



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Modularity and the spread of perturbations in complex dynamical systems

Artemy Kolchinsky, Alexander J. Gates, and Luis M. Rocha

Phys. Rev. E **92**, 060801 — Published 23 December 2015

DOI: [10.1103/PhysRevE.92.060801](https://doi.org/10.1103/PhysRevE.92.060801)

Modularity and the Spread of Perturbations in Complex Dynamical Systems

Artemy Kolchinsky¹, Alexander J. Gates^{1,2}, and Luis M. Rocha^{1,2,3}

¹*School of Informatics and Computing, Indiana University, Bloomington, Indiana 47408, USA*

²*Program in Cognitive Science, Indiana University, Bloomington, Indiana 47408, USA and*

³*Instituto Gulbenkian de Ciência, Oeiras, Portugal*

We propose a method to decompose dynamical systems based on the idea that modules constrain the spread of perturbations. We find partitions of system variables that maximize ‘perturbation modularity’, defined as the auto-covariance of coarse-grained perturbed trajectories. The measure effectively separates the fast intra-modular from the slow inter-modular dynamics of perturbation spreading (in this respect, it is a generalization of the ‘Markov stability’ method of community detection). Our approach captures variation of modular organization across different system states, time scales, and in response to different kinds of perturbations – aspects of modularity which are all relevant to real-world dynamical systems. It offers a principled alternative to detecting communities in networks of statistical dependencies between system variables (e.g. ‘relevance networks’ or ‘functional networks’). Using coupled logistic maps, we demonstrate that the method uncovers hierarchical modular organization planted in a system’s coupling matrix. Additionally, in homogeneously-coupled map lattices, it identifies the presence of self-organized modularity that depends on the initial state, dynamical parameters, and type of perturbations. Our approach offers a powerful tool for exploring the modular organization of complex dynamical systems.

Many complex systems are *modular*, in that their components are organized in tightly-integrated subsystems that are weakly-coupled to one another. Modularity has been argued to play many important roles, including increasing robustness [1–3], evolvability [1, 4], and functional differentiation [5, 6]. Thus, there is great interest in measures of modularity and methods for decomposing complex systems into weakly-coupled modules.

This problem is here considered in the domain of multivariate dynamics, a common formalism for modeling complex physical, biological, neural and social systems. We propose a method of identifying dynamical modules motivated by the intuition that, in a modular system, the spread of perturbations is characterized by two time scales: fast spreading within modules and slow spreading between modules [1, 7]. In our treatment, the spreading process is coarse-grained relative to a *partition* (a decomposition of system variables into disjoint subsystems) by measuring the magnitude of the perturbation’s effect within each subsystem over time. If a partition reflects underlying modular structure, initially perturbed subsystems remain affected as dynamics unfold, while initially unperturbed subsystems remain largely unaffected. In this case, the partition’s coarse-graining will capture the slow component of perturbation spreading dynamics, an effect quantified using a quality function called *perturbation modularity*. Our perturbation-based approach is related to many existing techniques for analyzing multivariate dynamics, including Lyapunov-exponent based methods [8–10] and impulse response analysis [11].

As will be elaborated below, our methodology can identify the dependence of optimal decompositions on initial states, time scales, and kinds of perturbations applied. These factors are all important aspects of modular organization in real-world dynamical systems. Dependence on the initial condition reflects that dynamical systems can exhibit different modular organizations in

different regions of their state-space; for example, distributed regions of the brain can couple into modular assemblies via oscillatory synchronization, with the same region participating in different assemblies depending on brain state [12, 13]. The choice of time scale affects optimal decompositions by determining the separation between intra-modular and inter-modular perturbation spreading; in real-world complex systems, longer time scales have often been argued to correspond to larger-scale modules [1, 14–17]. Finally, the dependence on the kinds of perturbations reflects that a dynamical system may be robust to some perturbations but highly-sensitive to others [18]; for example, in biological double-knockout experiments, cellular responses to the simultaneous deactivation of two genes can differ dramatically from responses to the individual deactivation of either gene [19].

