Learning interpretable collective variables for spreading processes on networks

Marvin Lücke ◉ and Stefanie Winkelmann ◉

Modeling and Simulation of Complex Processes, Zuse Institute Berlin, 14195 Berlin, Germany

Jobst Heitzig ◉

FutureLab on Game Theory and Networks of Interacting Agents, Potsdam Institute for Climate Impact Research, 14473 Potsdam, Germany
and Zuse Institute Berlin, 14195 Berlin, Germany

Nora Molkenthin ◉

Complexity Science Department, Potsdam Institute for Climate Impact Research, 14473 Potsdam, Germany

Péter Koltai ◉

Department of Mathematics, University of Bayreuth, 95447 Bayreuth, Germany

Collective variables (CVs) are low-dimensional projections of high-dimensional system states. They are used to gain insights into complex emergent dynamical behaviors of processes on networks. The relation between CVs and network measures is not well understood and its derivation typically requires detailed knowledge of both the dynamical system and the network topology. In this Letter, we present a data-driven method for algorithmically learning and understanding CVs for binary-state spreading processes on networks of arbitrary topology. We demonstrate our method using four example networks: the stochastic block model, a ring-shaped graph, a random regular graph, and a scale-free network generated by the Albert-Barabási model. Our results deliver evidence for the existence of low-dimensional CVs even in cases that are not yet understood theoretically.

Introduction. Networks of interacting agents are widely used to model sociodynamical phenomena [1] such as the spreading of a disease [2–4] or the diffusion of a (political) opinion within a society [5,6]. In such networks, nodes represent individual agents, and edges represent some form of social interaction. Each node has a state that evolves over time depending on the states of neighboring nodes. Often, stochastic effects are included to account for uncertainty in the dynamics and for the unpredictability of agents. These types of spreading processes are at the core of numerous open problems in a wide range of disciplines, such as understanding social collective behavior [7], assessing systemic risk in financial systems [8], or controlling modern power grids [9].

One approach to elevate our understanding of these models is to seek a low-dimensional representation of the system that captures the fundamental dynamics on timescales of interest. The projection into this low-dimensional space is called a *collective variable* (CV), and the projected dynamics is called the *effective dynamics*. Good CVs retain the essential information about the system's behavior, reducing the dimensionality and enabling a more efficient analysis and prediction.

For discrete-state spreading processes on certain simple networks, e.g., complete graphs and dense Erdős-Rényi random graphs, it is known that the shares of nodes in each state constitute a good CV, and its evolution is given by an ordinary differential equation in the mean-field limit [10]. There are many other results, valid for a varying range of networks and models, that describe the evolution of the shares of states in terms of (partial) differential equations in the mean-field or hydrodynamic limit [11–14]. Another popular choice of CVs are the counts of certain network motifs, e.g., the number of links between nodes of different states (typically called pair approximation [15–17]), and moment-closure methods can be used to approximate their evolution [18–22]. For the special case of binary-state dynamics, standard mean-field or pair-approximation theories have been complemented by higher-order master equations which improve accuracy by introducing fractions of neighbors of one or the other state, resolved with respect to the degrees of the nodes [23–27]. However, there is no all-encompassing theory relating any network topology and any process occurring on it to resulting CVs. Hence, a constructive computational approach—such as the one we will present—can elucidate cases that theoretical results do not yet cover. Moreover, the above-mentioned techniques *postulate* a candidate CV based on (physical) insights about the system, whereas our procedure does not require such intuition.

