

Multiplex Decomposition of Non-Markovian Dynamics and the Hidden Layer Reconstruction Problem

Lucas Lacasa,^{1,*} Inés P. Mariño,^{2,3,12} Joaquin Miguez,⁴ Vincenzo Nicosia,¹ Edgar Roldán,^{5,6} Ana Lisica,⁷ Stephan W. Grill,^{8,9} and Jesús Gómez-Gardeñes^{10,11}

¹*School of Mathematical Sciences, Queen Mary University of London, Mile End Rd, E14NS London, United Kingdom*

²*Department of Biology, Geology, Physics & Inorganic Chemistry, Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain*

³*Institute for Women's Health, University College London, London WC1E 6BT, London, United Kingdom*

⁴*Department of Signal Theory & Communications, Universidad Carlos III de Madrid, 28911 Leganés, Madrid, Spain*

⁵*The Abdus Salam International Center for Theoretical Physics, Strada Costiera 11, 34151 Trieste, Italy*

⁶*Max Planck Institute for the Physics of Complex Systems, Nöthnitzerstraße 38, 01187 Dresden, Germany*

⁷*London Center for Nanotechnology, University College London, London, United Kingdom*

⁸*Biotechnology Center, Technical University Dresden, Tatzberg 47/49, 01309 Dresden, Germany*

⁹*Max Planck Institute of Molecular Cell Biology and Genetics, Pfotenhauer Strasse 108, 01307 Dresden, Germany*

¹⁰*Department of Condensed Matter Physics, University of Zaragoza, 50009 Zaragoza, Spain*

¹¹*GOTHAM Lab, Institute for Biocomputation and Physics of Complex Systems (BIFI), University of Zaragoza, E-50018 Zaragoza, Spain*

¹²*Department of Applied Mathematics, Lobachevsky State University of Nizhny Novgorod, Nizhniy Novgorod, Russia*



(Received 2 February 2018; revised manuscript received 7 June 2018; published 7 August 2018)

Elements composing complex systems usually interact in several different ways, and as such, the interaction architecture is well modeled by a network with multiple layers—a multiplex network—where the system's complex dynamics is often the result of several intertwined processes taking place at different levels. However, only in a few cases can such multilayered architecture be empirically observed, as one usually only has experimental access to such structure from an aggregated projection. A fundamental challenge is thus to determine whether the hidden underlying architecture of complex systems is better modeled as a single interaction layer or if it results from the aggregation and interplay of multiple layers. Assuming a prior of intralayer Markovian diffusion, here we show that by using local information provided by a random walker navigating the aggregated network, it is possible to determine, in a robust manner, whether these dynamics can be more accurately represented by a single layer or if they are better explained by a (hidden) multiplex structure. In the latter case, we also provide Bayesian methods to estimate the most probable number of hidden layers and the model parameters, thereby fully reconstructing its architecture. The whole methodology enables us to decipher the underlying multiplex architecture of complex systems by exploiting the non-Markovian signatures on the statistics of a single random walk on the aggregated network. In fact, the mathematical formalism presented here extends above and beyond detection of physical layers in networked complex systems, as it provides a principled solution for the optimal decomposition and projection of complex, non-Markovian dynamics into a Markov switching combination of diffusive modes. We validate the proposed methodology with numerical simulations of both (i) random walks navigating hidden multiplex networks (thereby reconstructing the true hidden architecture) and (ii) Markovian and non-Markovian continuous stochastic processes (thereby reconstructing an effective multiplex decomposition where each layer accounts for a different diffusive mode). We also state and prove two existence theorems guaranteeing that an exact reconstruction of the dynamics in terms of these hidden jump-Markov models is always possible for arbitrary

*Corresponding author.

l.lacasa@qmul.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

finite-order Markovian and fully non-Markovian processes. Finally, we showcase the applicability of the method to experimental recordings from (i) the mobility dynamics of human players in an online multiplayer game and (ii) the dynamics of RNA polymerases at the single-molecule level.

DOI: [10.1103/PhysRevX.8.031038](https://doi.org/10.1103/PhysRevX.8.031038)

Subject Areas: Biological Physics, Complex Systems, Interdisciplinary Physics

I. INTRODUCTION

Network science has emerged as a powerful unifying framework for studying the emergence of collective phenomena in real complex systems from different domains [1,2] and has allowed us to increase the accuracy and predictive power of minimal models of complex dynamics, including epidemic spreading [3], synchronization [4], or social dynamics [5]. One of the most fascinating challenges faced in the last few years by network science is the need to incorporate and couple several network structures in order to correctly capture the inherently multidimensional nature of interaction patterns in real-world systems. As a result, much effort has been recently devoted to the definition and study of multilayer and multiplex networks [6–8]. The ubiquity of such structures in social, biological, and technological systems has required the revision of several canonical dynamical models that were previously studied only on isolated complex networks, including percolation [9–13], diffusion dynamics [14,15], navigation [16–18], epidemics [19–22], evolutionary games [23–25], synchronization [26], or opinion dynamics [27,28]. Notably, the collective behavior of such complex systems depends strongly on whether they can be described by isolated networks or by coupled networks [29], highlighting the importance of the multilayer architecture of real-world systems.

Multilayer network models of real-world systems face two fundamental and dual challenges. The first one is the necessity to assess, in a systematic way, whether a multilayer network model is adequate to represent the system and when such a model gives redundant information. This challenge was first addressed in Ref. [30] and constitutes an intense field of research today. The dual challenge aims at understanding whether an empirical network whose multilayer character is not directly observable is genuinely monolayer or is only an aggregated projection of a hidden multilayer network [see Fig. 1(a) for an illustration of a multiplex network with $L = 2$ layers and its aggregated projection]. Such a scenario has received much less attention despite being, for instance, central for networks arising in natural systems whose architecture is not directly observable, as in genetic networks or in brain functional networks where pairs of nodes modeling different brain areas can interact according to an *a priori* unknown range of different biological pathways [31].

In this article, we provide a method to identify the hidden multilayer structure of a complex system from coarse-grained dynamical measurements of its state. We show that,

by using only local information extracted from simple random-walk statistics, it is possible to discriminate whether the underlying structure of the system is actually a single-layer or a multilayer network and, in the latter case, to estimate the number of interacting layers in the system. Note that methods to infer network topological properties via random-walk statistics have been explored previously [32,33]. Notably, our discrimination method exploits the breaking of Markovianity occurring in a coarse-grained multilayer random walk, while the method to estimate the most probable number of layers is based on a maximum *a posteriori* (MAP) probabilistic criterion, which can be implemented via numerical integration methods, including conventional grid-based approximations [34] or more sophisticated Monte Carlo algorithms [35], which we show increase the computational efficiency.

Interestingly, this paper not only deals with a particular problem of network science. As a matter of fact, we show




FIG. 1. (a) A multiplex network with $L = 2$ layers (L^1 and L^2) and $K = 7$ nodes. A random walker diffusing over this structure generates a two-dimensional time series $\{X(t), \ell(t)\}_{t=1}^N$, where $X(t)$ and $\ell(t)$ are the vertex and layer locations of the walker at time t . In many real-world cases, the layer indicator $\ell(t)$ is hidden, and one has access only to $\{X(t)\}_{t=1}^N$, i.e., to the series of states of the walker on the projected network (P) shown in the bottom of the figure. For $L > 1$, the resulting trajectory is non-Markovian: We rely on this Markovianity-breaking phenomenon property to detect multiplexity and to provide an estimate of the number of layers in the system by observing only $\{X(t)\}_{t=1}^N$. (b) Simple canonical model where we fix a prior on the topology of each layer: a cycle graph with homogeneous transition rates, with uniform interlayer transition rates. This model serves as a basis for a generic stochastic decomposition of non-Markovian dynamics on the aggregated network.

that the mathematical formulation can indeed be applied to signals of arbitrary origin—not necessarily random walkers navigating a network—and the multiplexity estimation framework reduces in the general case to a stochastic decomposition of the signal in terms of an effective multiplex network, whose layers play the role of independent dynamical modes. More concretely, non-Markovian dynamics can thereby be decomposed into a stochastic combination of diffusive modes by projecting the dynamics into an appropriate hidden jump-Markov model.

We validate the proposed methodology with numerical simulations of (i) random walks navigating hidden multiplex networks (thereby reconstructing the true architecture of the hidden multiplex networks) and (ii) both Markovian and non-Markovian continuous stochastic processes (thereby reconstructing an effective multiplex decomposition where each layer accounts for a different dynamical mode). Furthermore, we state and prove two existence theorems, guaranteeing that such multiplex decomposition is always possible for any finite-order Markovian and infinite-order (fully non-Markovian) process. Specifically, we show that random sequences generated by those processes can be exactly reconstructed as a random walk over an effective multiplex network. Finally, we apply our method to experimental recordings of two complex systems of different nature and show that the method can be leveraged to decompose noisy, non-Markovian processes into alternating combinations of simpler dynamics, and we extract valuable information accordingly.

The article is structured as follows: In Sec. II, we propose the methods for multiplexity detection and layer estimation, and we explore their performance in a few examples. In Sec. III, we discuss the analogy between multiplexity unfolding and the decomposition of non-Markovian processes as multiple-layer Markovian processes; therefore, we extend the methodology to continuous time processes. We also state and discuss the implications of two existence theorems, whose proofs are put in two appendixes for readability. In Sec. IV, we address real-world scenarios, where we analyze two sets of experimental recordings, namely, human mobility in an online environment and traces of RNA polymerase, which further showcase the applicability of the method. In Sec. V, we provide a discussion of our results. Mathematical details and additional examples can be found in the appendixes.

II. DESCRIPTION OF THE METHODS AND ILLUSTRATIVE EXAMPLES

Multiplex networks are the most ubiquitous class of multilayer networks. They are a natural model for online social networks [36], where a given individual can communicate with others via different platforms (e.g., Facebook, Twitter, email, etc.) or transportation networks [37,38], where a set of locations can be connected in a multimodal way (e.g., bus, train, underground, etc.).

A multiplex network is defined by a set of $L \geq 1$ interaction layers (networks), all of them having the same set of K nodes but different topology (different edge set), with the peculiarity that each node has a replica in each layer [see Fig. 1(a) for an illustration]. This structure is thereby fully described by a set of adjacency matrices $\{\mathbf{A}^{(\ell)}\}_{\ell=1}^L$, where $A_{\alpha\beta}^{(\ell)} = 1$ if there is an edge between nodes α and β at layer ℓ and zero otherwise. For simplicity, we label the different layers of the multiplex network with Roman letters (i, j , etc.) and the nodes of each layer with Greek letters (α, β , etc.).

We consider a random walker navigating a multiplex [16] defined as follows: Jumps between layers are governed by a Markov chain with an $L \times L$ transition matrix \mathbf{R}_L (R_{ij} is the probability to jump from layer i to layer j), while the dynamics within each individual layer ℓ is also Markovian and determined by a $K \times K$ transition matrix $\mathbf{T}^{(\ell)}$ (where $T_{\alpha\beta}^{(\ell)}$ is the probability to walk from node α to node β at layer ℓ). For simplicity, we only consider diffusive dynamics, where at each time step, the walker at node α on layer ℓ (i) remains in the same layer with probability $R_{\ell\ell} = 1 - r$ or instantaneously jumps with uniform probability $R_{\ell\ell'} = r/(L - 1)$ to a different layer ℓ' and, subsequently, (ii) diffuses to one of the neighbors of node α in the chosen layer according to the layer's internal dynamics (given by $\mathbf{T}^{(\ell)}$ or $\mathbf{T}^{(\ell')}$). Notice that this type of dynamical model can be mathematically formalized in terms of a jump-Markov affine system, as defined in the field of control theory (see Ref. [39] and references therein for a review).

When $r \ll 1$, i.e., when walkers tend to remain in the same layer, this navigation model might mimic, for instance, human mobility in multilayered transportation networks [40,41], where multimodality is minimized to avoid waiting times related to connections between different modes.

In the particular case when layers have a simple cycle-graph topology, the jump-Markov model described above is also reminiscent of the so-called discrete flashing ratchet model [42,43], better known as a Parrondo game [44,45]. This is a paradoxical gambling strategy that allows winning in losing scenarios, where gamblers can alternate between two different strategies (game A, layer 1; game B, layer 2), each of them having different rules and winning probabilities. Our model for $L = 2$ can be seen as a variant of a Parrondo gambler that plays with probabilities r and $1 - r$ with two different biased coins.

More generally, Brownian ratchets are paradigmatic models used in nonequilibrium physics to describe the transport of Brownian particles embedded in periodic, asymmetric energy potentials, a paradigm originally proposed by Smoluchowski [46] and popularized by Feynman [47] in the context of thermodynamic engines, and further shown to be a minimal model system for molecular motors in biophysics [43]. Again, a Brownian particle subject to a

periodic asymmetric potential that is switched on and off stochastically is formally equivalent to our random walker navigating over a multiplex with $L = 2$ cycle graphs with different transition matrices [48].

The general navigation process on a multiplex network discussed above can be expressed as a stochastic process fully described by an infinite two-dimensional time series $\{X(t), \ell(t)\}_{t=1}^{\infty}$, where $X(t) \in \{1, \dots, K\}$ and $\ell(t) \in \{1, \dots, L\}$. As in real-world scenarios, the multiplex nature of the system is not always empirically accessible, the layer indicator $\ell(t)$ is usually hidden, and the only observable is the sequence of node locations $X(1), X(2), \dots$. In such a situation, we only have experimental access to *partial* information of the process, described by a finite sequence of observations of the variable X : $\mathcal{O} = \{X(t)\}_{t=1}^N$. This is formally equivalent to observing a dynamical process on the aggregated (projected) network. Hence, the question is as follows: Is it possible to discern if the system is multiplex and, in that case, to estimate the most probable number of layers if we only have access to \mathcal{O} ? We now propose and test a novel method to achieve this highly nontrivial task.