Our approach starts from a pre-specified dynamical system and thus differs fundamentally from existing treatments of modularity based on network representations of a system. Such methods are usually unable to capture the variation of modular organization across state-space or time scale, as well as other important dynamical aspects of modularity.

For instance, one standard approach applies graph-based *community detection* [20] to the structural network underlying a dynamical process (e.g. the social network over which an epidemic spreads). This treatment ignores the fact that the same structural network can support many different dynamical processes (for example, ‘complex contagion’ epidemics proceed differently from ‘simple contagion’ epidemics [21]). In contrast, our methodology is by definition sensitive to dynamical differences.

Another class of methods applies community structure to network representations of dynamics, defined either in terms of causal interactions or statistical dependencies between variables (e.g. *relevance networks* in systems biology [22] and *functional networks* in neuro-

science [23]). Unfortunately, constructing such networks involves a conversion of the dynamical system (defined in terms of transitions between multidimensional states) into a graph (defined in terms of nodes and edges). This conversion can affect modular decompositions in opaque ways as well as invalidate presumed graph-theoretic null models [24, 25]; statistical dependency networks in particular require a number of non-trivial decisions regarding the choice of dependency measure (correlation, transfer entropy, phase-locking measures, etc.), treatment of positive vs. negative interactions, and thresholding [23]. Furthermore, coupling between variables does not necessarily give rise to large values of correlation or other dependency measures [26] (also as shown in our first example below). Finally, community detection on dependency networks optimizes quality functions that are difficult to interpret in terms of the original system dynamics. Perturbation modularity does not require the construction of a network representation of a dynamical system and is interpretable in terms of the separation of slow and fast time scales of perturbation spreading.

Finally, because our approach is based on intrinsic system dynamics, it also differs from methods that identify modules by *imposing* a dynamical process onto a given network, such as diffusion of random walkers [27, 28] or coupled phase oscillators [15, 29]. However, as we discuss below, in certain cases our approach has connections to such methods. In particular, it can be seen as a generalization of the random-walk-based approach of *Markov stability* [28, 30, 31] to a broad class of dynamics.

To formally define perturbation modularity, consider a dynamical system with an N -dimensional state-space \mathbf{X} and evolution operator $f^t : \mathbf{X} \rightarrow \mathbf{X}$ at time scale t (both state and time can be continuous or discrete). Given a set \mathcal{E} of possible initial perturbations, $\boldsymbol{\varepsilon} \in \mathcal{E}$ is applied to an initial condition $\mathbf{x} \in \mathbf{X}$ to produce a perturbed initial condition $\mathbf{x} + \boldsymbol{\varepsilon} \in \mathbf{X}$. After time t , the size of the perturbation in the whole system is measured as the norm of the difference between the perturbed and unperturbed trajectories: $\|f^t(\mathbf{x} + \boldsymbol{\varepsilon}) - f^t(\mathbf{x})\|$. The relative size of the perturbation within a *subsystem* S (a subset of system variables) is:

$$m_S^t(\mathbf{x}, \boldsymbol{\varepsilon}) = \frac{\|(f^t(\mathbf{x} + \boldsymbol{\varepsilon}) - f^t(\mathbf{x}))_S\|}{\|f^t(\mathbf{x} + \boldsymbol{\varepsilon}) - f^t(\mathbf{x})\|} \quad (1)$$

where the subscript S on the right hand side indicates a projection onto the dimensions indexed by S . For simplicity, we only consider cases where the system's perturbed and unperturbed trajectories have not merged ($\|f^t(\mathbf{x} + \boldsymbol{\varepsilon}) - f^t(\mathbf{x})\| > 0$) and Eq. 1 is well-defined.

Assume a partition of the system $\pi = \{S_1, \dots, S_K\}$ into K disjoint subsystems. The *coarse-grained perturbation vector* $\mathbf{y}_\pi^t(\mathbf{x}, \boldsymbol{\varepsilon}) = \langle m_{S_1}^t(\mathbf{x}, \boldsymbol{\varepsilon}), \dots, m_{S_K}^t(\mathbf{x}, \boldsymbol{\varepsilon}) \rangle$ captures the relative size of the perturbation in each subsystem. Due to the normalization in Eq. 1, \mathbf{y}_π^t is invariant to the dynamical expansion of the whole-system phase-space, instead reflecting only the relative effects of perturbations on different subsystems.