In this Letter, we extend the *transition manifold approach* [28,29] to learn CVs based on simulation data and to systematically find the relationship of the learned CVs to topological features of the network (see Fig. 1). The transition manifold approach assumes and exploits that the transition density functions of the system accumulate around a low-dimensional manifold, from which a CV can be inferred. The approach is designed such that most information of the density propagation of the process is retained under projection onto
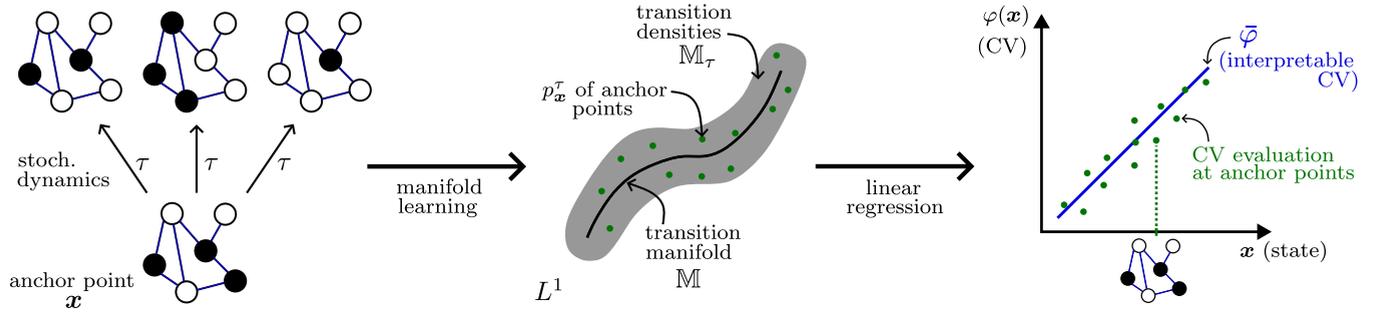
FIG. 1. Illustration of our method. Left: The random process is described by its distribution, sampled through *S* samples per initial network state (anchor). Middle: These distributions are used to learn a low-dimensional parametrization of the transition manifold. Right: A regression step allows for interpretability of the learned CV.

the CV [28]. To make it applicable to binary-state spreading processes on networks, we develop a technique for evenly sampling the state space and we add a linear regression step to produce interpretable CVs.

While goals similar to ours have been pursued in the literature before, to our best knowledge, this is the first work to learn CVs from data for networks of interacting agents. A crucial difference to previous works which learn reduced (also called surrogate or emergent) spaces [30–33] from data—even in the agent-based context [34]—is that they perform the reduction primarily based on the information of a state, and not on its dynamics. Once a good CV has been learned with our method, a surrogate dynamical model for its evolution could be determined and analyzed by tools also utilized in these references, in particular by linear transition operators associated with the dynamics [35–37]. Other approaches for data-based analysis and reduced modeling of interacting agent systems postulate CVs instead of learning them [38–42].

Recently, deep learning techniques have become popular for finding low-dimensional variables and surrogate dynamical models [35,36,43–48]. Artificial neural networks can represent coordinates from a large general class, but the dynamical conditions necessary for them to perform well remain implicit in these methods. Our approach, however, relies on explicit dynamical assumptions that are validated during the data-driven computation.

The transition manifold. Consider a fixed network of N nodes, on which each node $i \in \{1, \ldots, N\}$ has a discrete state $x_i \in \mathbb{S}$. The state space of the process is thus $\mathbb{X} := \mathbb{S}^N$, and its elements are system states $\boldsymbol{x} = (x_1, \ldots, x_N)$. Given a system state $\boldsymbol{x} \in \mathbb{X}$ and time $\tau \geqslant 0$, the *transition density function* $p_{\boldsymbol{x}}^{\tau} \in L^1(\mathbb{X}) =: L^1$ is defined such that $p_{\boldsymbol{x}}^{\tau}(\boldsymbol{y})$ is the probability that the system is in state $\boldsymbol{y}$ at time τ after having started in state $\boldsymbol{x}$ at time 0. (The term *density* comes from the original theory for continuous state spaces. We use L^1 to emphasize that these are vectors with entries summing to 1.) The *transition manifold approach* [28,29] exploits the observation that for certain systems and an appropriate choice of τ, the set

$$\mathbb{M}_{\tau} := \left\{ p_{\boldsymbol{x}}^{\tau} \mid \boldsymbol{x} \in \mathbb{X} \right\} \subset L^1 \tag{1}$$

is close to a d-dimensional submanifold $\mathbb{M} \subset L^1$ called the *transition manifold*. The lag time τ needs to be longer than the "relaxation time" towards $\mathbb{M}$ and shorter than the time it

takes to eventually converge to a stationary distribution (see Supplemental Material Sec. S2 [49] for details).

As a consequence, one can show that there exists a d-dimensional *collective variable* $\varphi : \mathbb{X} \to \mathbb{R}^d$, such that for all $\boldsymbol{x} \in \mathbb{X}$,

$$p_{\boldsymbol{x}}^{\tau} \approx \tilde{p}_{\varphi(\boldsymbol{x})}^{\tau}, \tag{2}$$

for some function $\tilde{p}_{(\cdot)}^{\tau}$. Hence, the essential information needed to characterize the dynamics is captured by the collective variable φ. A coordinate function φ satisfying this is for instance a "parametrization" of the manifold $\mathbb{M}$, in the sense that the assignment $\mathbb{M} \to \mathbb{X} \to \mathbb{R}^d$, $p_{\boldsymbol{x}}^{\tau} \mapsto \boldsymbol{x} \mapsto \varphi(\boldsymbol{x})$ is one to one (cf. Refs. [28,29]). Typically, the dimension d of the reduced state is significantly smaller than the dimension of the original state.

