can be updated recursively, and a statistic $\hat{\epsilon}$ mapping histories to these states.

If the true number of predictive states is finite, and some mild technical assumptions hold [33], a large deviations argument shows that $\Pr(\hat{\epsilon} \neq \epsilon) \rightarrow 0$ as the sample size $n \rightarrow \infty$. That is, CSSR will converge on the minimal sufficient statistic for the data-generating process, even though it lacks an explicit minimization step. Furthermore, once the right statistic has been discovered, the expected L_1 (total variation) distance between the actual predictive distribution, $\Pr(X_t^+|\epsilon(X_t^-))$ and that forecast by the reconstructed states, $\Pr(X_t^-|\hat{\epsilon}(X_t^-))$, goes to zero with rate $O(n^{-1/2})$, which is the same rate as for IID data. The time complexity of the algorithm is at worst $O(n) + O(k^{2L+1})$, where k is the number of discrete values possible for X_t , and L is the maximum length of suffixes considered in the reconstruction. Empirically, average-case time complexity is much better than this.

References

- [1] Simmel, G.: Conflict and the Web of Group Affiliation. The Free Press, Glencoe, Illinois (1955) Translated by Kurt H. Wolff and Reinhard Bendix, with a foreword by Everett C. Hughes.
- [2] Scott, J.: Social Network Analysis: A Handbook. Sage Publications, Thousand Oaks, California (2000)
- [3] Moody, J., White, D.R.: Social cohesion and embeddedness: A hierarchical conception of social groups. *American Sociological Review* **68** (2003) 103–127
- [4] Hand, D., Mannila, H., Smyth, P.: Principles of Data Mining. MIT Press, Cambridge, Massachusetts (2001)
- [5] Simon, H.A.: The architecture of complexity: Hierarchic systems. *Proceedings of the American Philosophical Society* **106** (1962) 467–482 Reprinted as chapter 8 of [53].
- [6] Alexander, C.: Notes on the Synthesis of Form. Harvard University Press, Cambridge, Massachusetts (1964)

⁶ In terms of automata theory, recursive updating corresponds to being a “deterministic” automaton, and non-deterministic automata always have deterministic equivalents.

- [7] Šijlak, D.S.: Decentralized Control of Complex Systems. Academic Press, Boston (1990)
- [8] Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. *Physical Review E* **69** (2003) 026113
- [9] Newman, M.E.J.: Finding community structure in networks using the eigenvectors of matrices. E-print, arxiv.org, physics/0605087 (2006)
- [10] Sporns, O., Tononi, G., Edelman, G.M.: Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex* **10** (2000) 127–141
- [11] Friston, K.: Beyond phrenology: What can neuroimaging tell us about distributed circuitry? *Annual Review of Neuroscience* **25** (2002) 221–250
- [12] Kaneko, K.: Clustering, coding, switching, hierarchical ordering, and control in a network of chaotic elements. *Physica D* **41** (1990) 137–172
- [13] Chazottes, J.R., Fernandez, B., eds.: Dynamics of Coupled Map Lattices and of Related Spatially Extended Systems. Volume 671 of Lecture Notes in Physics. Springer-Verlag, Berlin (2005)
- [14] Delgado, J., Solé, R.V.: Characterizing turbulence in globally coupled maps with stochastic finite automata. *Physics Letters A* **314** (2000) 314–319
- [15] Selverston, A.I., Moulins, M., eds.: The Crustacean Stomatogastric System: A Model for the Study of Central Nervous Systems. Springer-Verlag, Berlin (1987)
- [16] Luria, A.R.: The Working Brain: An Introduction to Neuropsychology. Basic Books, New York (1973)
- [17] Segel, L.A., Cohen, I.R., eds.: Design Principles for the Immune System and Other Distributed Autonomous Systems. Oxford University Press, Oxford (2001)
- [18] Chisholm, D.: Coordination without Hierarchy: Informal Structures in Multiorganizational Systems. University of California Press, Berkeley (1989)
- [19] Luhmann, N.: Social Systems. Stanford University Press, Stanford (1984/1995) Translated by John Bednarz, Jr., with Dirk Baecker. Foreword by Eva M. Knodt. First published as *Soziale Systeme: Grundriss einer allgemeinen Theorie*, Frankfurt am Main: Suhrkamp-Verlag.
- [20] Lindblom, C.E.: The Intelligence of Democracy: Decision Making through Mutual Adjustment. Free Press, New York (1965)
- [21] Young, H.P.: Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, Princeton (1998)
- [22] Eckmann, J.P., Moses, E., Sergi, D.: Entropy of dialogues creates coherent structures in e-mail traffic. *Proceedings of the National Academy of Sciences (USA)* **101** (2004) 14333–14337
- [23] Coser, R.L.: In Defense of Modernity: Role Complexity and Individual Autonomy. Stanford University Press, Stanford, California (1991)
- [24] Gellner, E.: Plough, Sword and Book: The Structure of Human History. University of Chicago Press, Chicago (1988)
- [25] Klinkner, K.L., Shalizi, C.R., Camperi, M.F.: Measuring shared information and coordinated activity in neuronal networks. In Weiss, Y., Schölkopf, B., Platt, J.C., eds.: *Advances in Neural Information Processing Systems 18 (NIPS 2005)*, Cambridge, Massachusetts, MIT Press (2006) 667–674
- [26] Reichardt, J., Bornholdt, S.: Detecting fuzzy community structures in complex networks with a Potts model. *Physical Review Letters* **93** (2004) 218701
- [27] Reichardt, J., Bornholdt, S.: Statistical mechanics of community detection. *Physical Review E* **74** (2006) 016110