We now define perturbation modularity $Q^t(\pi, \mathbf{x})$ as the vector autocovariance of the coarse-grained perturbation vector:

$$Q^t(\pi, \mathbf{x}) = \mathbb{E}[\mathbf{y}_\pi^0(\mathbf{x}, \boldsymbol{\varepsilon}) \cdot \mathbf{y}_\pi^t(\mathbf{x}, \boldsymbol{\varepsilon})] - \mathbb{E}[\mathbf{y}_\pi^0(\mathbf{x}, \boldsymbol{\varepsilon})] \cdot \mathbb{E}[\mathbf{y}_\pi^t(\mathbf{x}, \boldsymbol{\varepsilon})] \quad (2)$$

where the expected values are taken over $P(\boldsymbol{\varepsilon})$, a probability distribution over perturbations (i.e. the elements of \mathcal{E}). The first term of Eq. 2 measures the degree to which perturbations persist within a partition's subsystems (i.e. initially perturbed subsystems remain affected after time t , while initially unperturbed subsystems remain relatively unaffected). The second term of Eq. 2 provides a baseline expectation of perturbation effects that accounts for differences in subsystem sizes.

As stated, the spread of perturbations in a modular system will be constrained by module boundaries. The optimal modular decomposition is the partition that maximizes perturbation modularity: $\pi^* = \arg \max_\pi Q^t(\pi, \mathbf{x})$.

Perturbation modularity (Eq. 2), as well as optimal modular decompositions, are state-dependent in that they are defined relative to an initial condition \mathbf{x} . Different criteria may be used to determine the choice of this initial condition, such as dynamical importance (e.g. an equilibrium state), particular research interest, or random selection. Alternatively, the modularity of entire state-space regions, rather than individual states, can be measured as the expectation of perturbation modularity over a distribution of initial conditions (e.g. by averaging across the entire system state space). Similarly, stochastic dynamical systems can be accommodated by taking expectations over future state distributions. For simplicity, however, these extensions are not considered in the present work.

In addition to initial condition, perturbation modularity and optimal decompositions also depend on the time scale t , which, as mentioned, can act as a resolution parameter. When there is not a time scale of *a priori* interest, optimal decompositions can be identified at multiple resolutions by sweeping across a range of time scales. Finally, the measure also depends on the kinds of perturbations applied, as specified by \mathcal{E} and $P(\boldsymbol{\varepsilon})$. In practice, perturbations can be selected according to domain knowledge (e.g. typically-encountered environmental perturbations) or using 'neutral' options (e.g. small increments to single variables). In many cases, initial perturbations should be localized to a small number of variables (i.e. the elements of \mathcal{E} are sparse) because the spread of perturbations is more pronounced when only a few subsystems are initially perturbed. As we will show, perturbations that simultaneously affect many variables probe the system at larger scales and uncover larger modules, providing another way to explore decompositions at multiple resolutions.

Like other temporally-localized methods [32], perturbation modularity also depends on the norm used to mea-

sure perturbation magnitude. Below, the ℓ_1 norm is used because it performs well and permits connections to community detection methods in graphs (see Supplemental Material [33] for definition of ℓ_p norms).

Specifically, perturbation modularity is related to the *Markov stability* method of community detection in graphs, which identifies communities as subgraphs that trap random walkers [28, 30, 31]. Similarly to perturbation modularity, Markov stability separates diffusion dynamics into fast intra-community and slow inter-community components. As shown in the Supplemental Material [33], perturbation modularity is equivalent to Markov stability when the system of interest exhibits diffusion dynamics, perturbations are homogenous increases to single variables, and the ℓ_1 norm is used to measure perturbation magnitude. More broadly, our approach can be seen as a generalization of Markov Stability to other dynamics.