From now on, we consider binary-state dynamics, i.e., $\mathbb{S} = \{0, 1\}$. While this already covers a wide range of interesting dynamics, we expect that most of the following can be extended to an arbitrary number of states, at the cost of additional technicalities. For binary-state dynamics on a complete network, to give an example, a good collective variable is often given by the share of nodes in state 1, i.e., $d = 1$ and $\varphi(\boldsymbol{x}) = \sum_i x_i / N$ [10]. This is related to the quantity called magnetization in the similar Ising model [59].

Learning interpretable CVs. We propose the following method, which consists of three steps (Fig. 1).

Step 1. We choose a diverse set of dynamically relevant *anchor points* $\boldsymbol{x}^1, \ldots, \boldsymbol{x}^K \in \mathbb{X}$ in which the CV is going to be computed in the first instance. Diversity of the points is crucial in the sense that their respective transition densities cover $\mathbb{M}_{\tau}$ sufficiently well. Otherwise, the parametrization learned in the next step would yield a result biased by the insufficient coverage, that is potentially not a CV for the entire process. A precise quantification of sufficient coverage cannot be stated in generality as it depends on the system at hand and on the desired quality of the CVs. We employ an algorithm that prioritizes sampling anchor points containing communities of nodes with the same state, as this strategy produced the best results for the spreading processes we examined (see Supplemental Material Sec. S2 [49] for details).

Step 2. Next, we approximate a "parametrization" φ of the transition manifold $\mathbb{M}$ from simulation data. For each anchor point $\boldsymbol{x}^{\kappa}$ we conduct $S \in \mathbb{N}$ simulations of the process of length τ, yielding S samples for each transition density $p_{\boldsymbol{x}^{\kappa}}^{\tau}$. Using these data, we obtain evaluations of the collective vari-

able φ at the anchor points $\boldsymbol{x}^1, \ldots, \boldsymbol{x}^K$ by applying a manifold learning technique. Here, the dimension d of the CV is also an output. See Supplemental Material Sec. S2 [49] for further details on the approximation of the transition manifold.

The necessary number K of anchor points and S of simulations per anchor point depend on the size and complexity of the network. For the examples below ($N < 1000$), we found $K \approx 1000$ and $S \approx 100$ adequate. This very small number of anchor points compared to the number 2^N of all possible states is sufficient due to the targeted sampling method we employ (cf. step 1). We observed a substantial robustness of the results in varying the method's hyperparameters, but an entirely automatic procedure for their selection has yet to be designed. The output of this second step of the method is evaluations of the d-dimensional CV φ at the anchor points, $\varphi(\boldsymbol{x}^1), \ldots, \varphi(\boldsymbol{x}^K) \in \mathbb{R}^d$.

Step 3. The third step of the method aims at determining the meaning of the CV and finding a reasonable map $\bar{\varphi} : \{0, 1\}^N \to \mathbb{R}^d$ that extends it to states $\boldsymbol{x}$ outside of the original data set. Motivated by the fact that, for binary-state dynamics, the share of nodes in state 1 in (parts of) the network is known to be a good CV in specific cases [10], we propose maps $\bar{\varphi}$ of the form

$$\bar{\varphi}(\boldsymbol{x}) = \begin{pmatrix} \bar{\varphi}_1(\boldsymbol{x}) \\ \vdots \\ \bar{\varphi}_d(\boldsymbol{x}) \end{pmatrix}, \quad \bar{\varphi}_j(\boldsymbol{x}) = \sum_{i=1}^{N} \Lambda_{j,i} x_i, \qquad (3)$$

where $\Lambda \in \mathbb{R}^{d \times N}$ is a parameter matrix. For example, choosing $d = 1$ and $\Lambda = (1, \ldots, 1)$ yields a map describing the total count of state 1. In the different context of coupled ordinary differential equations, a CV similar to (3) was examined in Refs. [60,61].