A. Method to detect multiplexity

Consider the Markov switching walker discussed above, navigating on a (multiplex) network from which we only have access to coarse-grained information given by \mathcal{O} . Initially assuming $X(t)$ is a Markov process, we can estimate directly from \mathcal{O} the (monoplex) transition matrix \mathbf{Q} that would describe such a Markovian dynamics. Accordingly, we can define a Markov chain associated to \mathbf{Q} and generate a Markovian surrogate of the original process $\{Y(t)\}_{t=1}^N$. If the underlying network was truly monoplex, then $X(t)$ would actually be Markov, and $X(t)$ and $Y(t)$ would then have asymptotically equivalent statistics, $P_X(Z_1, \dots, Z_m) = P_Y(Z_1, \dots, Z_m)$, for all possible sequences Z_1, \dots, Z_m of any arbitrary length m . For multiplex structures however, losing information of $\ell(t)$, in general, breaks Markovianity, and therefore, $X(t)$ is typically non-Markovian. Accordingly, $X(t)$ and $Y(t)$ now share the same joint distributions only up to blocks of size $m = 2$. For blocks of size $m \geq 3$, their statistics may differ, $P_X(Z_1, \dots, Z_m) \neq P_Y(Z_1, \dots, Z_m)$. To quantify such a difference, we make use of the m th-order Kullback-Leibler divergence rate [49,50] between data blocks of size $m \geq 1$:

$$\mathcal{D}(m) := \frac{1}{m} \sum_{\mathcal{B}(m)} P_X(Z_1, \dots, Z_m) \log \frac{P_X(Z_1, \dots, Z_m)}{P_Y(Z_1, \dots, Z_m)}, \quad (1)$$

where $\mathcal{B}(m)$ enumerates all the blocks of size m . The statistic $\mathcal{D}(m)$ is semipositive definite and vanishes only when the joint probabilities coincide [51]. Thus, by construction, $\mathcal{D}(1) = \mathcal{D}(2) = 0$. The Markovianity-breaking criterion implies that if $\mathcal{D}(m) > 0$ for $m \geq 3$, then the underlying dynamics is multiplex [52].

As a proof of concept, we first consider the simple scenario where a random walker navigates over a two-layer multiplex ring (each layer is a cycle graph of K nodes), a model compatible with a discrete flashing ratchet as mentioned before. In the first layer, we define a Markov chain with homogeneous transition probabilities $T_{\alpha+1,\alpha}^{(1)} = 2/3$; $T_{\alpha,\alpha+1}^{(1)} = 1/3$ and $T_{\alpha\beta}^{(1)} = 0$ if $\beta \neq \alpha + 1 \bmod K$ or $\beta \neq (\alpha - 1) \bmod K$. A random walker diffusing in this layer will have an induced current in the direction of decreasing node indices. In the second layer, we define a different Markov chain with homogeneous transition probabilities $T_{\alpha+1,\alpha}^{(2)} = 1/2$; $T_{\alpha,\alpha+1}^{(2)} = 1/2$ and $T_{\alpha\beta}^{(2)} = 0$ otherwise, i.e., an unbiased random walk. While we can always estimate \mathbf{Q} numerically from the observed time series, in this simple case, it is easy to derive it analytically: $Q_{\alpha\beta} = W_1 T_{\alpha\beta}^{(1)} + W_2 T_{\alpha\beta}^{(2)}$, where W_1 (W_2) is the probability of finding the walker in layer $\ell = 1$ ($\ell = 2$). Now, since in this case $R_{ij} = R_{ji} = r$, the system is symmetric with respect to (w.r.t.) the switching process, and the walker spends, on average, the same amount of time in each of the two layers, $W_1 = W_2 = 1/2$, and then $Q_{\alpha\beta} = T_{\alpha\beta}^{(1)}/2 + T_{\alpha\beta}^{(2)}/2$. For this specific example, we thus find $Q_{\alpha+1,\alpha} = 7/12$, $Q_{\alpha,\alpha+1} = 5/12$.

The left panel of Fig. 2 shows numerical results of $\mathcal{D}(m)$ as a function of the block size m for different switching rates r . In order to deal with finite-size effects (which increase exponentially with m), we systematically increase the size of the walker series under study as a function of m , taking series of size $N(m) = N_0 \cdot 2^m$ data extracted from the original system and from the corresponding Markovian surrogates. We use $N_0 = 10^5$, although smaller values yield qualitatively equivalent results. As expected, $\mathcal{D}(1) = \mathcal{D}(2) = 0$, meaning that $Y(t)$ is a faithful Markovian surrogate of $X(t)$. Furthermore, $\mathcal{D}(m) > 0$ for $m \geq 3$, meaning that $X(t)$ is non-Markovian [49]; hence, the underlying network is correctly identified as a multiplex. This result is robust for a quite large range of values of r (see Appendix B, and note that $r = 0.5$ is a trivial exception), meaning that the method works even if the walker makes fast switches between layers. A similar scenario is found if we tune the transition probabilities such that no net-induced current is found (see Appendix B), pointing out that multiplexity can be unraveled even in that case. In the right panel in Fig. 2, we plot $\mathcal{D}(3)$ for different series sizes to showcase how finite-size effects vanish as the series gets larger. Notably, with a few thousand data points, we can already accurately detect multiplexity.

Furthermore, we have demonstrated the flexibility and robustness of our method to detect multiplexity in a range of additional scenarios, including (i) layers with increasingly different and disordered topologies controlled by both rewiring and edge addition, (ii) Erdos-Renyi graphs, and (iii) similar scenarios on larger graphs. For all cases, we

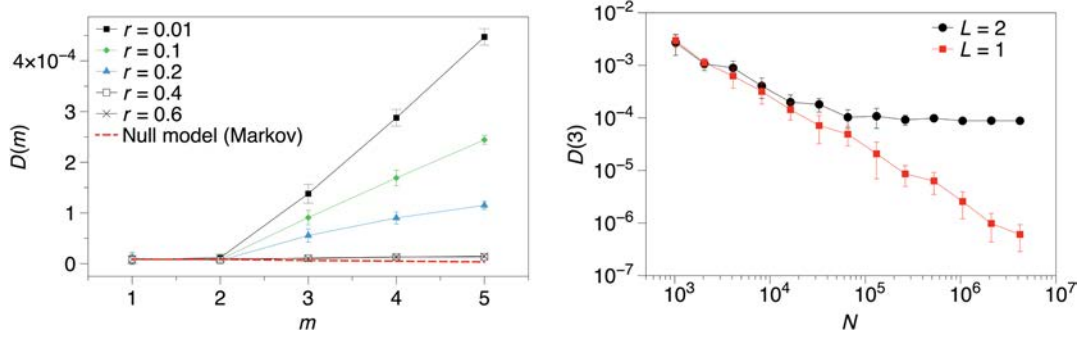


FIG. 2. (Left panel) Multiplexity detection statistic $\mathcal{D}(m)$ [see Eq. (1)] for $X(t)$, in a case where $X(t)$ shows an induced current, for different values of the switching rate r . The series $X(t)$ records the position of a walker diffusing over two layers, with transition probabilities in the layers given by $T_{a,a+1}^{(1)} = 1/3$ and $T_{a,a+1}^{(2)} = 1/2$, respectively, such that the transition probability in the Markovian surrogate series is $Q_{a,a+1} = 5/12$. We correctly find that $X(t)$ is non-Markovian even for large values of r as $\mathcal{D}(m)|_{m \geq 3} > 0$, which suggests an underlying multiplex structure. (Right panel) Finite-size-scaling analysis where we compute $\mathcal{D}(3)$ as a function of the series size N , for the multiplex considered in the left panel ($r = 0.1$, green dots) and a null model with equivalent monoplex dynamics (red squares) for which $\mathcal{D}(3)$ should vanish for large values of N . In both panels, the symbols represent the mean, and the error bars represent the standard deviation calculated over 10 different realizations.

find a correct multiplexity detection and good scalability (see Appendix B).

B. Method to quantify multiplexity

The Markovianity-breaking phenomenon that we have exploited only provides a means to discriminate whether hidden layers do exist in the multiplex, but not to quantify the number of underlying layers. In order to bridge this gap, we now make use of statistical inference tools to define a model selection scheme [53]. We assume that two models are different if they have a different number of layers. Accordingly, the number of layers L is now modeled as a random variable with prior probability mass function $P_0(L)$. Given the value of L , the motion of the random walker is determined by the Markov switching of layers, with transition matrix \mathbf{R}_L , and Markov walks within each layer characterized by $\mathbf{T}_L = \{\mathbf{T}^{(\ell)}\}_{\ell=1}^L$. Assuming prior probability density functions for these parameters $p_{0,R}(\mathbf{R}_L)$ and $p_{0,T}(\mathbf{T}_L)$, the likelihood of a given model with L layers conditional on the observed data $\{X(t)\}_{t=1}^N$ reads

$$P(\{X(t)\}_{t=1}^N | L) = \int P(\{X(t)\}_{t=1}^N | \mathbf{T}_L, \mathbf{R}_L) \times p_{0,T}(\mathbf{T}_L) p_{0,R}(\mathbf{R}_L) \mu(d\mathbf{T}_L \times d\mathbf{R}_L), \quad (2)$$

where μ is a suitable reference measure for \mathbf{T}_L and \mathbf{R}_L (see Appendixes C and D for technical details). In general, this multidimensional integral cannot be computed exactly, and it needs to be approximated numerically. Then, the number of layers L in the system that generated the data $\{X(t)\}_{t=1}^N$ can be detected using a MAP criterion, i.e.,

$$\hat{L}^{\text{MAP}} = \arg \max_{L \in \{1, 2, \dots, L_+\}} P(\{X(t)\}_{t=1}^N | L) P_0(L), \quad (3)$$

where \hat{L}^{MAP} is the estimate of the actual number of layers, $2 \leq L_+ < \infty$ is the (assumed) maximum possible value of L , and $P(L | \{X(t)\}_{t=1}^N) \propto P(\{X(t)\}_{t=1}^N | L) P_0(L)$ is the *a posteriori* probability mass function of L given the observations $\{X(t)\}_{t=1}^N$.

The practical computation of the MAP estimator in Eq. (3) can be addressed in different ways. The classical literature on hidden Markov models (HMMs) [54–56] suggests the use of the expectation-maximization (EM) algorithm (in various forms) to compute approximate maximum likelihood (ML) estimators, $\hat{\mathbf{T}}_L$ and $\hat{\mathbf{R}}_L$, of the parameters and then assume that $P(\{X(t)\}_{t=1}^N | L) \approx P(\{X(t)\}_{t=1}^N | \hat{\mathbf{T}}_L, \hat{\mathbf{R}}_L)$ in order to compare the models (i.e., one tries to optimize the parameters instead of averaging over them). This approach has also been applied to jump-Markov affine systems [57], and it relies on standard techniques; however, it has several drawbacks and is therefore not adopted here. Actually, the equation of the *model* likelihood with the *parameter* likelihood easily breaks down when the parameter estimates are poor (e.g., because of overfitting). Another major disadvantage of EM is that it converges locally and thus performs badly when the parameter likelihood is multimodal or when the parameter dimension varies significantly for different models. More sophisticated parametric schemes have been proposed (see, e.g., Ref. [58]); however, they are still subject to these fundamental limitations. Integration in Eq. (2), which we adopt here, has been favored theoretically but criticized practically because of the computational cost of approximating $P(\{X(t)\}_{t=1}^N | L)$ numerically [56]. However, we have found that state-of-the-art variational

Bayes [59] or adaptive importance sampling [60] methods can be applied effectively up to moderate values of L . See Appendixes C 3 and D for further discussion, including examples of using both deterministic integration and the adaptive Monte Carlo sampler from Ref. [61].

To illustrate the MAP model selection method given by Eq. (3), we consider again the discrete flashing ratchet model formed now by $L = 3$ ring-shaped layers with $K = 3$ nodes and homogeneous transition probabilities [see Fig. 1(b) for an illustration and Appendix D for other examples]. In this example, the probability to stay in the same layer is $R_{\ell\ell} = 1 - r = 0.84$, and the probabilities for the walker to move from node $\alpha \rightarrow \alpha + 1$ are $T_{\alpha,\alpha+1}^{(1)} = 0.16$, $T_{\alpha,\alpha+1}^{(2)} = 0.76$, and $T_{\alpha,\alpha+1}^{(3)} = 0.24$, respectively ($T_{\alpha+1,\alpha}^{(\ell)} = 1 - T_{\alpha,\alpha+1}^{(\ell)}$ for every ℓ , and $T_{\alpha\beta}^{(\ell)} = 0$ for any other α and β).

In order to evaluate the likelihood function $P(\{X(t)\}_{t=1}^N | L)$, we need to approximate the integral in Eq. (2) for each value $L = 1, \dots, 4$ as discussed previously. Here, we assume $L_+ = 4$. A very simple strategy is to evaluate it via numerical integration over a deterministic grid of 19 points on the interval (0,1) for each unknown parameter. For the case $L = 1$, this reduces to a single unknown, $T_{\alpha,\alpha+1}^{(1)}$, and for $L > 1$, there are $L + 1$ unknowns: r and $T_{\alpha,\alpha+1}^{(\ell)}$ for $\ell = 1, \dots, L$. Note that the particular choice of transition probabilities was taken to make the problem more challenging, as these values are not commensurate with the integration grid points. We use an unbiased prior probability density function $p_{0,R}(\mathbf{R}_L)$ given by a uniform probability density function on (0,1) for the unknown parameter r , while the prior $p_{0,T}(\mathbf{T}_L)$ is used to penalize system configurations with two or more identical layers [62]. For this particular numerical experiment, the penalizing prior is $p_{0,T}(\mathbf{T}_L) \propto \min_{(\ell,\ell')} |T_{\alpha,\alpha+1}^{(\ell)} - T_{\alpha,\alpha+1}^{(\ell')}|$; i.e., the prior probability density function of a given configuration \mathbf{T}_L is proportional to the minimum distance between any pair of matrices $\mathbf{T}^{(\ell)}$ and $\mathbf{T}^{(\ell')}$. Since in this example each scalar $\mathbf{T}^{(\ell)}$ fully characterizes layer ℓ , this prior simply penalizes configurations where two or more layers are very similar and thus avoids overfitting. In general, the prior $p_{0,T}(\mathbf{T}_L)$ is set to penalize models where pairs of layers have similar transition matrices, to avoid redundancy, so $p_{0,T}(\mathbf{T}_L) \propto \min_{(\ell,\ell')} \|\mathbf{T}^{(\ell)} - \mathbf{T}^{(\ell')}\|$, for some chosen norm $\|\cdot\|$.

Figure 3 shows that the true model is easily estimated with the proposed scheme, as the posterior probability emphatically peaks at $L = 3$, which implies that $\hat{L}^{\text{MAP}} = 3$. Without penalization—i.e., with a uniform prior $p_{0,T}$ —we obtain multiple equivalent solutions involving layers with identical values of the estimated parameters.