- [28] Olufsen, M.S., Whittington, M.A., Camperi, M., Kopell, N.: New roles for the gamma rhythm: Population tuning and preprocessing for the beta rhythm. *Journal of Computational Neuroscience* **14** (2003) 33–54
- [29] Shalizi, C.R., Crutchfield, J.P.: Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics* **104** (2001) 817–879
- [30] Knight, F.B.: A predictive view of continuous time processes. *Annals of Probability* **3** (1975) 573–596
- [31] Littman, M.L., Sutton, R.S., Singh, S.: Predictive representations of state. In Dietterich, T.G., Becker, S., Ghahramani, Z., eds.: *Advances in Neural Information Processing Systems 14*, Cambridge, Massachusetts, MIT Press (2002) 1555–1561
- [32] Kantz, H., Schreiber, T.: *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge, England (1997)
- [33] Shalizi, C.R., Klinkner, K.L.: Blind construction of optimal nonlinear recursive predictors for discrete sequences. In Chickering, M., Halpern, J., eds.: *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference*, Arlington, Virginia, AUAI Press (2004) 504–511
- [34] Quian Quiroga, R., Kraskov, A., Kreuz, T., Grassberger, P.: Performance of synchronization measures in real data: case studies on electroencephalographic signals. *Physical Review E* **65** (2002) 041903
- [35] Newman, M.E.J.: Mixing patterns in networks. *Physical Review E* **67** (2003) 026126
- [36] Smyth, P.: Clustering sequences using hidden Markov models. In Mozer, M.C., Jordan, M.I., Petsche, T., eds.: *Advances in Neural Information Processing 9*, Cambridge, Massachusetts, MIT Press (1997) 648–654
- [37] Oates, T., Firiou, L., Cohen, P.R.: Clustering time series with hidden Markov models and dynamic time warping. In Giles, C.L., Sun, R., eds.: *Proceedings of the IJCAI-99 Workshop on Neural, Symbolic and Reinforcement Learning Methods for Sequence Learning*. (1999) 17–21
- [38] Cadez, I.V., Gaffney, S., Smyth, P.: A general probabilistic framework for clustering individuals and objects. In Ramakrishnan, R., Stolfo, S., Bayardo, R., Parsa, I., eds.: *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, New York, ACM Press (2000) 140–149
- [39] Hirata, Y., Judd, K., Kilminster, D.: Estimating a generating partition from observed time series: Symbolic shadowing. *Physical Review E* **70** (2004) 016215
- [40] Clauset, A.: Finding local community structure in networks. *Physical Review E* **72** (2005) 026132
- [41] Amari, S.i.: Information geometry on hierarchy of probability distributions. *IEEE Transactions on Information Theory* **47** (2001) 1701–1711
- [42] Schneidman, E., Still, S., Berry, M.J., Bialek, W.: Network information and connected correlations. *Physical Review Letters* **91** (2003) 238701
- [43] Guimera, R., Sales-Pardo, M., Amaral, L.A.N.: Modularity from fluctuations in random graphs and complex networks. *Physical Review E* **70** (2004) 025101
- [44] Reichardt, J., Bornholdt, S.: When are networks truly modular? E-print, arxiv.org, cond-mat/0606220 (2006)
- [45] Shalizi, C.R., Haslinger, R., Rouquier, J.B., Klinkner, K.L., Moore, C.: Automatic filters for the detection of coherent structure in spatiotemporal systems. *Physical Review E* **73** (2006) 036104
- [46] Moreira, A.A., Mathur, A., Diermeier, D., Amaral, L.A.N.: Efficient system-wide coordination in noisy environments. *Proceedings of the National Academy of Sciences (USA)* **101** (2004) 12085–12090

- [47] Varn, D.P., Crutchfield, J.P.: From finite to infinite range order via annealing: The causal architecture of deformation faulting in annealed close-packed crystals. *Physics Letters A* **324** (2004) 299–307
- [48] Ray, A.: Symbolic dynamic analysis of complex systems for anomaly detection. *Signal Processing* **84** (2004) 1115–1130
- [49] Padró, M., Padró, L.: Applying a finite automata acquisition algorithm to named entity recognition. In: *Proceedings of 5th International Workshop on Finite-State Methods and Natural Language Processing (FSMNLP'05)*. (2005)
- [50] Padró, M., Padró, L.: A named entity recognition system based on a finite automata acquisition algorithm. *Procesamiento del Lenguaje Natural* **35** (2005) 319–326
- [51] Padró, M., Padró, L.: Approaching sequential NLP tasks with an automata acquisition algorithm. In: *Proceedings of International Conference on Recent Advances in NLP (RANLP'05)*. (2005)
- [52] Kullback, S.: *Information Theory and Statistics*. 2nd edn. Dover Books, New York (1968)
- [53] Simon, H.A.: *The Sciences of the Artificial*. Third edn. MIT Press, Cambridge, Massachusetts (1996) First edition 1969.