We find optimal parameters Λ by employing linear regression to fit $\bar{\varphi}$ to the computed CV values in the anchors, $\varphi(\boldsymbol{x}^1), \ldots, \varphi(\boldsymbol{x}^K)$, from step 2. To prevent overfitting we use the *graph total variation* regularizer, which penalizes variation of Λ along the edges (see Supplemental Material Sec. S3 [49]). The reason for this choice is that each node in densely interconnected clusters is expected to contribute similarly to the CV. Even for networks containing hubs we can maintain this regularizer by bootstrapping our result to modify the ansatz functions (cf. Example 4 below). There we use linear basis functions for a first regression step, observe a strong correlation in Λ to the network's degrees, and modify the ansatz functions accordingly. If suggested by physical intuition, one can also entirely deviate from the linear ansatz functions we propose.

In the following examples binary-state spreading processes are studied, in which each node experiences memoryless random evolution in continuous time, with transition rates determined by a constant exploration rate (noise) and an "influence" rate based on the states in the node's neighborhood (see Supplemental Material Sec. S1 [49] for details). The investigations below refer to the *noisy voter model* [62,63], yet the Supplemental Material also includes findings related to alternative dynamics.

Example 1: Stochastic block model. We examine a network of $N = 900$ nodes that is constructed using the *stochastic block model*. The network consists of three clusters such that
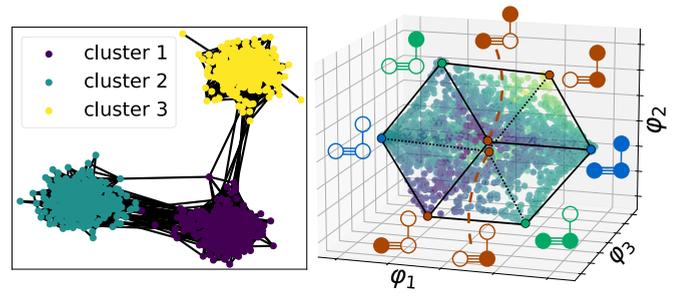
FIG. 2. For the stochastic block model network (left), the transition manifold is a three-dimensional cuboid (right). The vertices of the cuboid correspond to extreme states $\boldsymbol{x}$ where for each cluster either all (solid circle) or no nodes (open circle) have state 1.