Now, it is well known that direct, grid-based deterministic integration of the posterior probability is intuitive but




FIG. 3. Posterior probability $P(L | \{X(t)\}_{t=1}^N)$ as a function of the number of layers L , computed from a trajectory of 2×10^4 time steps generated using a model with $L = 3$ cycle graphs ($K = 3$) and parameters $R_{ii} = 1 - r = 0.84$ for $i = 1, 2, 3$, $T_{\alpha,\alpha+1}^{(1)} = 0.16$, $T_{\alpha,\alpha+1}^{(2)} = 0.76$, and $T_{\alpha,\alpha+1}^{(3)} = 0.24$ (see the text and Appendix B for details). The algorithm easily estimates the correct model $L = 3$. (Inset) A linear-log plot of the same graph. Note that the probability for $L = 1$ is zero up to the computer's accuracy; hence, this point is not shown in logarithmic scales.

computationally inefficient. Accordingly, we have further considered alternative approximations of the integral in Eq. (2) using a nonlinear population Monte Carlo (NPMC) algorithm [61], and we have effectively reduced the runtime by a factor close to 100 on the same computer for the example of Fig. 3 (see Appendixes C and D for details). Actually, our NPMC is a sophisticated importance sampling procedure by which an efficient grid of the parameter space is obtained. This procedure meshes, in a tight way, the regions where the posterior probability density of the parameters is high and, in a sparse way, in the regions where the posterior probability density is low. Accordingly, the nodes of this mesh cover all the regions of the parameter space where the likelihood is sufficiently large (this approach ensures a global exploration of the parameter space, at odds with the EM algorithm, which can only perform a local search). A simple inspection of the likelihoods of each node in the grid allows us to robustly decide which is the one with highest likelihood. Accordingly, we are able to infer that the maximum of the posterior probability density $P(\{X(t)\}_{t=1}^N | \mathbf{T}_L, r) p_{0,T}(\mathbf{T}_L)$ [i.e., the integrand of Eq. (2) with uniform prior for r] is indeed attained at the true value of \mathbf{T}_3 (see Appendix D 5 for details). We therefore conclude that our model selection scheme correctly estimates not only the number of layers but also the transition probabilities within each layer, i.e., the full architecture. Finally, we have also verified that the posterior probability density is smooth close to its maximum, as perturbations $\tilde{\mathbf{T}}_3 = \mathbf{T}_3 + \delta\mathbf{T}$ systematically yield a smaller posterior probability density function for sufficiently large N (see Fig. 21).

III. MULTIPLEX DECOMPOSITION OF NON-MARKOVIAN DYNAMICS

As it happens for any Bayesian inference-based method [53,63,64], our approach, in principle, would provide a conclusive indication of the hidden multiplexity in the case where the prior—intralayer dynamics are diffusive—is a reasonable assumption. Notwithstanding, the methodology proposed here actually extends above and beyond the reconstruction of hidden multiplex architectures using walkers with partial information. As a matter of fact, a similar approach can be considered even when the architecture is *truly* single layered. Suppose, for instance, that a given observed time series is truly the outcome of a non-Markovian dynamics running on a physical single-layered network. In that case, our multiplex reconstruction method would still provide the most probable multiplex model, with $L > 1$ due to lack of the walker's Markovianity. The key difference is that now layers in the hidden multiplex would be providing the most probable *effective* multiplex reconstruction with Markovian intralayer dynamics, which would yield such complex dynamics. Incidentally, note that this brand of *effective* models is used in community detection in single-layered networks, which result from finding the optimal number of effective groups of nodes which maximize a certain likelihood function [53].

In what follows, we first capitalize on this new interpretation to extend our previous analysis on random walks on graphs to continuous stochastic processes, and then we present two existence theorems that guarantee that this stochastic decomposition is universally applicable to random sequences with arbitrary memory.

A. Extension to continuous processes

When the original dynamics is continuous, we can discretize motion and embed the original time series into a simple graph topology, for example, via the transformation $\{X(t)\} \rightarrow \{\tilde{X}(t)\}$ given by $\tilde{X}(0) = 0$ and

$$\tilde{X}(t+1) = \begin{cases} \tilde{X}(t) + 1 \mod K, & \text{if } X(t+1) > X(t) \\ \tilde{X}(t) - 1 \mod K, & \text{if } X(t+1) < X(t), \end{cases} \quad (4)$$

and apply our multiplexity detection methods to the discretized trace $\{\tilde{X}(t)\}$. To illustrate and validate this extension, we have considered two continuous stochastic processes: (i) an Ornstein-Uhlenbeck process (Markovian) governed by the stochastic differential equation

$$\dot{x} = -x + \xi, \quad (5)$$

where ξ is a Gaussian white noise with zero mean $\langle \xi(t) \rangle = 0$, amplitude σ , and autocorrelation $\langle \xi(t)\xi(t') \rangle = \sigma^2 \delta(t - t')$, and (ii) a generalization of the preceding

Langevin equation, where the noise term is not white anymore but has some color. Such a process can be described by the following Langevin equation:

$$\dot{x} = -x + \eta, \quad (6)$$

with $\eta(t)$ being defined by an Ornstein-Uhlenbeck process such that $\langle \eta(t) \rangle = 0$ and $\langle \eta(t)\eta(t') \rangle \propto \exp(-|t - t'|/\tau)/\tau$, with $\tau > 0$ the correlation time of the noise. Note that in Eqs. (5) and (6), the dot denotes the time derivative. When the noise term η is not white anymore, the Langevin equation (6) generates non-Markovian trajectories for the variable x . Interestingly, Eq. (6) has been attracting considerable attention in soft matter as a minimal model for the non-Markovian dynamics of the position of a passive particle immersed in an active (e.g., bacterial) bath [65,66].

We perform numerical simulations of these processes using a Euler-Mayurama algorithm and subsequently embed the resulting trajectories of 10^4 data in a cycle-graph topology [see Fig. 1(b)] via Eq. (4) with $K = 4$; then, we apply our multiplex detection and estimation protocol. In the top panels of Fig. 4, we plot $\mathcal{D}(m)$ and the outcome of a layer estimation. For the Ornstein-Uhlenbeck process, the method does not detect multiplexity (and a layer estimation confirms that a model with an $L = 1$ layer is the most probable one), in good agreement with the fact that the Ornstein-Uhlenbeck process is indeed a Markovian process. On the other hand, the Langevin equation with colored noise has a nontrivial memory kernel and generates non-Markovian dynamics, which is correctly captured by the fact that $\mathcal{D}(3) > 0$. In this case, the dynamics optimally decomposes into a Markov switching combination of $L = 2$ layers.

To further validate these results, we now generate synthetic trajectories from the estimated multiplex models with L layers. Importantly, we should recall that our layer estimation method makes use of an importance sampling algorithm to concentrate the Monte Carlo search in the regions of the parameter space with large likelihood. As a by-product, we are able to estimate—given the number of layers L —the parameters of the most likely model, and therefore, we can now generate a synthetic series from this fitted model. We then compare the statistics of a synthetic series from the most likely model with L layers with those of the original Ornstein-Uhlenbeck and non-Markovian processes, respectively. In the bottom panels of Fig. 4, we plot the m th-order Kullback-Leibler divergence between the original (discretized) series and the synthetic series generated by the reconstructed multiplex model. We confirm that the series generated by the model with $L = 1$ and $L = 2$, respectively, are the ones that show higher similarity with the Ornstein-Uhlenbeck (Markovian) and non-Markovian series, respectively. Finding $L = 2$ layers in the non-Markovian case is indeed reasonable, having in mind that Eq. (6) allows a well-known decomposition as

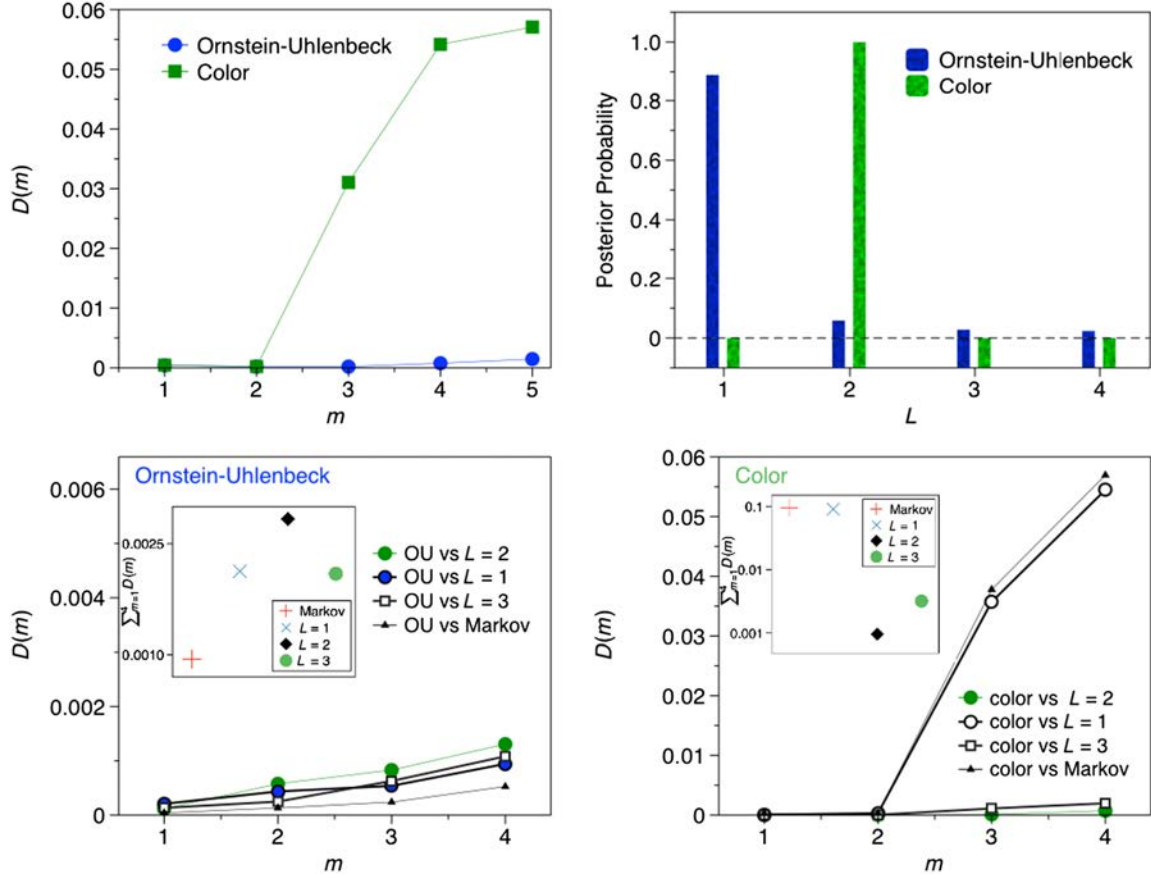


FIG. 4. (Top-left panel) Multiplexity detection statistic $D(m)$, applied to a trajectory generated by (blue) an Ornstein-Uhlenbeck process (Markovian, see the text) and (green) a Langevin equation with colored noise (non-Markovian, see the text) after embedding the trajectories in a ring topology via Eq. (4). The multiplexity detection is negative in the Markovian case and positive for the Langevin equation with a nonwhite noise term, as expected. (Top-right panel) Posterior probability of a multiplex model with L layers, confirming that $L = 1$ (i.e., monoplex) is the most probable model for the Markovian case (Ornstein-Uhlenbeck) and that a model with $L = 2$ is the most likely model for the case of a Langevin equation with correlated noise. (Bottom-left panel) $D(m)$ (order- m Kullback Leibler divergence) between the Ornstein-Uhlenbeck (Markovian) process and a synthetic series generated by the most likely model with L layers (for $L = 1, 2, 3$). Estimation of the model parameters (the architecture) is possible thanks to the importance sampling procedure, which preselects parameter configurations with high likelihood, in such a way that one can find the global optimum by searching in this reduced set. The synthetic series that shows the highest similarity with the original process is found for the model $L = 1$ (only improved by the comparison with a Markovianized time series) coinciding with the prediction found by the model selection scheme. (Bottom-right panel) Similar measures as the left panel, performed on the series extracted from the Langevin equation with colored noise (non-Markovian). In this case, the synthetic series that shows more similarity with the actual non-Markovian series is the one generated by the most likely model with $L = 2$ layers.

shown in the following 2D (Markovian) stochastic differential equation:

$$\dot{x} = -x + y, \quad \dot{y} = -\beta y + \xi, \quad (7)$$

with ξ a Gaussian white noise process and $\beta > 0$ a positive constant.

In even more general terms, our methodology provides a mathematically sound solution for the stochastic projection of non-Markovian dynamics onto a base of simple diffusive dynamical modes. For random sequences that were originally generated by a random walker navigating a hidden multiplex network, the method easily reconstructs such hidden architecture, whereas for general random sequences

[such as the ones generated by Eqs. (5) and (6)], the effective multiplex model provides a good reconstruction of the original dynamics. Hence, the question is as follows: Is this type of hidden jump-Markov model able to fully (i.e., exactly) reconstruct any random sequence (i.e., high-order Markovian and fully-non Markovian)? These questions are answered affirmatively in the next subsection, where we state and prove two theorems that address these matters.

B. Exact decomposition of Markovian and non-Markovian dynamics via multiplex models

In this subsection, we pose the question of whether the statistics of arbitrary random sequences $\{X(t)\}_{t \geq 0}$, possibly with long memory, can be recovered *exactly* using a hidden

jump-Markov model of the class described in Sec. II (from now on, we refer to it as a multiplex model). The answer is positive (in a probabilistic sense, to be made precise), as summarized by two representation theorems.

In particular, we first state and prove that every Markov model of finite (but arbitrarily large) order h can be recast into a multiplex model. Then, we address the representation of models with infinite memory, by letting $h \rightarrow \infty$, and show that, under mild regularity assumptions, they admit a compact representation in the form of a model with an uncountable (continuous) set of hidden layers. For readability, we have relegated as many technical details as possible (including all the mathematical proofs) to Appendixes F and G, and we only discuss the key results and implications here.