In addition, ℓ_1 perturbation modularity on a dynamical system is equivalent to *directed weighted Newman's modularity* [34, 35] on a specially-constructed graph (see Supplemental Material [33]). In this graph, nodes correspond to system variables and the edge from node i to node j has weight:

$$w_{ij} = \mathbb{E} \left[m_{\{i\}}^0(\mathbf{x}, \boldsymbol{\varepsilon}) m_{\{j\}}^t(\mathbf{x}, \boldsymbol{\varepsilon}) \right]$$

where the expectation is over $P(\boldsymbol{\varepsilon})$ and the subscripts $\{i\}$ and $\{j\}$ indicate single-variable subsystems. This mapping permits perturbation modularity to be maximized using highly-efficient existing community detection algorithms (such as the Louvain method [36, 37] used for the examples below; code available online [38]).

Several criteria can be used to measure the quality of identified decompositions. High-quality decompositions have large perturbation modularity values (e.g. near 1 for ℓ_1 or ℓ_2 perturbation modularity, see Supplemental Material [33] for derivation of bounds on perturbation modularity). Additionally, high-quality decompositions are robust to small changes in system and optimization parameters. This can be quantified by measures of partition similarity like normalized mutual information (NMI) [39], an information-theoretic measure that ranges from 0 (maximally dissimilar partitions) to 1 (identical partitions). In several of the examples below, we plot NMI similarity between optimal decompositions identified at close values of t ; high NMI values indicate modular organization robust to small changes in time scale. Similar techniques are used in the Markov stability literature to identify time scales with robust decompositions [40].

We demonstrate our method on several examples of *coupled logistic maps*, non-linear discrete-time dynamical systems that have been used to explore spatially-extended chaos and pattern formation [41]. Assume a system of N variables, with $x_i(t)$ indicating the state of variable i at time t , and the transition function:

$$x_i(t+1) = (1 - \gamma)g(x_i(t)) + \gamma \sum_{j \neq i} \frac{k_{ji}}{d_i} g(x_j(t)) \quad (3)$$




FIG. 1. (Color online) System of 80 coupled logistic maps ($\alpha=2, \gamma=0.04$) (a) The coupling matrix exhibits hierarchical modularity at three levels. (b) The system is chaotic and no strong correlations between variables are observed over 10,000 time steps. (c) Perturbation modularity of optimal decompositions (PM, solid black) at different time scales t and NMI between optimal decompositions at successive times (dashed blue). Three time scale regions are robust (NMI=1, grey), corresponding to the three hierarchical levels of the coupling matrix (insets).

where $g(x) = 1 - \alpha x^2$ is the logistic map, parameter $\alpha \in [1, 2]$ controls the chaoticity, parameter $\gamma \in [0, 1]$ controls the coupling strength, ‘coupling matrix’ elements k_{ji} determine the influence of variable j on variable i , and $d_i = \sum_{j \neq i} k_{ji}$ normalizes the coupling strengths. When variables are homogeneously coupled to nearest neighbors on a 1-dimensional ring lattice, these systems are called *coupled map lattices* (CML) [41]. Coupled logistic maps display a rich variety of spatiotemporal patterns in different parameter regimes due to the interplay between inter-variable coupling (which ‘homogenizes’ variable states) and chaos (which injects variation into variable states).

We consider several examples of coupled logistic maps. Unless otherwise noted, perturbations consist of a uniform distribution over small increases to single variables: $\mathcal{E} = \{0.0001 \cdot \mathbf{e}_i : i = 1..N\}$, where \mathbf{e}_i is the i^{th} N -dimensional standard basis vector. The ℓ_1 norm is used to measure perturbation size.

In **Example 1**, we uncover modular organization that is present in a system’s coupling matrix, though not apparent in the correlation statistics. Consider an $N=80$ variable system with chaotic dynamics ($\alpha=2, \gamma=0.04$) and a hierarchically-modular coupling matrix (Fig. 1a). The system is composed of 8 tightly-coupled low-level modules ($k_{ji}=1$) with 10 variables each, pairs of which are nested within 4 mid-level modules ($k_{ji}=0.01$), pairs of which are in turn nested within 2 weakly-coupled high-level modules ($k_{ji}=0.0001$). A random state is used as the initial condition.