cluster 1 and 2 are densely connected, cluster 1 and 3 are connected only sparsely, and cluster 2 and 3 are not connected at all (cf. Fig. 2). We expect the optimal CV to be $d = 3$ dimensional and contain the counts of 1's in each cluster. This CV is exact in the sense that for $N \to \infty$ it satisfies a mean-field equation [10]. Applying our method and plotting the resulting CV point cloud $\{\varphi(\boldsymbol{x}^1), \ldots, \varphi(\boldsymbol{x}^K)\} \subset \mathbb{R}^3$ yields an approximately cuboid-shaped transition manifold. We found that the vertices of this cuboid correspond to extreme states $\boldsymbol{x}$ in which for each cluster either all or no nodes have state 1 (cf. Fig. 2). To discover the meaning of the three coordinates $\varphi_1, \varphi_2, \varphi_3$, we calculate the optimal fit according to the regression problem proposed in step 3 of the method, which yields a collective variable $\bar{\varphi}(\boldsymbol{x}) = \Lambda \boldsymbol{x}$ with optimal parameters Λ shown in Fig. 3. The entries of the first row $\Lambda_{1,:}$ are approximately equal and thus $\bar{\varphi}_1$ describes the count of 1's in the whole network. The optimal $\Lambda_{2,:}$ is positive and constant within cluster 3 and negative and constant within clusters 1 and 2. Thus, $\bar{\varphi}_2$, calculates how the 1's are distributed between clusters $\{1, 2\}$ and $\{3\}$. Finally, $\Lambda_{3,:}$ is positive in cluster 1, negative in cluster 2, and approximately 0 in cluster 3, which implies that $\bar{\varphi}_3$ measures how the 1's are distributed between clusters 1 and 2, regardless of the number of 1's in cluster 3. Hence, the learned CV $\bar{\varphi}$ includes exactly the information that was predicted by theory [10] for large N, i.e., the counts of 1's
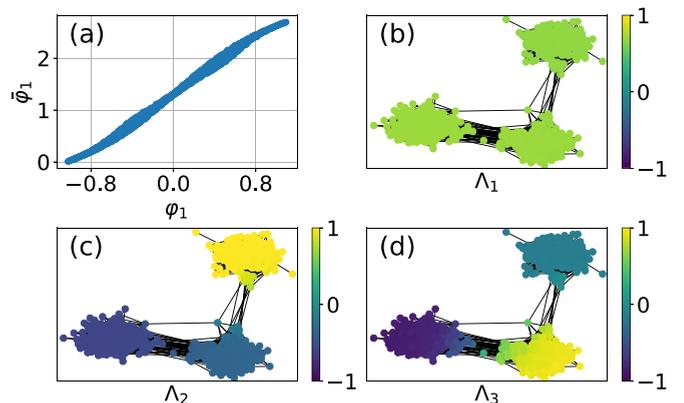
FIG. 3. Optimal Λ from (3) for Example 1. (a) Data $\varphi_1(\boldsymbol{x}^k)$ vs optimal fit $\bar{\varphi}_1$. [(b)–(d)] Optimal Λ entries for the respective coordinates plotted as color values on the network.
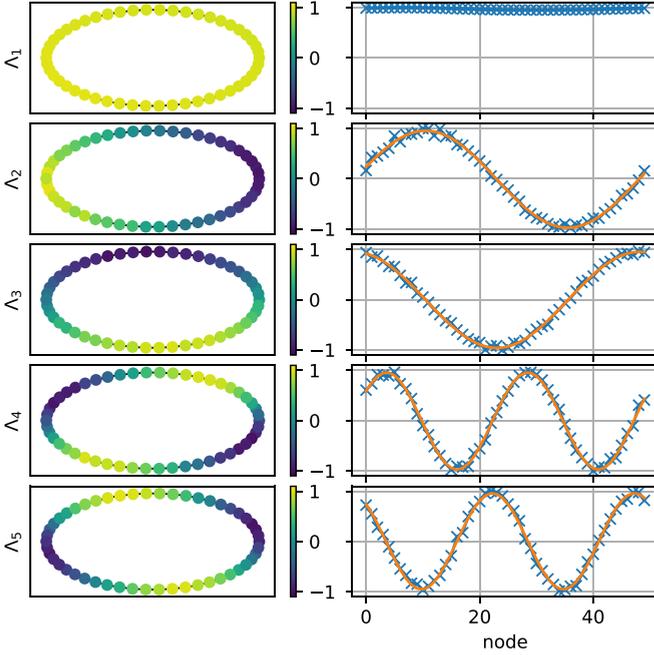
FIG. 4. Left: Optimal $\Lambda_{i,:}$ plotted as color values on the ring-shaped network. Right: $\Lambda_{i,:}$ (blue crosses) and a sine fit (orange line). The collective variables φ_i represent the real Fourier coefficients of the distribution of 1's on the ring, since $\varphi_i(\boldsymbol{x}) \approx \Lambda_{i,:} \boldsymbol{x}$ with the $\Lambda_{i,:}$ being sines and cosines of increasing frequencies.

for each cluster, but the coordinates are ordered by dynamical prevalence. (For instance, coordinate 3 is the least prevalent because information flows quickly between the two densely connected clusters 1 and 2.)

An interesting question would be to consider the change of the CV and especially its dimension with increasing edge density between the clusters. In particular, do structural transitions in the CV coincide with the so-called *detectability threshold* of the stochastic block model [64–66], where the edge statistics become indistinguishable from an Erdős-Rényi random graph model? This will be addressed in future work.

Example 2: Ring-shaped network. We apply our method to a ring-shaped network of $N = 50$ nodes. Examining the point cloud $\{\varphi(\boldsymbol{x}^1), \ldots, \varphi(\boldsymbol{x}^K)\}$ for different choices of d, we cannot identify a low-dimensional transition manifold as increasing d keeps adding valuable information. To keep the CV dimension reasonably small, we choose $d = 5$. (This yields CVs of reasonable quality; cf. Supplemental Material Sec. S5 [49].)

Solving the regression problem in step 3 yields a Λ that is constant in the first coordinate, i.e., the most important information is again the total count of 1's (see Fig. 4). The subsequent $\Lambda_{j,i}$ are pairs of sine and cosine functions of the node index i, starting with one oscillation for coordinates $j = 2, 3$ and then doubling the frequency for coordinates $j = 4, 5$. Hence, the collective variable $\bar{\varphi}$ measures the distribution of 1's on the ring, with increasing precision as we let d increase. This structure mimics Fourier coefficients, which suggests that (in the limit of infinitely many nodes) the optimal collective variable measures the position-dependent concentration of 1's as a density function on the ring. This result agrees well with other works considering ring-shaped or lattice networks
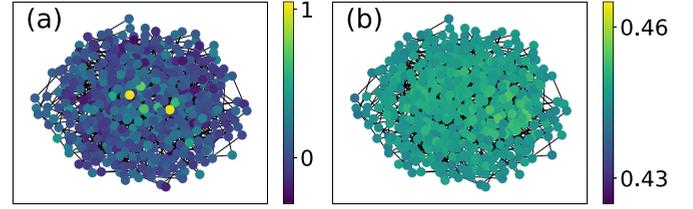
FIG. 5. (a) For the Albert-Barabási network, the optimal Λ as in (3) assigns a large weight to nodes with a high degree. (b) After preweighting with node degree [cf. (4)], the optimal Λ is constant. Hence, the collective variable describes the degree-weighted count of 1's.