1. Representation of Markov models of order h

Let us consider a discrete state space with K elements, $\mathbb{K} = \{1, \dots, K\}$ (this will be the node set of the multiplex). A Markov sequence of (integer) order $h \geq 1$ is defined by the set of K^{h+1} probability masses

$$P_{\mathbb{K}}^h(i_t | i_{t-h:t-1}) := P(X(t) = i_t | X(t-h:t-1) = i_{t-h:t-1}),$$

where $X(t-h:t-1) = \{X(t-h), \dots, X(t-1)\}$ and $i_{t-h:t-1} = \{i_{t-h}, \dots, i_t\} \in \mathbb{K}^{h+1}$ is a sequence of $h+1$ state-space observations. If we fix $i_{t-h:t-1}$, then $P_{\mathbb{K}}^h(i_t | i_{t-h:t-1})$ is a probability mass function (PMF) and, as a consequence, $P_{\mathbb{K}}^h(i_t | i_{t-h:t-1}) \in [0, 1]$ and $\sum_{i_t \in \mathbb{K}} P_{\mathbb{K}}^h(i_t | i_{t-h:t-1}) = 1$.

Now, any Markov model of finite order h can be transformed into an *equivalent* multiplex model with a sufficiently large, but finite, number of layers L , as formally stated below.

Theorem 1: For every Markov model of order $h < \infty$ on the state space $\mathbb{K} = \{1, \dots, K\}$, there exists an equivalent multiplex model, with observation space \mathbb{K} and $L = K^h$ layers.

The proof of Theorem 1 is technical and is therefore put in Appendix F. Theorem 1 guarantees that every random Markov sequence with finite memory h can be represented by a multiplex model with K^h layers (i.e., this is an existence theorem). The theorem does not state that this representation is unique, though, and it does not state that it is minimal either. According to the numerical evidence given in the previous sections, we conjecture that a suitable selection of \mathbf{R}_L and \mathbf{T}_L (e.g., using the estimation methods described in this paper) can yield an accurate representation of a sequence of order h with considerably less than K^h layers.

Let us also remark that there is no contradiction between Theorem 1 and our earlier claim that multiplex models can represent systems with infinite memory. Indeed, depending on the choice of its parameters (L , \mathbf{R}_L , and \mathbf{T}_L), a multiplex system can yield random sequences with either finite or infinite memory. For example, in the proof of Theorem 1, we explicitly construct a multiplex model that matches

the transition probabilities of a Markov model of order h . On the other hand, multiplex models where the transitions between layers are independent of the walker's past trajectory yield sequences with infinite memory (except for pathological cases).

2. Representation of sequences with infinite memory

In the second part, we extend the previous existence theorem to a wide class of infinite memory (i.e., fully non-Markovian) models. Let $\{X_h(t)\}_{t \geq 0}$ denote a random sequence on $\mathbb{K} = \{1, \dots, K\}$ that can be represented exactly by a Markov model of order h . We consider here the class of sequences with infinite memory that can be obtained from Markov models as $h \rightarrow \infty$ and refer to them as “Markov- ∞ ” sequences. To be precise, we say that $\{X(t)\}_{t \geq 0}$ is the limit of $\{X_h(t)\}_{t \geq 0}$ as $h \rightarrow \infty$, and we write

$$X \stackrel{d}{=} \lim_{h \rightarrow \infty} X_h \quad (8)$$

when we can approximate the transition probabilities of the sequence $\{X(t)\}_{t \geq 0}$ with an arbitrarily small error using a Markov model of sufficiently large order. Specifically, we need to introduce the following technical definition:

Definition 1: The random sequence $\{X(t)\}_{t \geq 0}$ is Markov- ∞ if it satisfies the regularity conditions below:

(C1) The joint probability of any sequence of states vanishes uniformly with the length of the sequence, i.e.,

$$\lim_{t \rightarrow \infty} \sup_{i_{0:t} \in \mathbb{K}^{t+1}} P(X(0:t) = i_{0:t}) = 0.$$

(C2) There exists a sequence of Markov models $\{X_h(t)\}_{t \geq 0}$, $h = 1, 2, \dots$, such that for any $\epsilon > 0$ (arbitrarily small), there exists $h' \in \mathbb{N}$ (sufficiently large) that guarantees

$$\sup_{i_{0:t} \in \mathbb{K}^{t+1}} |P(X(t) = i_t | X(0:t-1) = i_{0:t-1}) - \dots - P_{\mathbb{K}}^h(i_t | i_{t-h:t-1})| < \epsilon,$$

for every $t > h$ and

$$\sup_{i_{0:t} \in \mathbb{K}^{t+1}} |P(X(t) = i_t | X(0:t-1) = i_{0:t-1}) - \dots - P_{\mathbb{K}}^h(i_t | i_{t-h:t-1})| = 0,$$

for every $t \leq h$, whenever $h > h'$.

Let $\{X_L(t)\}$ represent a random sequence generated by a multiplex model with L layers. Since every Markov- ∞ model is the limit of a sequence of Markov systems with increasing order, $h \rightarrow \infty$, they can also be obtained (via Theorem 1) as the limit of a sequence of multiplex systems as the number of layers grows, $L \rightarrow \infty$. Moreover, it turns out that the limit $\lim_{L \rightarrow \infty} \{X_L(t)\}_{t \geq 0}$ can be interpreted

itself as a multiplex model with an uncountable set of layers and a first-order Markov system on \mathbb{K} associated with each layer. This is made precise by the following theorem:

Theorem 2 Let $\{X(t)\}_{t \geq 0}$ be a Markov- ∞ random sequence on \mathbb{K} . There exists a Markov kernel

$$\mathcal{M}: \mathcal{B}([0, 1)) \times [0, 1) \rightarrow [0, 1], \quad (9)$$

where $\mathcal{B}([0, 1))$ denotes the Borel σ -algebra of subsets of $[0, 1)$, a probability measure $\mathcal{M}_0: \mathcal{B}([0, 1)) \rightarrow [0, 1]$, and an uncountable family of $K \times K$ transition matrices

$$\mathbf{T}(y) = \begin{bmatrix} T_{11}(y) & \dots & T_{1K}(y) \\ \vdots & \ddots & \vdots \\ T_{K1}(y) & \dots & T_{KK}(y) \end{bmatrix}, \quad y \in [0, 1), \quad (10)$$

such that

$$\begin{aligned} P(X(t) = i_t | X(0:t-1) = i_{0:t-1}) \\ = P_{i_0} \int \dots \int \prod_{k=1}^{t-1} T_{i_{k+1}i_k}(y_k) \mathcal{M}(dy_k | y_{k-1}) \mathcal{M}_0(dy_0), \end{aligned}$$

for every $i_{0:t} \in \mathbb{K}^{t+1}$, where $P_{i_0} = P(X(0) = i_0)$.

See Appendix G for a proof. Theorem 2 indicates that multiplex network models can be generalized to obtain probabilistic systems with an infinite and uncountable number of layers, which are flexible enough to represent (i.e., exactly recover) a broad class of random sequences with infinite memory.

IV. APPLICATIONS TO EXPERIMENTAL DATA

We now apply our methodology to two experimental recordings of completely different nature: the mobility dynamics of human players of an online video game and the dynamics of polymerases during RNA transcription.

A. Application to mobility: Pardus universe

Human and animal mobility [67–69] is often described by dynamical processes on single-layer networks and, interestingly, has been found to display signatures of memory [70–72], which can be interpreted as a deviation from Markovianity. Can such lack of Markovianity be interpreted as being the result of a Markovian dynamics taking place on a hidden multiplex network? In the case where mobility takes place across (hidden) multimodal transportation systems (as when we collect GPS traces of urban mobility), layers could be physical (underground, bus, car, etc.). On the other hand, animal foraging dynamics are clearly different and alternate during day and night [73]. In a similar vein, human mobility patterns change when switching from work to leisure styles [74]: These would be cases where layers would be effective rather than physical.

There are a large variety of problems involving the aforementioned scenarios, which could be amenable to our approach. To guarantee computational efficiency, we would only require that the network over which the agents move is not too large, something that in the general case can be achieved by coarse graining the network via community detection. Mathematically, within this context, a non-Markovian process running on a monoplex and a Markov switching process running on a multiplex are equally valid models, much in the same way a function is equivalent to its Fourier series representation. Still, we consider this new interpretation not just suggestive but also parsimonious from a cognitive point of view, and therefore, it might be of relevance in the study of memory in search processes.

To illustrate this type of application, we consider experimental mobility trajectories performed by players (agents) in a virtual environment: The Pardus universe (see Ref. [75] for details). The Pardus universe is a multiplayer online role-playing video game, which is used as a large-scale socioeconomic laboratory to study mobility in a controlled way. It consists of an (online) physical network with $K = 400$ nodes, a networked universe with social and economic activities, where the players of the game move around. The mobility traces of these players can then be naturally symbolized by coarse graining the original network via community detection into a network of 20 nonoverlapping communities wired by a 20×20 weighted adjacency matrix with complex topology [75].

Szell *et al.* analyzed online player's trajectories and reported evidence of long-term memory in the diffusing patterns of agent mobility in this universe [75]. Here, we apply the multiplexity detection statistic to a long time series obtained by concatenating individual traces. In Fig. 5, we show the numerical results (black dots), along with the null case of a Markovian diffusion (red squares) over the same network. We find $\mathcal{D}(3) > 0$ as a footprint of multiplexity, which can be interpreted here as the hidden presence of




FIG. 5. Multiplexity detection statistic $\mathcal{D}(m)$, applied to experimental mobility trajectories on the Pardus universe [75]. We find $\mathcal{D}(3) > 0$, which suggests that mobility series are non-Markovian, in agreement with independent evidence [75].

several *effective* layers. Szell *et al.* introduced a long-term memory model to account for the mobility dynamics and the heterogeneity of players. An alternative interpretation is that players are performing simple diffusion dynamics but are switching stochastically between different effective dynamical regimes. The challenge, which we leave as an open question for future work, would be to assign a social or perhaps cognitive meaning to the different layers. In this sense, our method here is closer to an unsupervised clustering paradigm than to a supervised one. Since our experimental trace is given by the concatenation of traces of different walkers, the different effective layers could also reveal a taxonomy of different game strategies.

B. Application to biology: Transcription by eukaryotic RNA polymerases

It has been shown that the ratchet mechanism plays an important role in active transport by molecular motors in living cells [43]. A paradigmatic example of a molecular motor that can be well described by a ratchet mechanism is RNA polymerase (RNAP). RNAPs are macromolecular enzymes responsible for the transcription of genetic information encoded in the DNA into RNA [76]. During

transcription, the spatial location of RNAPs along the DNA template exhibits noise due to thermal fluctuations. In addition, the dynamics of polymerases exhibits switching between two different dynamical regimes: an active polymerization state called elongation and a passive or diffusive state called backtracking [77]. While elongating, RNAP moves along the DNA template with a net velocity of the order of about 20 nucleotides per second, with 0.34 nm being the distance between two nucleotides. In the backtracking regime, the polymerization reaction is stalled, and RNAPs perform passive Brownian diffusion due to several types of noise (e.g., thermal and chemical). Optical tweezers enable measuring the motion of a single RNAP during transcription of a single DNA template at a base-pair resolution [77,78]. Single-molecule experimental RNAP traces are, however, extremely noisy, and the problem of identifying the possible mixture of different underlying dynamical processes—including transitions from elongation to backtracking as well as the hidden presence of other regimes—from a single experimental trace is a challenging task.

At the single-molecule level, the RNAP dynamics corresponds to a molecular motor producing a non-Markovian dynamics, which can thus be modeled as a Markov switching random-walk dynamics in a biochemical

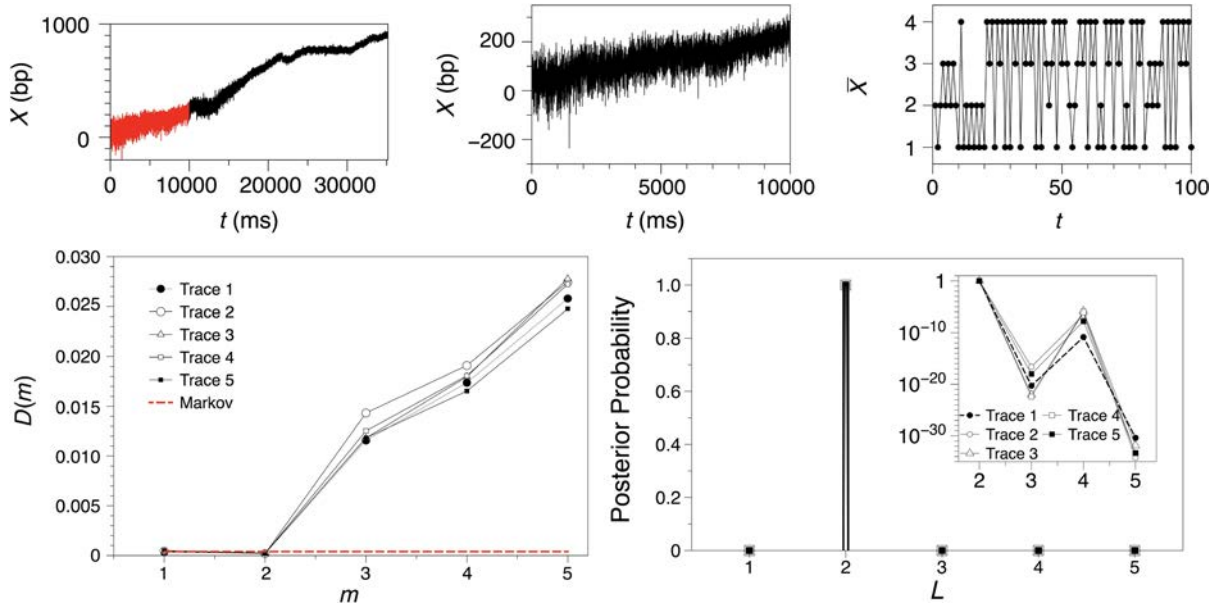


FIG. 6. (Upper panels) Experimental RNA polymerase I trace $X(t)$ (in base-pair units) from which we extract an excerpt of the first 10^4 time steps (highlighted in red in the original trace, shown in the middle panel) and a sample of the first 100 data of its symbolized trace $\tilde{X}(t)$ (right). Experimental traces of RNA polymerase I are extremely noisy, and one cannot always easily distinguish different dynamical regimes (see Appendix E for additional traces studied in this work). (Bottom-left panel) Multiplexity detection statistic $\mathcal{D}(m)$ applied to five (symbolized) experimental RNAP traces. We consistently find $\mathcal{D}(3) > 0$, which suggests that to correctly describe the dynamics of RNA polymerase I, we need at least two diffusive layers. (Bottom-right panel) Posterior probability for the layer estimation applied to five experimental series. We confirm that, with overwhelming probability, the most likely number of layers is $L = 2$. One possibility is that the identification of $L > 1$ is the consequence of the presence of colored noise in the single-molecule optical tweezer transcription experiment, whereas another possibility is that $L = 2$ corresponds to the elongation (active state) and backtracking (passive state) dynamical modes (in the inner panel, the same results are shown in log-linear scales, and the probability associated with $L = 1$ is zero so is not defined in a log scale).