Because the system is strongly chaotic for these values of α and γ , there is no obvious ‘order parameter’ for identifying modular organization from system trajectories [15]; for instance, variable states are largely uncorrelated over 10,000 time steps (Fig. 1b). However, be-




FIG. 2. (Color online) Two 100-variable CMLs are compared: one ‘modular’ (top row; $\alpha=1.7, \gamma=0.1$) and one ‘diffusive’ (bottom row; $\alpha=1.9, \gamma=0.6$). (a,d) Spacetime plots of the effect of a single-variable perturbation. A pixel is colored black if the absolute difference between perturbed and unperturbed trajectory at a variable (vertical axis) exceeds 1% of the size of the system-wide perturbation at a given time (horizontal axis). (b,e) Spacetime plots of the optimal decompositions at different time scales; color indicates each variable’s subsystem. (c,f) Perturbation modularity of optimal decompositions (PM, solid black) at different time scales t and NMI between optimal decompositions at successive times (dashed blue). Stable decompositions are observed in the modular CML (top row).

cause perturbations first spread within low-level modules, then mid-level modules, and finally high-level modules, our method easily uncovers the hierarchical modular organization. Fig. 1c shows the perturbation modularity (black) and NMI (dashed blue) for optimal decompositions at different time scales. There are three robust time scale regions, corresponding to each of the three hierarchical levels of the coupling matrix (insets in Fig. 1c). Beyond time scale ~ 50 , perturbations have spread between the high-level modules; at this point, optimal decompositions reflect random fluctuations in initial conditions, and perturbation modularity and NMI values are near 0.

In **Example 2**, we investigate a more interesting case in which modularity emerges in a homogeneously-coupled CML. In some parameter regimes, spatial variation in initial conditions breaks the lattice coupling symmetry and leads to the emergence of modular *domains* (contiguous lattice regions) that constrain the spread of perturbations [42]. Such a ‘modular’ regime is investigated using a CML with $N=100$ variables and weak coupling-strength ($\alpha=1.7, \gamma=0.1$). The initial condition is set by iterating a random state for 10,000 time steps. Fig. 2a shows the spacetime plot of the effect of a single-variable perturbation to this initial condition: the perturbation spreads to several nearby variables until $t\sim 50$ but then remains confined within its domain. When computed over a uniform distribution of single-variable perturbations, our method uncovers robust modular organization for a large range of time scales (Fig. 2b), with optimal decompositions exhibiting high values of perturbation modularity and NMI (Fig. 2c).

The above system can be compared to a CML in a ‘dif-




FIG. 3. (Color online) State-dependence in the modular CML. For a 12,000 step trajectory starting from a random state, optimal decompositions (time scale $t=300$) are identified using states along this trajectory as initial conditions. (a) Optimal perturbation modularity (PM, solid black) grows with increasing trajectory steps (horizontal axis), indicating the emergence of robust modular structures. Trajectory step 10,000 (dotted green line) is the initial condition in examples 2 and 5. (b) Optimal decompositions identified at different trajectory steps; color indicates the subsystem of each variable (vertical axis).

fusive’ regime ($\alpha=1.9, \gamma=0.6$). For these parameters, the effects of perturbations spread freely across the lattice, as shown in the spacetime plot of Fig. 2d (initial condition is the same random state as in the modular CML iterated for 10,000 time steps). This system does not exhibit robust modular organization: optimal decompositions are not stable (Fig. 2e) and optimal perturbation modularity and NMI values are low (Fig. 2f). Once the effects of perturbations spread completely around the ring lattice at $t\sim 100$, both optimal perturbation modularity and NMI values are near 0.

In **Example 3**, we demonstrate state-dependent modularity by tracking the gradual emergence of modular organization over the course of a CML trajectory. The modular CML of example 2 ($\alpha=1.7, \gamma=0.1$) was started from a random state and iterated over a 12,000 step trajectory. The state encountered after 10,000 time steps was previously used as the initial condition in example 2. Here we find optimal decompositions (time scale $t=300$) when different states along the aforementioned trajectory are used as initial conditions. Over the course of the trajectory, optimal perturbation modularity grows through a series of plateaus (Fig. 3a), indicating the appearance of modular structures. Fig. 3b shows the optimal decompositions identified at different trajectory steps. Variables $\sim 1-40$ quickly form modular structures (by trajectory step ~ 2000), while variables $\sim 40-100$ need almost 7,000 steps to do so. This provides an example of *self-organized modularity*, or modular organization that emerges during a system’s dynamical evolution.