(e.g., Refs. [11,13,14]), which find that the concentration of 1's is governed by a diffusive partial differential equation in the hydrodynamic limit. The CV of the system thus being a function on the ring, any finite-dimensional approximation has a truncation error. However, orthogonal trigonometric polynomials are a natural (and in an L^2-sense optimal) choice, found by our method.

Example 3: Random 3-regular network. One challenge in reduced modeling of spreading processes on random regular graphs is that edges are correlated. If the degree grows indefinitely with the network size, it was shown for the voter model that the share of state 1 is an asymptotically perfect CV [10]. In the same study it was observed numerically that for small degrees this CV still seems to support an effective dynamics, which deviates from the one obtained by mean-field approximation. Our method applied to a random 3-regular network validates the observation by reproducing this CV; see Supplemental Material Sec. S4 [49].

Example 4: Albert-Barabási network. Finally, we apply our method to a network generated by the Albert-Barabási model [67]. In the preferential attachment algorithm each new node is connected to $m = 2$ existing nodes, that are randomly picked with probability proportional to their degree. This procedure yields (asymptotically) a scale-free network. Applying our method results in a point cloud $\{\varphi(\boldsymbol{x}^1), \ldots, \varphi(\boldsymbol{x}^K)\}$ that indicates a $d = 1$-dimensional transition manifold (see Supplemental Material Fig. S3 [49]). The optimal $\Lambda \in \mathbb{R}^N$ according to the linear regression problem in step 3 assigns a large positive weight to nodes of a high degree, whereas nodes with a small degree have small or even negative weight (cf. Fig. 5). This conflicts with our choice of graph total variation regularizer that favors solutions for which Λ is equal for neighboring nodes. We tackle this issue by applying a preweighting of each node i with its degree d_i:

$$\bar{\varphi}(\boldsymbol{x}) = \sum_{i=1}^{N} \Lambda_i \, d_i \, x_i. \tag{4}$$

The optimal Λ for (4) becomes approximately constant, and hence the CV measures the degree-weighted count of state 1 in the system (cf. Fig. 5). Multiple experiments for varying parameters confirmed this result, provided the preferential attachment parameter is chosen $m \geqslant 2$. (For $m = 1$ the resulting networks exhibit a significantly larger diameter [68]. As a consequence, the degree-weighted count does not seem to sufficiently characterize the dynamics.) We are not aware of

any theoretical work showing that the degree-weighted count is a good CV for (binary-state or other) spreading processes on Albert-Barabási networks, although Refs. [69,70] hint at the significance of this observable.

Validation. A numerical validation of the CVs learned in the above examples is presented in Supplemental Material Sec. S5 [49].

Conclusion. We propose a method to learn interpretable CVs for spreading processes on networks without the need for prior expert knowledge about the network topology or the dynamical process. This method consists of the following steps: First, we sample anchor (network) states, from which we start many short simulations. Then we approximate the transition manifold and extend the learned CVs to unseen data using (total-variation-regularized) linear regression. The CVs are interpretable since the inferred parameters indicate the function and significance of features of the network structure. We have demonstrated this method for four different network topologies and two types of spreading dynamics (see Supplemental Material Sec. S4 [49]) and have thus shown its flexibility and usefulness. Although out of scope for the current Letter, we expect that the method can be generalized to processes with more than two discrete states as well as inhomogeneous agent dynamics.

[1] C. Castellano, S. Fortunato, and V. Loreto, Rev. Mod. Phys. **81**, 591 (2009).

[2] I. Z. Kiss, J. C. Miller, and P. L. Simon, *Mathematics of Epidemics on Networks* (Springer, Berlin, 2017).

[3] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, Rev. Mod. Phys. **87**, 925 (2015).