multiplex with $L = 2$ layers (corresponding to elongation and backtracking, respectively) [79]. This suggests that our methodology (using cycle graphs with homogeneous transition rates as the prior topology) is applicable, and under the aforementioned assumptions, we should find that both $\mathcal{D}(3) > 0$ and the model with $L = 2$ should have a clear maximum likelihood. To test such a prediction, we have applied our complete methodology to five single-molecule experimental traces of the position X of RNA polymerase I (Pol I) from yeast *S. cerevisiae*, obtained with a dual-trap optical tweezer setup in the assisting force mode [80]. We systematically choose the first 10^4 data points at a 1-kHz sampling rate (see Fig. 6 and Appendix E) to keep the time series short and make the inference problem harder (for instance, for the series shown in Fig. 6, it is not easy to visually distinguish the elongation and backtracking regimes). Since the original traces have continuous state space, we discretize these series by embedding the experimental recordings into a cycle graph via the transformation proposed in Eq. (4) for $K = 4$. An example is shown in the top-right panel in Fig. 6. In the bottom-left panel of the same figure, we plot $\mathcal{D}(m)$ applied to $\{\tilde{X}(t)\}$ for each of the five experimental series, consistently showing $\mathcal{D}(3) > 0$. For comparison and control of finite-size effects, a similar measure is computed on a null model: a time series of 10^4 data points generated by a (monoplex) Markov chain whose transition matrix has been estimated from $\tilde{X}(t)$ (red dashed line). We can conclude that the underlying dynamics requires at least two alternating dynamics—a stochastic alternation between two diffusive layers—hence the projection onto an effectively multiplex model.

Subsequently, in the bottom-right panel of the same figure, we provide the results on the layer estimation using our nonlinear population Monte Carlo algorithm (convergence after 12 iterations, 10^2 samples per iteration). The algorithm directly provides $\log(P(\mathcal{O}|L))$, so assuming a uniform prior on the number of layers, $\log(P(\mathcal{O}|L)) \propto \log(P(L|\mathcal{O}))$, i.e., the logarithm of the *a posteriori* probability of model L . The true probability of the model (in natural units) is subsequently extracted and plotted accordingly. We find that the probability is essentially 1 for the model with $L = 2$ layers and negligible for the rest. Our algorithm thus reveals that the optimal hidden multiplex has $L = 2$ effective layers. One possibility is that the identification of $L > 1$ is the consequence of the presence of colored noise in the single-molecule optical tweezer transcription experiment. Quite intriguingly, additional evidence points to the presence of non-Markovianity also in the backtracking regime (see Appendix E), which might suggest the presence of colored noise even in the backtracking mode, as in the example in Sec. III. Another possibility is, as previously discussed, that the two effective layers correspond to the biochemical mechanisms of elongation (active state) and backtracking (passive state) dynamical modes. Finally, we expect our approach to also

be applicable to more complicated scenarios such as identifying the number of different nucleotides in copolymerization processes of templates with strong disorder [81].

V. DISCUSSION AND CONCLUSION

In this work, we have introduced a method that both detects and quantifies the degree of multiplexity in the hidden underlying structure of a networked system by only having access to local and partial statistics of a random walker. Our working hypothesis (prior) is that there is a hidden multiplex where walkers diffuse, switching layers stochastically and diffusing over each layer. Under these circumstances, any random walker for which we only see a projection of such trajectory in the aggregated network is necessarily non-Markovian if the number of layers is larger than one. Hence, our algorithm for multiplexity detection exploits such breaking of Markovianity as a means to detect multiplexity.

Incidentally, here we have focused on the specific case of multiplex networks, where every layer has the same number of (replica) nodes. Actually, in a multiplex, one can even have the same topology in each layer, where only transition weights differ, as in the case of the multiplex cycle graphs considered above. On the other hand, in a generic multilayer network, each layer will have, in general, a different number of nodes and different topology. This latter situation can be reinterpreted as having a multiplex where, in each layer, we can have isolated nodes that are never reached by a walker, or forbidden transitions. Accordingly, we envisage that layers in a multiplex would be, in general, harder to distinguish via our method than layers in a generic multilayer network, and as such, we expect the detection method to be easily generalizable to the multilayer case, something that should be studied in the future.

In a second step, we have introduced a probabilistic scheme to estimate the most probable number of layers composing the hidden multiplex. Note that probabilistic model selection is not new in network inference; for instance, in Ref. [82], a similar concept was used to estimate the most probable combination of basic block models that accounts for a certain network topology (see, also, Refs. [83–86]), whereas in Ref. [53], a probabilistic framework was developed to estimate the most probable number of communities in a single-layer network. Formally similar strategies to estimate model parameters based on ϵ -machines or jump-Markov system identification have also been put forward in the nonlinear dynamics [87] and control theory [57] communities, respectively. In our case, the posterior probabilities quantify the likelihood of having a hidden multiplex with L layers; i.e., our approach for multiplex model estimation is purely Bayesian. We were able to demonstrate the validity and accuracy of this second part in simple synthetic networks (multiplex cycles) up to $L = 5$ layers, as reported in the main text and appendices. Since the model selection protocol can be seen as a multidimensional Bayesian inference problem, these schemes—similarly to hidden Markov models and other methods in statistical inference—suffer from poor

scalability: Essentially, the computation of the model posterior probability explodes with the number of unknowns. This is a limitation of the method, and its optimization is therefore an open problem for future work. As a matter of fact, a simple and intuitive (although inefficient) way to estimate these posteriors is to use a (deterministic) grid integration scheme. In an effort to improve scalability and optimize such calculations, we have proposed a nonlinear population Monte Carlo algorithm (described in full detail in Appendixes C and D), which reduces the computer runtime by a factor close to 10^2 , without performance loss, for all the examples where we had previously used deterministic integration. Interestingly enough, our model selection scheme not only selects the most probable number of hidden layers: By capitalizing on an importance sampling routine that focuses on a small region of the parameter space where the likelihood is concentrated, we can also provide a Bayesian estimation of the model parameters (i.e., the topology and transition rates), thereby estimating not only the most probable number of layers but, given that number of layers, the full architecture.

We have thoroughly illustrated the validity and scalability of the whole method by addressing several synthetic systems of varying complexity, as depicted in Sec. II and the appendixes, where we show that we can reconstruct the full architecture of the hidden multiplex network by analyzing the statistics of random walks over the projected network.

Interestingly, our method can be applied to signals of arbitrary origin, extending the formalism to also deal with continuous processes—i.e., time series that are not necessarily random walkers navigating a network—after a simple graph embedding (series discretization). Under this extension, our hidden jump-Markov model provides a decomposition of a given non-Markovian dynamics into a Markov switching combination of diffusive modes and thus enjoys larger generality: The reconstructed multiplex model embeds the originally non-Markovian signal into a random walk navigating an effective multiplex network, where each layer accounts for a different type of diffusive dynamics. We validated this extension by analyzing canonical continuous processes (the Ornstein-Uhlenbeck process and a Langevin equation with colored noise). Furthermore, we have proved two existence theorems which guarantee that an exact reconstruction is always possible, for any type of (arbitrary) finite-order Markovian and fully non-Markovian (i.e., infinite memory) process.

Finally, we applied this methodology in two experimental scenarios (a case of human mobility in an online universe and the analysis of the dynamics of RNA polymerase) and hence showcased its applicability in real, experimental data.

To conclude, starting from the question of whether it is possible to disentangle the hidden multiplex architecture of a complex system if one only has experimental access to a projection of this architecture, in this work we have elaborated a mathematical and computational framework that actually

deals more generally with the decomposition of non-Markovian dynamics. When these series are indeed traces of walkers navigating a network, under the premise of having intralayer diffusion, our approach provides a workable solution for the unfolding of a multiplex network from its aggregated projection. We should make it clear that, generally speaking, it is not possible to assert that the effective multiplex representation is the *true* underlying architecture, much like one cannot typically claim that there exist true communities in an observed network. However, one can determine in a principled way whether the available observations are more reliably represented by a multiplex model or by a monoplex one, in the same vein as models of networks with community structure sometimes reproduce the observations more faithfully than models that lack such community structure.

In the general case, our method provides a potentially useful approach to disentangle the combination of dynamical regimes that appear intertwined in noisy dynamics; for instance, this is the case of RNA polymerase moving in a noisy environment and stochastically switching between an active and a passive state. This suggests that applications of this work not only include network science but extend to other fields in biophysics, condensed matter physics, gambling, or mathematical finance, where non-Markovian signals pervade.

ACKNOWLEDGMENTS

We sincerely thank Michael Szell, Roberta Sinatra, and Vito Latora for sharing data on the Pardus universe and for fruitful discussions, and we thank anonymous referees for useful comments. L. L. acknowledges funding from EPSRC Grant No. EP/P01660X/1. I. P. M. acknowledges the Spanish Ministry of Economy and Competitiveness (Projects No. TEC2015-69868-C2-1-R ADVENTURE and No. TEC2017-86921-C2-1-R CAIMAN) for financial support. J. M. acknowledges the Spanish Ministry of Economy and Competitiveness (Project No. TEC2015-69868-C2-1-R ADVENTURE) and the Office of Naval Research (ONR) Global (Grant No. N62909-15-1-2011) for financial support. I. P. M. also acknowledges support from the grant of the Ministry of Education and Science of the Russian Federation Agreement No. 074-02-2018-330. J. G. G. acknowledges financial support from MINECO (Projects No. FIS2014-55867-P and No. FIS2017-87519-P) and from the Departamento de Industria e Innovación del Gobierno de Aragón y Fondo Social Europeo (FENOL group E36_17R).

L. L., I. P. M., and J. M. contributed equally to this work.

APPENDIX A: A FEW NUMERICAL CONSIDERATIONS

Bounds on switching rate.—Numerically, the problem of detecting multiplexity via statistical differences between the Markovian surrogates $Y(t)$ and $X(t)$ should, in principle, be easier when the switching rate r is small enough

such that we allow trends associated with each layer to build up in the series, but large enough such that $\mathcal{D}(1) = \mathcal{D}(2) = 0$ (that is, large enough such that the one- and two-step joint distributions are still equivalent). A simple theoretical lower bound is $r < 1/m$ (as the average size of a trend is $1/r$, and this should be at least as large as the block size). Note that this bound is not tight in what we refer to as real-world scenarios (as for human mobility), where the switching rate is normally low in order to avoid unnecessary delays or has characteristic timescales that are much lower than the diffusion timescales (as for changing the foraging mode between day and night, in animal mobility).

Finite size bias in KLD.—As a technical remark, note that $\text{KLD}(p||q)$ diverges if the distributions p and q have different supports [i.e., if $q(m) = 0$, $p(m) \neq 0$ or $p(m) = 0$, $q(m) \neq 0$ for some value m]. In order to appropriately weight this possibility while maintaining the finite measure in pathological cases, a common procedure [49] is to introduce a small bias that allows for the possibility of having a small uncertainty for every contribution. Here, we introduce a bias of order $\mathcal{O}(1/n^2)$, where n is the series size (i.e., we replace all vanishing frequencies with $1/n$, and we normalize the frequency histogram appropriately).

APPENDIX B: INFERRING MULTIPLEXITY IN SOME ADDITIONAL SCENARIOS

We start by providing, in Fig. 7, an additional analysis similar to Fig. 2 but where there is no induced current. The




FIG. 7. The normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovian surrogate $Y(t)$, in the case where $X(t)$ does not show any induced current, for different values of the switching rate r . The series $X(t)$ records the position of a walker diffusing over two layers, where transition probabilities in the layers are $T_{i,i+1}^{(1)} = 1/3$ and $T_{i,i+1}^{(2)} = 2/3$. We correctly find that $X(t)$ is non-Markovian even for large values of r as $\mathcal{D}(m > 2) > 0$, which suggests an underlying multiplex structure.

method correctly detects multiplexity in this arguably more complicated scenario.

In the next sections, we extend the initial study on inferring multiplexity to the case where layers have different complex topologies, departing from the situation where each layer is a cycle graph (ring). Note that when each layer has a different topology, we expect the algorithm to more easily detect the underlying multiplex character (in this sense, extension of this formalism to the more general case of a multilayer network is very promising). In particular, we explore the following additional scenarios:

- (i) Scenario 2: Two complete graphs with different transition matrices.
- (ii) Scenario 3: Two layers with cycle graphs, where in one of the layers we introduce a shortcut, which is crossed with a probability ϵ . This scenario allows us to explore small perturbations in the transition matrices with respect to the original scenario studied in the main text.
- (iii) Scenario 4: Each layer has a different topology and dynamics, the first being a cycle graph (ring) with positive net current and the second layer being a complete graph with null net current.
- (iv) Scenario 5: Initially having two identical layers (two cycle graphs), we introduce a number of additional edges (shortcuts) in the second one to explore topological perturbations on our initial scenario. We also explore the scalability of this general scenario by considering the effect of increasing the number of nodes.
- (v) Scenario 6: Initially, we consider two identical layers formed by Erdos-Renyi graphs and then rewire a percentage of the edges in one of the layers. Transition matrices are obtained here by unbiasing the walker $T_{ij} = A_{ij}/k_i$. We also explore the scalability of this scenario by considering the effect of increasing the size of the graphs.