Previously, we showed that perturbation modularity captures the presence of stable modular structures in different CML parameter regimes (Example 2), and that it can uncover modular organization in a state-dependent manner (Example 3). In **Example 4**, we use perturbation modularity to characterize regions in the CML parameter space with respect to the onset of modularity.

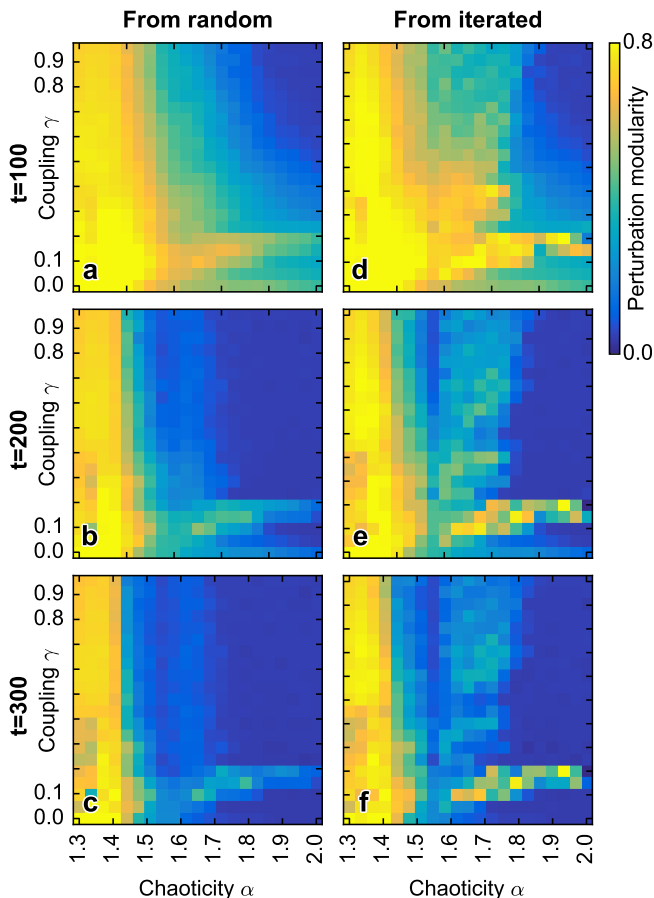


FIG. 4. (Color online) Parameter phase map of CML. Perturbation modularity (color) for optimal decompositions of 100-variable CMLs with different values of chaoticity (α ; horizontal axes) and coupling strength (γ ; vertical axes). Perturbation modularity is computed at three time scales for two different classes of initial conditions: random states [(a) $t=100$, (b) $t=200$, and (c) $t=300$] as well as random states iterated for 10,000 time steps [(d) $t=100$, (e) $t=200$, and (f) $t=300$].

Specifically, we construct 100-variable CMLs with different values of chaoticity (α) and coupling (γ) parameters. For these different parameter values, Fig. 4 shows values of optimal perturbation modularity computed at three time scales ($t=100$, $t=200$, and $t=300$) and two different classes of initial conditions: random states (Fig. 4a-c) and random states iterated for 10,000 time steps (Fig. 4d-f). In all cases, optimal perturbation modularity values were averaged across 10 random samples of initial conditions.

Several regimes of spatial organization can be identified in the parameter phase maps. For $\alpha \lesssim 1.44$, spatial domains, which form even when the system is started from random initial conditions, constrain the spread of perturbations over long time scales; we call this the *modular* regime. For other parameter values (e.g. $1.6 \lesssim \alpha \lesssim 1.95$, $\gamma \approx 0.1$, the yellow ‘tongue’ in Fig. 4d-f), modular domains only appear when random




FIG. 5. (Color online) Perturbation-dependence in the modular CML. (a) Spacetime plot of the effect of a perturbation to a 20-variable window (red arrows), as in Fig. 2a,d. (b) Optimal decompositions (time scale $t=300$) for different perturbation sizes (horizontal axis).

states are iterated for many steps before being used as initial conditions. This regime, which includes the case studied in Example 3, we call *self-organized modular*. Finally, for parameter values corresponding to the blue regions in Fig. 4, which we call the *diffuse* regime, modular domains are not present and perturbations spread freely. Here, different parameter values give rise to different diffusion speeds [42]: for example, $\alpha=1.9, \gamma=0.7$ exhibits no modular organization at time scale $t=100$; on the other hand, $\alpha=1.9, \gamma=0.2$ maintains some modularity at $t=100$, but this organization disintegrates at $t=200$.