[4] S. Mauras, V. Cohen-Addad, G. Duboc, M. D. la Tour, P. Frasca, C. Mathieu, L. Opatowski, and L. Viennot, PLoS Comput. Biol. **17**, e1009264 (2021).

[5] S. Banisch and E. Olbrich, J. Math. Sociol. **43**, 76 (2019).

[6] A. Das, S. Gollapudi, and K. Munagala, in *Proceedings of the 7th ACM International Conference on Web Search and Data Mining, WSDM '14* (Association for Computing Machinery, New York, 2014), p. 403412.

[7] J. B. Bak-Coleman, M. Alfano, W. Barfuss, C. T. Bergstrom, M. A. Centeno, I. D. Couzin, J. F. Donges, M. Galesic, A. S. Gersick, J. Jacquet, A. B. Kao, R. E. Moran, P. Romanczuk, D. I. Rubenstein, K. J. Tombak, J. J. V. Bavel, and E. U. Weber, Proc. Natl. Acad. Sci. USA **118**, e2025764118 (2021).

[8] F. Caccioli, P. Barucca, and T. Kobayashi, J. Comput. Soc. Sci. **1**, 81 (2018).

[9] D. Witthaut, F. Hellmann, J. Kurths, S. Kettemann, H. Meyer-Ortmanns, and M. Timme, Rev. Mod. Phys. **94**, 015005 (2022).

[10] M. Lücke, J. Heitzig, P. Koltai, N. Molkenthin, and S. Winkelmann, Stoch. Process. Appl. **166**, 104220 (2023).

[11] E. Presutti and H. Spohn, Ann. Probab. **11**, 867 (1983).

[12] D. Keliger, I. Horváth, and B. Takács, Stoch. Process. Appl. **148**, 324 (2022).

[13] R. Durrett and W.-T. L. Fan, Ann. Appl. Probab. **26**, 3456 (2016).

[14] W.-T. L. Fan, Bernoulli **27**, 1899 (2021).

[15] M. J. Keeling, D. Rand, and A. Morris, Proc. R. Soc. London, Ser. B **264**, 1149 (1997).

[16] F. Vazquez and V. M. Eguíluz, New J. Phys. **10**, 063011 (2008).

[17] E. Pugliese and C. Castellano, Europhys. Lett. **88**, 58004 (2009).

[18] A.-L. Do and T. Gross, *Adaptive Networks: Theory, Models and Applications* (Springer, Berlin, 2009), pp. 191–208.

[19] M. Taylor, P. L. Simon, D. M. Green, T. House, and I. Z. Kiss, J. Math. Biol. **64**, 1021 (2012).

[20] P. Moretti, S. Y. Liu, A. Baronchelli, and R. Pastor-Satorras, Eur. Phys. J. B **85**, 88 (2012).

[21] C. Kuehn, *Understanding Complex Systems* (Springer, Berlin, 2016), pp. 253–271.

[22] A. F. Peralta, A. Carro, M. S. Miguel, and R. Toral, New J. Phys. **20**, 103045 (2018).

[23] K. T. Eames and M. J. Keeling, Proc. Natl. Acad. Sci. USA **99**, 13330 (2002).

[24] V. Marceau, P.-A. Noël, L. Hébert-Dufresne, A. Allard, and L. J. Dubé, Phys. Rev. E **82**, 036116 (2010).

[25] J. P. Gleeson, Phys. Rev. Lett. **107**, 068701 (2011).

[26] J. Lindquist, J. Ma, P. van den Driessche, and F. H. Willeboordse, J. Math. Biol. **62**, 143 (2011).

[27] J. P. Gleeson, Phys. Rev. X **3**, 021004 (2013).

[28] A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz, and C. Schütte, J. Nonlinear Sci. **28**, 471 (2018).

[29] A. Bittracher, S. Klus, B. Hamzi, P. Koltai, and C. Schütte, J. Nonlinear Sci. **31**, 3 (2021).

[30] R. R. Coifman, I. G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, Multiscale Model. Simul. **7**, 842 (2008).

[31] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, J. Chem. Phys. **134**, 124116 (2011).

[32] T. Berry, D. Giannakis, and J. Harlim, Phys. Rev. E **91**, 032915 (2015).

[33] P. Koltai and S. Weiss, Nonlinearity **33**, 1723 (2020).

[34] F. Kemeth, T. Bertalan, T. Thiem, F. Dietrich, S. J. Moon, C. Laing, and I. Kevrekidis, Nat. Commun. **13**, 3318 (2022).