1. Scenario 2: Complete graphs

In this scenario, we consider two layers with identical topology but different transition matrices. Here, we build two replicas of a complete graph. In the first layer, we define an unbiased random walker with transition matrix

$$\mathbf{T}^{(1)} = \begin{bmatrix} 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \\ 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \end{bmatrix},$$

whereas in the second layer, we introduce a parametric deviation




FIG. 8. Scenario 2. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$, in the case where both layers are complete graphs with different transition matrices and the switching rate is $r = 0.1$. The difference is parametrized by ϵ (when $\epsilon = 0$, both matrices are identical, and they increasingly differ with increasing values of ϵ). The method detects multiplexity [$\mathcal{D}(m > 2) > 0$] with larger values as ϵ increases.

$$\mathbf{T}^{(2)} = \begin{bmatrix} 0 & 1/3 + 2\epsilon & 1/3 - \epsilon & 1/3 - \epsilon \\ 1/3 + 2\epsilon & 0 & 1/3 - \epsilon & 1/3 - \epsilon \\ 1/3 + 2\epsilon & 1/3 - \epsilon & 0 & 1/3 - \epsilon \\ 1/3 + 2\epsilon & 1/3 - \epsilon & 1/3 - \epsilon & 0 \end{bmatrix}.$$

The statistics of a Markov chain over each layer separately differ more for larger ϵ . For $\epsilon = 0$, both layers are identical, and a walker diffusing over the multiplex (switching layers at a constant rate r) reduces to a Markov chain over one layer, whereas for $\epsilon > 0$, the process is non-Markovian if we only have access to the state $X(t)$. In Fig. 8, we plot the results for $\mathcal{D}(m)$, which show that multiplexity can always be detected, and such detection is easier as ϵ increases.

2. Scenario 3: Controlled perturbation on one layer

In this scenario, we explore the effect of a controlled perturbation in the topology of one of the layers. We start by defining $L = 2$ identical layers (two rings with the same transition matrix, with a homogeneous probability to flow from $i \rightarrow i + 1$ $T_{i,i+1} = 1/3$). In the second replica, we introduce a shortcut between two nodes, weighting the probability of traversing this node as 2ϵ and biasing the rest of the edges accordingly. Thus, the transition matrices of both layers read

$$\mathbf{T}^{(1)} = \begin{bmatrix} 0 & 1/3 & 0 & 2/3 \\ 2/3 & 0 & 1/3 & 0 \\ 0 & 2/3 & 0 & 1/3 \\ 1/3 & 0 & 2/3 & 0 \end{bmatrix},$$




FIG. 9. Scenario 3. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$ in the case where the second layer has a shortcut, which is traversed with probability ϵ (scenario 3) and the switching rate is $r = 0.1$. For $\epsilon = 0$, both layers are identical, and the walker is essentially navigating over a monoplex; therefore, $X(t)$ is Markovian, and $\mathcal{D}(m) = 0 \forall m$. For $\epsilon > 0$, the graph is multiplex; therefore, $\mathcal{D}(m > 2) > 0$, and such a feature is detected with quantitatively larger fingerprints as ϵ increases.

$$\mathbf{T}^{(2)} = \begin{bmatrix} 0 & 1/3 - \epsilon & 2\epsilon & 2/3 - \epsilon \\ 2/3 + \epsilon & 0 & 1/3 - \epsilon & 0 \\ 2\epsilon & 2/3 - 3\epsilon & 0 & 1/3 + \epsilon \\ 1/3 + \epsilon & 0 & 2/3 - \epsilon & 0 \end{bmatrix}.$$

When $\epsilon = 0$, this extra edge has no effect, and we should expect that the multiplex is effectively monoplex; however, for $\epsilon > 0$, both layers show gradually different structures, and as such, $X(t)$ is non-Markovian. Similar to the previous case, our methodology predicts that the network is a multiplex when $\mathcal{D}(m \leq 2) = 0$ and $\mathcal{D}(m > 2) > 0$.




FIG. 10. Scenario 4. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$ in the case where the second layer is a complete graph.

Intuitively, as ϵ increases, the effect of the shortcut should be higher, and thus, detecting the multiplex nature of the network should be easier. In Fig. 9, we plot $\mathcal{D}(m)$ for a Markov chain walking over this topology with a constant switching rate $r = 0.1$ and different values of ϵ . The results support the accuracy of the method; even for small values of ϵ , the multiplex nature of the network is detected.

3. Scenario 4: Ring versus complete graph

In this scenario, we focus on a multiplex with $L = 2$, where each layer is totally different. The first layer is a ring with a net current described by the transition matrix

$$\mathbf{T}^{(1)} = \begin{bmatrix} 0 & 1/3 & 0 & 2/3 \\ 2/3 & 0 & 1/3 & 0 \\ 0 & 2/3 & 0 & 1/3 \\ 1/3 & 0 & 2/3 & 0 \end{bmatrix},$$

whereas the second layer is a complete graph with no net current and detailed balance everywhere:

$$\mathbf{T}^{(2)} = \begin{bmatrix} 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \\ 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \end{bmatrix}.$$

Again, we find that the method can clearly detect the multiplex nature of the network as $\mathcal{D}(m > 2) > 0$ (see Fig. 10).

4. Scenario 5: Sequentially introducing shortcuts

In this scenario, we consider two identical replicas (a multiplex with $L = 2$ layers), where in the second layer, we add a certain number of additional edges. Originally, both replicas are rings with detailed balance (unbiased walker),

$$\mathbf{T}^{(1)} = \mathbf{T}^{(2)} = \begin{bmatrix} 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \end{bmatrix}.$$




FIG. 11. Scenario 5. Adjacency matrices of each layer for three cases: (i, left panel) No shortcuts are introduced, and both layers are identical; (ii, middle panel) one shortcut has been introduced; (iii, right panel) two shortcuts have been rewired. The transition matrices for each case are $T_{ij} = A_{ij}/k_i$, where k_i is the degree of node i .




FIG. 12. Scenario 5. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$ for a two-layer multiplex, where the second layer is a replica of the first layer in which a number of shortcuts have been introduced. When no edges have been introduced, both replicas are identical, $X(t)$ is Markovian, and $\mathcal{D}(m) = 0 \forall m$; otherwise, the process is non-Markovian, and the algorithm detects multiplexity by finding $\mathcal{D}(m > 2) > 0$. Such detection improves as the topology of both layers is increasingly different.

In this initial case, we expect $\mathcal{D}(m) = 0, \forall m$ (the multiplex is effectively monoplex). We then add to this benchmark a different number of edges [and, accordingly, we expect $\mathcal{D}(m > 2) > 0$, and larger values when the number of added edges increases], effectively interpolating between scenario 2 and scenario 3. The adjacency matrices for three concrete cases (with no added edges, one added edge, and two added edges—this latter case being equivalent to a complete graph) are depicted in Fig. 11, and in Fig. 12, we show the values of $\mathcal{D}(m)$.

5. Scenario 5b: Introducing edges on larger graphs

Here, we investigate the scalability of scenario 5, and in particular, we explore (i) the effect of increasing the number of nodes K in each layer and (ii) the effect of increasing the number of rewired edges on the detectability. We start by defining two identical replicas of a ring with K nodes with unbiased transition matrix $T_{ij} = 1/2$ for $i \neq j$. In the




FIG. 13. Scenario 5b. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$ for a two-layer multiplex where each layer has K nodes, and the second layer is a replica of the first layer, where $p = 2$ shortcuts have been introduced.




FIG. 14. Scenario 5b. Detectability measure $\mathcal{D}(3)$ for a two-layer multiplex with $K = 10$ nodes per layer, where the second layer is a replica of the first layer in which a number of shortcuts p have been introduced. Results have been averaged over 10 network realizations for each case.

second layer, we introduce a number of shortcuts, and we analyze two particular behaviors as follows.

Effect of node increase.—We fix the number of shortcuts $p = 2$ and vary the number of nodes K , and we explore the dependence of $\mathcal{D}(m)$ on K (see Fig. 13). As we keep the series size $\mathcal{N}(m) = 10^5 \times 2^m$ independent from K , we expect that as K increases, the statistics are poorer as we need larger series to capture an equivalent number of transitions.

Detectability as a function of the number of shortcuts introduced.—Here, we fix $K = 10$ and explore the multiplex detectability as the number of shortcuts p is increased in the second layer. Multiplexity is detected when $\mathcal{D}(3) > 0$, and such detection is easier with larger $\mathcal{D}(3)$. In Fig. 14, we plot $\mathcal{D}(3)$ as a function of the number of shortcut edges p .

6. Scenario 6: Rewiring edges

In this scenario, we initially consider two replicas of the same Erdos-Renyi graph (where nodes i and j are connected with probability $p = 0.8$, above the percolation threshold, to have a connected graph). We consider three different situations: (i) Both layers are kept identical, (ii) we rewire one edge at random, and (iii) we rewire two edges at random. The adjacency matrices of each layer for these three cases are represented in Fig. 15, and we choose unbiased random walkers with layer transition matrices $P_{ij} = A_{ij}/k_i$. In Fig. 16, we plot the values of $\mathcal{D}(m)$ for these three cases. As expected, when the layers are identical, we find $\mathcal{D}(m) = 0 \forall m$, whereas when we rewire edges from a layer, the network is converted into a multiplex one, and $\mathcal{D}(m > 2) > 0$. Also, $\mathcal{D}(m > 2)$ takes larger values—and thus multiplex detection is easier—when the layers are increasingly different.

7. Scenario 6b: Rewiring edges on larger graphs

Finally, we consider Erdos-Renyi graphs (linking probability 0.65) with $K = 10$ nodes per layer and explore the multiplex detectability as we rewire a percentage of nodes

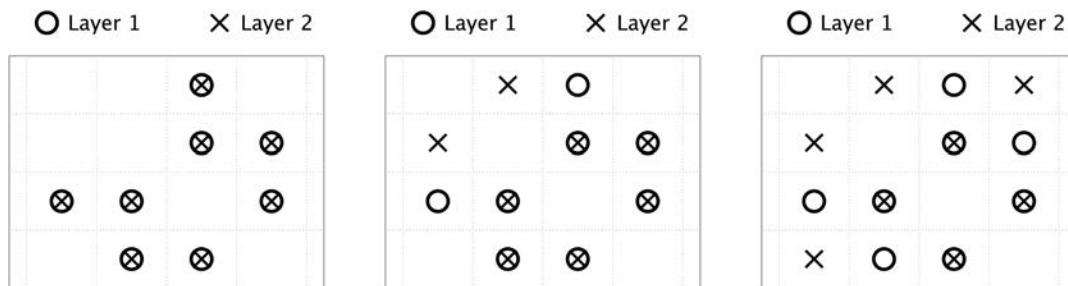


FIG. 15. Scenario 6. Adjacency matrices of each layer for three cases: (i, left panel) No rewiring takes place, and both layers are identical; (ii, middle panel) one edge has been rewired; (iii, right panel) two edges have been rewired. The transition matrices for each case are $T_{ij} = A_{ij}/k_i$, where k_i is the degree of node i .




FIG. 16. Scenario 6. Normalized Kullback-Leibler divergence $\mathcal{D}(m)$ between $X(t)$ and its Markovianized surrogate $Y(t)$ for a two-layer multiplex, where the second layer is a replica of the first layer in which a number of edges have been rewired. When no edges have been rewired, both replicas are identical, $X(t)$ is Markovian, and $\mathcal{D}(m) = 0 \forall m$; otherwise, the process is non-Markovian, and the algorithm detects multiplexity by finding $\mathcal{D}(m > 2) > 0$. Such detection improves as the topology of both layers is increasingly different.

p in the second layer. Multiplexity is detected when $\mathcal{D}(3) > 0$, and the larger $\mathcal{D}(3)$, the easier such detection is. In Fig. 17, we plot $\mathcal{D}(3)$ as a function of p . Dots are the result of an ensemble average over ER graph realizations. For an ensemble, we keep the number of rewired edges fixed and compute the effective average percentage of rewired edges (which fluctuates because each realization of an ER graph will have a different total number of edges).




FIG. 17. Scenario 6b. Detectability measure $\mathcal{D}(3)$ for a two-layer multiplex with $K = 10$ nodes per layer, where each layer is a connected Erdos-Renyi graph (link probability 0.65). The second layer is a replica of the first layer, where a percentage p of nodes have been rewired. Results have been averaged over 10^2 network realizations for each case.

APPENDIX C: MATHEMATICAL AND ALGORITHMIC FRAMEWORK FOR LAYER ESTIMATION

Here, we formalize the problem and provide a detailed derivation of the model, the probabilistic model selection scheme, and some possible algorithmic implementations of this scheme. Let us remark that the probabilistic framework described herein includes the case in which the path followed by the walker across the multiplex network cannot be observed exactly (i.e., there are observation errors) even if this is not addressed in the main text.

We use an argument-wise notation to denote PMFs and probability density functions (PDFs). If X and Y are discrete random variables (RVs), then $P(X)$ and $P(Y)$ are the PMF of X and the pmf of Y , respectively, and are possibly different. Similarly, $P(X, Y)$ and $P(X|Y)$ denote the joint PMF of the two RVs and the conditional PMF of X given Y , respectively. We use lowercase p for PDFs. If X and Y are continuous RVs, then $p(X)$ and $p(Y)$ are the corresponding densities, which are possibly different, and $p(X, Y)$ and $p(X|Y)$ denote the joint and conditional PDFs. We may have a PDF of a continuous RV X conditional on a discrete RV Y , $p(X|Y)$, as well as a PMF of Y given X , $P(Y|X)$. Most RVs are indicated with uppercase letters, e.g., X . If we need to denote a specific realization of the RV, then we use the same letter but in lowercase, e.g., $X = x$ or $Y = y$. Matrices and vectors are indicated with a boldface font, e.g., \mathbf{T} .

1. The model

We assume that a walker travels through a multiplex network, taking random moves between neighboring nodes and, occasionally, between layers. Let L denote the number of layers in the multiplex network, and let K be the number of nodes per layer. At discrete time t , the RV $X(t)$ denotes the in-layer walker position. Therefore, $X(t) \in \{0, \dots, K-1\}$, and $X(t) = k$ means that the particle is located at node number k at time t , irrespective of the layer. The RV $\ell(t)$ indicates the layer at time t , i.e., $\ell(t) \in \{1, 2, \dots, L\}$, and $\ell(t) = l$ means that the walker is found in layer l at time t . The state of the walker, therefore, is given by the 2×1 vector $\mathbf{Z}(t) = [X(t), \ell(t)]^\top$.