Finally, in **Example 5**, we explore modularity’s dependence on the kinds of perturbations applied. We again consider the modular CML ($\alpha=1.7, \gamma=0.1$) and initial condition of example 2. Instead of perturbing single variables, we now simultaneously perturb multiple variables in lattice-contiguous ‘windows’ of different sizes (variables simultaneously incremented by 0.0001; all N windows are perturbed with uniform probability); for illustration, Fig. 5a shows the effect of a perturbation to a window of 20 variables. Fig. 5b shows that optimal decompositions (time scale $t=300$) depend on perturbation size. As more variables are perturbed, smaller subsystems merge into larger subsystems in a hierarchical manner.

Future work can pursue several extensions to our approach. First, estimating perturbation modularity from real-world datasets is of great practical interest; this can be investigated by applying the method to fitted dynamical models (e.g. vector autoregressive or dynamical causal modeling [43]) or using non-parametric approaches. Second, it is possible to explore other measures of decomposition quality beyond robustness to time scale, including robustness to changes in initial conditions and kinds of perturbations; alternatively, decomposition quality may be evaluated by testing the statistical significance of optimal perturbation modularity against null-model ensembles of non-modular dynamical systems [24]. Third, it is of interest to identify possible limitations of our method, such as for example whether it is susceptible to the kinds of resolution limits [44] and detectability

thresholds [45] encountered by graph-based community detection methods. Finally, future research can investigate other measures of perturbation magnitude (e.g. different norms or divergence functions), kinds of decompositions (e.g. overlapping subsystems), and cost functions (beyond vector autocovariance). For example, cost functions that capture the invertibility or sparsity of coarse-grained dynamics could be used to decompose a system into a mesoscopic ‘control diagram’, in which each subsystem controls a small number of others.

To summarize, we identify modular decompositions of multivariate dynamical systems based on the idea that modules constrain the spread of perturbations. We propose a quality function, called perturbation modularity, which can be used to identify optimal coarse-grainings that capture the slow component of perturbation spreading dynamics. The method generalizes graph-based com-

munity detection to a broad class of nonlinear dynamical systems and provides a principled alternative to detecting communities in network representations of dynamics. The method captures variation in modular organization across different time scales, initial conditions, and kinds of perturbations and offers a powerful tool for exploring modularity in complex systems.

ACKNOWLEDGMENTS

We thank Y-Y Ahn, Filippo Radicchi, and Daniel Damineli for helpful feedback. This work was supported in part by a NSF IGERT fellowship to AJG, Fundação para a Ciência e a Tecnologia (Portugal) grant PTDC/EIA-CCO/114108/2009, and the FLAD Computational Biology Collaboratorium (Portugal).