[35] B. Lusch, J. N. Kutz, and S. L. Brunton, Nat. Commun. **9**, 4950 (2018).

[36] A. Mardt, L. Pasquali, H. Wu, and F. Noé, Nat. Commun. **9**, 5 (2018).

[37] G. Froyland, D. Giannakis, B. R. Lintner, M. Pike, and J. Slawinska, Nat. Commun. **12**, 6570 (2021).

[38] F. Lu, M. Zhong, S. Tang, and M. Maggioni, Proc. Natl. Acad. Sci. USA **116**, 14424 (2019).

[39] M. Zhong, J. Miller, and M. Maggioni, Phys. D Nonl. Phenom. **411**, 132542 (2020).

[40] J.-H. Niemann, S. Klus, and C. Schütte, PLoS One **16**, e0250970 (2021).

[41] N. Wulkow, P. Koltai, and C. Schütte, J. Nonlinear Sci. **31**, 19 (2021).

[42] L. Helfmann, J. Heitzig, P. Koltai, J. Kurths, and C. Schütte, Eur. Phys. J. Spec. Top. **230**, 3249 (2021).

[43] S. Brandt, F. Sittel, M. Ernst, and G. Stock, J. Phys. Chem. Lett. **9**, 2144 (2018).

[44] C. Wehmeyer and F. Noé, J. Chem. Phys. **148**, 241703 (2018).

[45] M. Raissi, P. Perdikaris, and G. E. Karniadakis, arXiv:1801.01236.

[46] W. Chen, H. Sidky, and A. L. Ferguson, J. Chem. Phys. **150**, 214114 (2019).

[47] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, Proc. Natl. Acad. Sci. USA **116**, 22445 (2019).

[48] S. E. Otto and C. W. Rowley, SIAM J. Appl. Dyn. Syst. **18**, 558 (2019).

[49] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.109.L022301 for additional details on the method and the experiments, as well as a numerical study of the quality of the learned collective variables, which includes Refs. [50–58].

[50] M. Porter and J. Gleeson, *Dynamical Systems on Networks* (Springer, Berlin, 2016).

[51] A. Bittracher, M. Mollenhauer, P. Koltai, and C. Schütte, Multiscale Model. Simul. **21**, 449 (2023).

[52] K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf, Found. Trends Mach. Learn. **10**, 1 (2017).

[53] R. R. Coifman and S. Lafon, Appl. Comput. Harmon. Anal. **21**, 5 (2006).

[54] A. Bittracher and C. Schütte, *Advances in Dynamics, Optimization and Computation: A Volume Dedicated to Michael Dellnitz on the Occasion of his 60th Birthday* (Springer, Berlin, 2020), pp. 132–150.

[55] R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer, IEEE Trans. Image Process. **17**, 1891 (2008).

[56] R. J. Tibshirani and J. Taylor, Ann. Stat. **39**, 1335 (2011).

[57] T. Hastie, J. Friedman, and R. Tibshirani, *The Elements of Statistical Learning* (Springer, New York, 2001).

[58] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, Math. Program. Comput. **12**, 637 (2020).

[59] B. A. Cipra, Am. Math. Mon. **94**, 937 (1987).

[60] J. Gao, B. Barzel, and A.-L. Barabási, Nature (London) **530**, 307 (2016).

[61] E. Laurence, N. Doyon, L. J. Dubé, and P. Desrosiers, Phys. Rev. X **9**, 011042 (2019).

[62] B. L. Granovsky and N. Madras, Stoch. Process. Appl. **55**, 23 (1995).

[63] A. Carro, R. Toral, and M. S. Miguel, Sci. Rep. **6**, 24775 (2016).

[64] J. Reichardt and M. Leone, Phys. Rev. Lett. **101**, 078701 (2008).

[65] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Phys. Rev. Lett. **107**, 065701 (2011).

[66] J. Banks, C. Moore, J. Neeman, and P. Netrapalli, Proc. Mach. Learn. Res. **49**, 383 (2016).

[67] R. Albert and A.-L. Barabási, Rev. Mod. Phys. **74**, 47 (2002).

[68] B. Bollobás and O. Riordan, Combinatorica **24**, 5 (2004).

[69] K. Suchecki, V. M. Eguíluz, and M. S. Miguel, Europhys. Lett. **69**, 228 (2005).

[70] G. Bianconi, Phys. Lett. A **303**, 166 (2002).