At each time step, the walker may jump across layers. This motion is assumed to be Markov, and hence it can be characterized by an $L \times L$ stochastic transition matrix \mathbf{R}_L , where the entry R_{ij} , $(i, j) \in \{0, \dots, L-1\}^2$, represents the probability of moving from layer i to layer j . Subsequently, the particle diffuses within the new layer to one of its neighbors. Within each single layer, the motion of the walker is also assumed to be Markovian. Hence, in the l th layer, it is governed by a $K \times K$ transition matrix $\mathbf{T}^{(l)}$, such that $T_{ij}^{(l)}$ is the probability of a particle lying in layer l to diffuse from node i to node j . These probabilities are constant over time t . The complete Markov model is

characterized by the set of matrices $\{\mathbf{R}_L, \mathbf{T}^{(0)}, \dots, \mathbf{T}^{(L-1)}\}$, and we denote $\mathcal{T}_L = \{\mathbf{T}^{(l)}\}_{l=0}^{L-1}$ in the sequel for convenience.

We can think of this model as a discrete-time state-space dynamical system, where the state variables at time t are

$$\mathbf{X}(t) = \{\mathbf{R}_L, \mathcal{T}_L, X(t), \ell(t)\}, \quad t = 0, 1, 2, \dots, \quad (\text{C1})$$

and we assume there is a sequence of observations $\{Y(t)\}_{t \geq 1}$ taking values in the set of node labels $\{0, 1, \dots, K-1\}$. In the main text, we assume that the observations are exact and, therefore, $Y(t) = X(t)$. However, we can handle a more general class of models in which $Y(t)$ is a RV with conditional PMF $P(Y(t)|X(t))$ (independently of the current or past layers). If the observation is exact, then

$$P(Y(t) = j | X(t) = i) = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise.} \end{cases}$$

However, the proposed model (and related numerical methods) admit the cases in which observation errors may occur, and hence, $P(Y(t)|X(t))$ is a nondegenerate PMF. We assume there are known and independent *a priori* PMFs for the node and layer at time $t = 0$, $P_0(X(0), \ell(0)) = P_0(X(0))P_0(\ell(0))$. In practical problems, the parameters \mathcal{T}_L and \mathbf{R}_L are unknown, and we also endow them with prior PDFs $p_0(\mathcal{T}_L, \mathbf{R}_L) = p_0(\mathcal{T}_L)p_0(\mathbf{R}_L)$ w.r.t. a suitable reference measure $\mu(d\mathcal{T}_L \times d\mathbf{R}_L)$. Most often, and for a general scenario, μ can be the Lebesgue measure on $\mathbb{R}^{L \times K \times K} \times \mathbb{R}^{L \times L}$, but other choices may be possible if we wish to impose constraints on \mathcal{T}_L and \mathbf{R}_L . For the case of the network with ring-shaped layers in the main text, μ can be reduced to the Lebesgue measure on \mathbb{R}^{L+1} .

2. Bayesian model selection

Assume that we have collected a sequence of N observations, which we now label

$$Y(1:N) = \{Y(1), Y(2), \dots, Y(N)\} \in \{0, 1, \dots, K-1\}^N.$$

We wish to make a decision as to *what model* is the best fit for that sequence. We adopt the view that two models are different when they have a different number of layers; hence, if model A has L layers and model B has L' layers, $A = B \Leftrightarrow L = L'$. A convenient way to tackle this problem is to model the total number of layers L as a RV, in such a way that each possible value of L corresponds to a different model. If we define

- (i) a prior probability mass function for L , say, $P_0(L)$, for $L \in \{1, 2, \dots, L_+\}$, where L_+ is the maximum admissible number of layers, and

- (ii) a likelihood function

$$P(Y(1:N)|L) = \int P(Y(1:N)|\mathcal{T}_L, \mathbf{R}_L) p_0(\mathcal{T}_L) \times p_0(\mathbf{R}_L) \mu(d\mathcal{T}_L \times d\mathbf{R}_L), \quad (\text{C2})$$

then we can aim at computing the posterior PMF of the number of layers,

$$P(L|Y(1:N)) \propto P(Y(1:N)|L)P_0(L),$$

and choose the model according to the MAP criterion

$$\begin{aligned} \hat{L}_{MAP} &= \arg \max_{L \in \{1, 2, \dots, L_+\}} P(L|Y(1:N)) \\ &= \arg \max_{L \in \{1, 2, \dots, L_+\}} P(Y(1:N)|L)P_0(L); \end{aligned} \quad (\text{C3})$$

i.e., we choose the value of L that is more probable given the available observations.

Expression (C3) yields the optimal solution to the problem of selection of the number of layers in the multiplex from a probabilistic Bayesian point of view. As discussed below, one can find alternatives to this approach in the literature on HMMs [54,56]; however, the latter suffer from a number of theoretical and practical limitations, and we strongly advocate for the Bayesian solution (C3).

3. Connections with hidden Markov model estimation theory

The problem of selecting the number of layers L in the multiplex model can be cast as one of selecting a HMM where the complete state is $\mathbf{Z}(t) = [X(t), \ell(t)]^\top$ and the transition from time t to time $t+1$ is governed by the (unknown) stochastic matrices $\mathbf{R}_L, \mathbf{T}^{(0)}, \dots, \mathbf{T}^{(L-1)}$. Let us adapt the notation a bit in order to make it closer to the classical HMM theory. We only have partial observations of the Markov chain, namely, the sequence $\mathcal{O} = Y(1:N)$, with “emission probabilities” $P(Y(t)|X(t))$, while the sequence of layer labels $\{\ell(t)\}_{t=1}^N$ remains unobserved. The goal is to estimate the total number of layers L from the observations \mathcal{O} .

The theory of HMMs has received considerable attention in the literature since the 1970s because of their application in a variety of fields, including speech processing, molecular biology, data compression, or artificial intelligence. The problem of fitting a HMM, i.e., estimating its unknown parameters, has been thoroughly researched [54]. The classical technique is the Baum-Welch algorithm, which is actually an instance of the expectation-maximization (EM) method [54,88]. Indeed, the general EM methodology, in several forms, is the standard approach to the problem of fitting HMMs, often combined with other techniques for its implementation, such as the Viterbi

algorithm, the forward-backward algorithm, or the Kalman smoother (see Ref. [56] for an excellent survey). Many of these techniques adopt the general form of a space-alternating EM algorithm, where the unobserved states and the unknown parameters are iteratively estimated, one at a time. The space-alternating generalized EM (SAGE) methodology was introduced in Ref. [55], and it provides a common framework for many current algorithms for fitting HMMs.

However, the estimation of the number of layers in the proposed multiplex scheme does not amount to HMM fitting. Modeling L as a random variable, in order to solve problem (C3), we aim at computing the model posterior probabilities given the available data \mathcal{O} , i.e.,

$$P(L|\mathcal{O}) \propto P(\mathcal{O}|L)P_0(L), \quad (\text{C4})$$

where $P_0(L)$, $L = 1, 2, \dots$, are the *a priori* probabilities we attribute to models with different numbers of layers (e.g., we may deem models with many layers less probable than simpler models with a few layers) and $P(\mathcal{O}|L)$ is the model likelihood. The latter is an integral with respect to the probability distribution of the matrix parameters $\mathcal{T}_L = \{\mathbf{T}^{(0)}, \dots, \mathbf{T}^{(L-1)}\}$ and \mathbf{R}_L for an L -layer system, namely,

$$P(\mathcal{O}|L) = \int P(\mathcal{O}|\mathcal{T}_L, \mathbf{R}_L) p_{0,T}(\mathcal{T}_L) p_{0,R}(\mathbf{R}_L) \times \mu(d\mathcal{T}_L \times d\mathbf{R}_L), \quad (\text{C5})$$

which is equivalent to Eq. (C2). Recall that $p_{0,T}(\mathcal{T}_L)$ and $p_{0,R}(\mathbf{R}_L)$ are *a priori* PDFs w.r.t. a reference measure μ . These PDFs can be chosen differently for different values of L . EM methods for HMM fitting are tools to address the problem of estimating \mathcal{T}_L and \mathbf{R}_L via the maximization of the parameter likelihood $P(\mathcal{O}|\mathcal{T}_L, \mathbf{R}_L)$ that appears in the integrand of Eq. (C5).

We see from Eqs. (C4) and (C5), however, that what we need is to be able to *integrate* the likelihood $P(\mathcal{O}|\mathcal{T}_L, \mathbf{R}_L)$, rather than *maximizing* it. Nevertheless, most methods in the HMM literature tackle the model selection problem (in our case, selection of the number of layers L) by computing estimates of the parameters via the EM method and then comparing the likelihoods of the optimized parameters [56,58]. In our setup, this means that, given two choices L_1 and L_2 , we would estimate $\hat{\mathcal{T}}_{L_1}$, $\hat{\mathbf{R}}_{L_1}$ and $\hat{\mathcal{T}}_{L_2}$, $\hat{\mathbf{R}}_{L_2}$ [using an EM scheme to maximize $P(\mathcal{O}|\mathcal{T}_{L_i}, \mathbf{R}_{L_i})$ for $i = 1, 2$] and then compare the likelihoods $P(\mathcal{O}|\hat{\mathcal{T}}_{L_1}, \hat{\mathbf{R}}_{L_1})$ and $P(\mathcal{O}|\hat{\mathcal{T}}_{L_2}, \hat{\mathbf{R}}_{L_2})$. This approach has several problems:

- (i) There is no guarantee that $\hat{\mathcal{T}}_{L_1}$, $\hat{\mathbf{R}}_{L_1}$ and $\hat{\mathcal{T}}_{L_2}$, $\hat{\mathbf{R}}_{L_2}$ are accurate estimates (e.g., they may be overfitted). It may well happen that, e.g., $\hat{\mathcal{T}}_{L_1}$, $\hat{\mathbf{R}}_{L_1}$ are poor estimates and, hence, $P(\mathcal{O}|\hat{\mathcal{T}}_{L_1}, \hat{\mathbf{R}}_{L_1}) < P(\mathcal{O}|\hat{\mathcal{T}}_{L_2}, \hat{\mathbf{R}}_{L_2})$, while $P(\mathcal{O}|L_1) > P(\mathcal{O}|L_2)$.

- (ii) The EM framework yields *local* optimization algorithms. Even if the EM scheme converges, it may yield a local maximizer of the likelihood for L_1 and, perhaps, a global maximizer for L_2 . In this case, we may again have $P(\mathcal{O}|\hat{\mathcal{T}}_{L_1}, \hat{\mathbf{R}}_{L_1}) < P(\mathcal{O}|\hat{\mathcal{T}}_{L_2}, \hat{\mathbf{R}}_{L_2})$, while $P(\mathcal{O}|L_1) > P(\mathcal{O}|L_2)$.
- (iii) Even if we manage to obtain accurate maximum likelihood estimates of \mathcal{T}_{L_1} , \mathbf{R}_{L_1} and \mathcal{T}_{L_2} , \mathbf{R}_{L_2} , there is no guarantee that $P(\mathcal{O}|\hat{\mathcal{T}}_{L_1}, \hat{\mathbf{R}}_{L_1}) < P(\mathcal{O}|\hat{\mathcal{T}}_{L_2}, \hat{\mathbf{R}}_{L_2})$ must imply $P(\mathcal{O}|L_1) < P(\mathcal{O}|L_2)$.

Many authors have aimed at mitigating these flaws by introducing different heuristics in the way the models to be fitted are chosen (typically, heuristics for merging and splitting candidate states, in our case candidate layers) and producing sophisticated EM parameter estimation algorithms. See Ref. [58] for examples. This approach does not attack the core of the problem, though.

Instead, Ref. [56] advocates Bayesian model selection as a framework to address problem (C4), which automatically handles overfitting (by imposing prior probability distributions on the parameters) and the comparison of models of different complexity [by integrating over the parameters as in Eq. (C5)]. In Ref. [56], the term used for the MAP model selection method of Eq. (C3) is, actually, Bayesian integration, which makes reference to the need to solve, or numerically approximate, the integral in Eq. (C5). Some candidate methods to tackle this computation include the following:

- (i) The Laplace approximation [89], which consists in searching the maximum of $P(\mathcal{O}|\hat{\mathcal{T}}_L, \hat{\mathbf{R}}_L)$ and then approximating the integrand $P(\mathcal{O}|\mathcal{T}_L, \mathbf{R}_L) \times p_{0,T}(\mathcal{T}_L) p_{0,R}(\mathbf{R}_L)$ by a Gaussian with the adequate covariance structure. This approach ignores the fact that $P(\mathcal{O}|\hat{\mathcal{T}}_L, \hat{\mathbf{R}}_L)$ is, in our case, multimodal.
- (ii) The variational Bayes method [59,90] is an approximation scheme that relies on the use of surrogate probability distributions for the parameters (which need to be analytically tractable) in order to design an EM method that tackles the maximization of the model likelihood $P(\mathcal{O}|L)$, i.e., the integral in Eq. (C5). It is a relatively “inexpensive” method in terms of computational cost, comparable to classical EM-based model fitting techniques. However, like any EM scheme, it performs a local optimization and does not guarantee an optimal solution.
- (iii) Deterministic integration of Eq. (C5) using either deterministic regular grids on the space of the parameters $\{\mathcal{T}_L, \mathbf{R}_L\}$ or specific cubature methods for some convenient family of functions [91]. While accurate, the complexity of these methods typically grows exponentially with the dimension of the parameters; hence, they can be prohibitive for larger-scale models. Examples for multiplex models with up to $L \geq 4$ layers are shown.

- (iv) Conventional Monte Carlo integration, which suffers from a similar complexity limitation. Classical Markov chain Monte Carlo (MCMC) samplers [92,93] could be well suited to solve integrals with respect to the posterior PDF $p(\mathcal{T}_L, \mathbf{R}_L | \mathcal{O})$; however, the integral in Eq. (C5) is actually the normalizing constant of this posterior, which turns out to be hard to estimate via MCMC, which limits its application to model selection, in general [93].

The classical alternative to MCMC in Monte Carlo integration is importance sampling (IS) [93]. While conventional IS suffers from a problem called weight degeneracy, which translates into poor scaling with the dimension of the parameters in the integral, recently, families of much more efficient *adaptive* IS schemes have been introduced [60,61,94,95]. These techniques yield estimates of the integral in Eq. (C5) in a simple way (unlike MCMC) and can potentially work in high dimensions [95].