-
- [1] H. A. Simon, Proc. Am. Phil. Soc. **106**, 467 (1962).
 - [2] H. Kitano, Nat. Rev. Genet. **5**, 827 (2004).
 - [3] Y. Wu, P. Li, M. Chen, J. Xiao, and J. Kurths, Physica A **388**, 2987 (2009).
 - [4] G. P. Wagner and L. Altenberg, Evolution **50**, 967 (1996).
 - [5] G. Tononi, O. Sporns, and G. M. Edelman, PNAS **91**, 5033 (1994).
 - [6] G. P. Wagner, M. Pavlicev, and J. M. Cheverud, Nat. Rev. Genet. **8** (2007).
 - [7] R. K. Pan and S. Sinha, EPL **85**, 68006 (2009).
 - [8] B. Luque and R. V. Solé, Physica A **284**, 33 (2000).
 - [9] A. Pomerance, E. Ott, M. Girvan, and W. Losert, PNAS **106**, 8209 (2009).
 - [10] S. H. Strogatz, *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering* (Westview press, 2014).
 - [11] G. Koop, M. H. Pesaran, and S. M. Potter, J Econometrics **74**, 119 (1996).
 - [12] P. Fries, Trends in cognitive sciences **9**, 474 (2005).
 - [13] T. Womelsdorf, J.-M. Schoffelen, R. Oostenveld, W. Singer, R. Desimone, A. K. Engel, and P. Fries, Science **316**, 1609 (2007).
 - [14] J. Wu and O. L. Loucks, Q Rev Biol, 439 (1995).
 - [15] A. Arenas, A. Díaz-Guilera, and C. J. Pérez-Vicente, Phys. Rev. Lett. **96**, 114102 (2006).
 - [16] G. Palla, A.-L. Barabási, and T. Vicsek, Nature **446**, 664 (2007).
 - [17] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton, PNAS **108**, 7641 (2011).
 - [18] J. M. Carlson and J. Doyle, PNAS **99**, 2538 (2002).
 - [19] D. Deutscher, I. Meilijson, S. Schuster, and E. Ruppin, BMC Systems Biology **2**, 50 (2008).
 - [20] S. Fortunato, Physics Reports **486** (2010).
 - [21] D. Centola and M. Macy, Am. J. Sociol. **113**, 702 (2007).
 - [22] B. Zhang and S. Horvath, Stat. Appl. Genet. Molec. Biol. **4** (2005).
 - [23] M. Rubinov and O. Sporns, Neuroimage **52**, 1059 (2010).
 - [24] A. Zalesky, A. Fornito, and E. Bullmore, NeuroImage **60**, 2096 (2012).
 - [25] M. MacMahon and D. Garlaschelli, Phys. Rev. X **5**, 021006 (2015).
 - [26] M. C. Soriano, G. Van der Sande, I. Fischer, and C. R. Mirasso, Physical review letters **108**, 134101 (2012).
 - [27] M. Rosvall and C. T. Bergstrom, PNAS **105**, 1118 (2008).
 - [28] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, PNAS **107**, 12755 (2010).
 - [29] S. Boccaletti, M. Ivanchenko, V. Latora, A. Pluchino, and A. Rapisarda, PRE **75**, 045102 (2007).
 - [30] R. Lambiotte, J.-C. Delvenne, and M. Barahona, arXiv preprint arXiv:0812.1770 (2008).
 - [31] R. Lambiotte, J. Delvenne, and M. Barahona, IEEE TNSE (2015).
 - [32] A. Pikovsky and A. Politi, Nonlinearity **11**, 1049 (1998).
 - [33] See Supplemental Material at [URL will be inserted by publisher] for derivation of bounds on perturbation modularity, connection to Markov stability community detection, and mapping to Newman’s modularity.
 - [34] A. Arenas, J. Duch, A. Fernández, and S. Gómez, New J. Phys **9**, 176 (2007).
 - [35] E. A. Leicht and M. E. J. Newman, Phys. Rev. Lett. **100**, 118703 (2008).
 - [36] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, J. Stat. Mech. Theor. Exp. **2008**, P10008 (2008).
 - [37] R. Aldecoa and I. Marín, Bioinformatics **30**, 1041 (2014).
 - [38] Example code for identifying optimal decompositions at github.com/artemyk/perturbationmodularity.
 - [39] L. Danon, A. Diaz-Guilera, J. Duch, and A. Arenas, J. Stat. Mech. Theor. Exp. **2005**, P09008 (2005).
 - [40] M. T. Schaub, J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, PloS one **7**, e32210 (2012).
 - [41] K. Kaneko, Physica D **34**, 1 (1989).
 - [42] K. Kaneko, Physica D **23**, 436 (1986).
 - [43] K. J. Friston, L. Harrison, and W. Penny, Neuroimage **19**, 1273 (2003).
 - [44] S. Fortunato and M. Barthélemy, PNAS **104**, 36 (2007).
 - [45] R. R. Nadakuditi and M. E. J. Newman, Phys. Rev. Lett. **108**, 188701 (2012).