Below, we present a detailed description of the nonlinear population Monte Carlo (PMC) scheme of Ref. [61] and show an example of model selection with up to 10 layers ($L \leq 10$). While conventional (and even state-of-the-art) importance samplers are based on the computation of weights of the form $w(z) \propto \{[p(z)]/[q(z)]\}$, where $p(z)$ is the target PDF and $q(z)$ is a proposal density, the key feature of the nonlinear PMC scheme is to compute transformed weights $\tilde{w}(z) \propto \phi\{[p(z)]/[q(z)]\}$, where $\phi(\cdot)$ is a nonlinear function, in order to reduce the variance of the weights [if z is a random variable, then $U = w(Z)$ is random as well]. This very simple transformation, if properly chosen, significantly improves the numerical stability of the algorithm when the dimension of Z grows, while preserving the convergence properties of conventional IS. The examples presented below, for the nonlinear PMC and a deterministic scheme based on regular grids, show that this Monte Carlo integration scheme can be as effective as a deterministic integrator with just a fraction of the running time.

4. Computation of the posterior probabilities via Monte Carlo integration

Let us return to the original notation where $Y(1:N)$ denotes the sequence of observations. In order to select the number of layers L in the multiplex according to the Bayesian criterion in Eq. (C3), we need the ability to evaluate the posterior probability

$$P(L|Y(1:N)) \propto P(Y(1:N)|L)P_0(L),$$

where the prior $P_0(L)$ is known (chosen by design) but the model likelihood $P(Y(1:N)|L)$ is an integral given by Eq. (C2), namely,

$$P(Y(1:N)|L) = \int P(Y(1:N)|\mathcal{T}_L, \mathbf{R}_L) p_0(\mathcal{T}_L) \times p_0(\mathbf{R}_L) \mu(d\mathcal{T}_L \times d\mathbf{R}_L). \quad (\text{C6})$$

Using the Bayes theorem, we realize that the integrand in Eq. (C6) is proportional to the posterior density of the parameters, given the observations $Y(1:N)$, i.e.,

$$p(\mathcal{T}_L, \mathbf{R}_L | Y(1:N)) \propto P(Y(1:N)|\mathcal{T}_L, \mathbf{R}_L) p_0(\mathcal{T}_L) p_0(\mathbf{R}_L). \quad (\text{C7})$$

Taken together, Eqs. (C6) and (C7) indicate that the model likelihood $P(Y(1:N)|L)$ is the normalization constant of the posterior PDF of the parameters, $p(\mathcal{T}_L, \mathbf{R}_L | Y(1:N))$. This normalization constant is often termed the model evidence in the Bayesian terminology.

An efficient way of computing the normalization constant of a target PDF via Monte Carlo integration is by using the importance sampling (IS) method.

a. Importance sampling in a nutshell

Let $p(z)$ be a target PDF that we can evaluate up to a normalization constant c ; i.e., we have the ability to compute

$$\tilde{p}(z) = c p(z)$$

pointwise, but c is unknown. The IS method [93] enables the estimation of c [actually, it enables the estimation of integrals of the form $\int f(z)p(z)dz$ in general, for any integrable test function f] by sampling from an alternative PDF, $q(z)$, often called the proposal density or importance function. We assume that $q(z)$ is chosen to satisfy

$$w(z) = \frac{\tilde{p}(z)}{q(z)} < \infty, \quad (\text{C8})$$

where $w(z)$ is the weight function. The inequality in Eq. (C8) typically implies, at least, that $q(z) > 0$ whenever $p(z) > 0$.

The basic IS algorithm proceeds as follows:

- (1) Draw M independent samples z^1, \dots, z^M from $q(z)$.
- (2) Compute weights

$$\tilde{w}^i = w(z^i) = \frac{\tilde{p}(z^i)}{q(z^i)}, \quad \text{for } i = 1, \dots, M.$$

- (3) Normalize the weights,

$$w^i = \frac{\tilde{w}^i}{\sum_{m=1}^M \tilde{w}^m}.$$

It is a straightforward application of the strong law of large numbers [93] to prove that

$$\lim_{M \rightarrow \infty} \sum_{i=1}^M w^i f(z^i) = \int f(z) p(z) dz \quad \text{almost surely (a.s.)}$$

for any square-integrable test function f .

However, the most relevant result for the purpose of this paper is that

$$\hat{c}^M = \frac{1}{M} \sum_{i=1}^M \tilde{w}^i \quad (\text{C9})$$

is an unbiased, consistent estimator of the normalization constant c . By the strong law of large numbers again,

$$\lim_{M \rightarrow \infty} \hat{c}^M = \int w(z) q(z) dz = \int \tilde{p}(z) dz = c \quad \text{a.s.} \quad (\text{C10})$$

since $\tilde{p}(z) = c p(z)$ and $p(z)$ is a PDF (hence, it integrates to 1).

In the Bayesian model selection problem at hand, the target non-normalized function is given by $\tilde{p}(z) \equiv P(Y(1:N) | \mathcal{T}_L, \mathbf{R}_L) p_0(\mathcal{T}_L) p_0(\mathbf{R}_L)$, which we can evaluate (as will be shown below), and $c = \int \tilde{p}(z) dz \equiv P(Y(1:N) | L)$ is the model likelihood.

b. Nonlinear population Monte Carlo technique

The main drawback of standard IS is that, whenever there is a significant mismatch between $\tilde{p}(z)$ and $q(z)$, the variance of the weights becomes very large. As a consequence, estimators converge very slowly (with M), and they become of little use. This issue is usually referred to as weight degeneracy [96]. It typically happens when the dimension of the random variable Z is large or when, simply, the target PDF is very narrow.

To mitigate degeneracy, a number of adaptive IS have been proposed, especially since the publication of Ref. [97]. Here, we resort to one such method that introduces a nonlinear transformation of the weights to control their variability and, hence, degeneracy. The technique is called the NPMC method, and it was originally proposed in Ref. [61]. It consists of J iterations, each involving the computation of both conventional importance weights (IW) and transformed importance weights (TIWs). The transformation is a clipping or truncation operation, denoted $\phi(\cdot, \cdot)$. For a set of M ordered IWs, $\tilde{w}^{i_1} > \tilde{w}^{i_2} > \dots > \tilde{w}^{i_M}$, we obtain a set of M TIWs, with clipping of order $M_c < \sqrt{M}$, as

$$\tilde{w}^i = \phi(i, \{\tilde{w}^m\}_{m=1}^M) = \begin{cases} \tilde{w}^{i_{M_c}} & \text{if } \tilde{w}^i \geq \tilde{w}^{i_{M_c}} \\ \tilde{w}^i & \text{otherwise.} \end{cases}$$

This operation truncates the bigger M_c weights. A general algorithm is outlined below, with M samples per iteration and an instrumental Markov sampling kernel $K(\cdot, \cdot | \mathcal{T}_L^i, \mathbf{R}_L^i)$ centered at $\mathcal{T}_L^i, \mathbf{R}_L^i$.

(1) Initialization.

- (a) Draw M independent samples $(\tilde{\mathcal{T}}_{L,0}^i, \tilde{\mathbf{R}}_{L,0}^i)$, $i = 1, \dots, M$, from the prior PDFs $p_0(\mathcal{T}_L)$ and $p_0(\mathbf{R}_L)$.
- (b) Compute non-normalized IWs, $\tilde{w}_0^i = P(Y(1:N) | \tilde{\mathcal{T}}_{L,0}^i, \tilde{\mathbf{R}}_{L,0}^i)$, $i = 1, \dots, M$.
- (c) Compute non-normalized TIWs, $\tilde{w}_0^i = \phi(i, \{\tilde{w}_0^m\}_{m=1}^M)$, $i = 1, \dots, M$.
- (d) Normalize the TIWs,

$$w_0^i = \frac{\tilde{w}_0^i}{\sum_{m=1}^M \tilde{w}_0^m}, \quad i = 1, \dots, M.$$

- (e) Resample M times the set $\{\tilde{\mathcal{T}}_{L,0}^i, \tilde{\mathbf{R}}_{L,0}^i\}_{i=1}^M$, with replacement and using the normalized TIWs as probability masses, to yield an unweighted sample set $\{\mathcal{T}_{L,0}^i, \mathbf{R}_{L,0}^i\}_{i=1}^M$.

(2) Iteration. For $j = 1:J$:

- (a) Draw M independent samples

$$(\tilde{\mathcal{T}}_{L,j}^i, \tilde{\mathbf{R}}_{L,j}^i) \sim K(\mathcal{T}_L, \mathbf{R}_L | \mathcal{T}_{L,j-1}^i, \mathbf{R}_{L,j-1}^i),$$

$$i = 1, \dots, M.$$

- (b) Compute non-normalized IWs,

$$\tilde{w}_j^i = \frac{P(Y(1:N) | \tilde{\mathcal{T}}_{L,j}^i, \tilde{\mathbf{R}}_{L,j}^i) p_0(\tilde{\mathcal{T}}_{L,j}^i) p_0(\tilde{\mathbf{R}}_{L,j}^i)}{K(\tilde{\mathcal{T}}_{L,j}^i, \tilde{\mathbf{R}}_{L,j}^i | \mathcal{T}_{L,j-1}^i, \mathbf{R}_{L,j-1}^i)},$$

$$i = 1, \dots, M.$$

- (c) Compute non-normalized TIWs, $\tilde{w}_j^i = \phi(i, \{\tilde{w}_j^m\}_{m=1}^M)$, $i = 1, \dots, M$.
- (d) Normalize the TIWs,

$$w_j^i = \frac{\tilde{w}_j^i}{\sum_{m=1}^M \tilde{w}_j^m}, \quad i = 1, \dots, M.$$

- (e) Resample M times the set $\{\tilde{\mathcal{T}}_{L,j}^i, \tilde{\mathbf{R}}_{L,j}^i\}_{i=1}^M$, with replacement and using the normalized TIWs as probability masses, to yield an unweighted sample set $\{\mathcal{T}_{L,j}^i, \mathbf{R}_{L,j}^i\}_{i=1}^M$.

After the J th iteration, we have the IS estimator of the model likelihood

$$P^M(Y(1:N) | L) = \frac{1}{M} \sum_{i=1}^M \tilde{w}_J^i.$$

This is a consistent estimator that converges with optimal Monte Carlo error rates. In particular, assuming that the IWs are bounded away from zero, Theorem 1 of Ref. [98] states that for any arbitrarily small $\epsilon < \frac{1}{2}$, there exists an a.s. finite random variable U^ϵ such that

$$|P^M(Y(1:N)|L) - P(Y(1:N)|L)| < \frac{U^\epsilon}{M^{\frac{1}{2}-\epsilon}} P(Y(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) \quad (\text{C12})$$

and, therefore,

$$\lim_{M \rightarrow \infty} P^M(Y(1:N)|L) = P(Y(1:N)|L) \quad \text{a.s.}$$

All that remains is to show that the parameter likelihood function $P(Y(1:N)|\mathcal{T}_{L,j}, \mathbf{R}_{L,j})$, and hence the IWs and the TIWs, can be evaluated exactly.

c. Exact calculation of the parameter likelihood $P(Y(1:N)|\mathcal{T}_L, \mathbf{R}_L)$

We start with the factorization

$$P(Y(1:N)|\mathcal{T}_L, \mathbf{R}_L) = \prod_{t=1}^N P(Y(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L), \quad (\text{C11})$$

where each factor

is a PMF that can be computed exactly using a recursive algorithm. We can write the function $P(Y(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L)$ as a marginal of the mass $P(Y(t), X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L)$, namely,

$$P(Y(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) = \sum_{X(t)} \sum_{\ell(t)} P(Y(t)|X(t)) P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L), \quad (\text{C13})$$

where we have used the fact that $P(Y(t)|\mathbf{X}(t)) = P(Y(t)|X(t))$. The conditional observation PMF $P(Y(t)|X(t))$ is one of the building blocks of the state-space model, so it can be readily evaluated. The predictive PMF $P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L)$, on the other hand, can be decomposed as

$$P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) = \sum_{X(t-1)} \sum_{\ell(t-1)} P(X(t)|X(t-1), \ell(t-1), \mathcal{T}_L) \times P(\ell(t)|\ell(t-1), \mathbf{R}_L) \times P(X(t-1), \ell(t-1)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L), \quad (\text{C14})$$

which depends on the filtering PMF at time $t-1$:

$$P(X(t-1), \ell(t-1)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) \propto P(Y(t-1)|X(t-1)) \times P(X(t-1), \ell(t-1)|Y(1:t-2), \mathcal{T}_L, \mathbf{R}_L). \quad (\text{C15})$$

Note that Eqs. (C13)–(C15) are recursively related. If we start from the prior PMF $P_0(X(0), \ell(0)) = P_0(X(0))P_0(\ell(0))$, then we can recursively compute the required sequence of PMFs for $t = 1, 2, \dots, T$ as

$$P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) = \sum_{X(t-1)} \sum_{\ell(t-1)} P(X(t)|X(t-1), \ell(t-1), \mathcal{T}_L) \times P(\ell(t)|\ell(t-1), \mathbf{R}_L) \times P(X(t-1), \ell(t-1)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L), \quad (\text{C16})$$

$$P(Y(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L) = \sum_{X(t)} \sum_{\ell(t)} P(Y(t)|X(t)) \times P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L), \quad (\text{C17})$$

$$P(X(t), \ell(t)|Y(1:N), \mathcal{T}_L, \mathbf{R}_L) \propto P(Y(t)|X(t)) P(X(t), \ell(t)|Y(1:t-1), \mathcal{T}_L, \mathbf{R}_L). \quad (\text{C18})$$

Substituting Eq. (C17) into Eq. (C11) for each $t = 1, \dots, N$ yields the parameter likelihood.

APPENDIX D: QUANTIFYING COMPLEXITY: ADDITIONAL RESULTS

In this section, we present some additional computer simulation results that complement those shown in the main text. For the simulations, we assume (the same as in the

main text) a simple ring topology for the layers of the multiplex, as well as a particular structure for the matrix \mathbf{R}_L , in such a way that it can be parametrized by a single scalar $r \in (0, 1)$. This is described in Appendix D 1. In Appendix D 2, we show how, given the ring topology of the layer and assuming exact observations $Y(t) = X(t)$, it is possible to construct a relatively simple deterministic approximation to the posterior probabilities $P(L|Y(1:N))$. A NPMC algorithm for this specific model is also detailed.

Finally, our numerical results are displayed and discussed in Appendix D 5.

1. Multiplex of ring-shaped layers

From now on, we assume a ring topology (cycle graph) for every layer. In this particular case, the particle may jump from one node to its neighbors (left or right on the ring) and from one layer to another layer. This motion is random, and we model it by way of the following probabilities,

$$R_{ij} = P(\ell(t) = j | \ell(t-1) = i), \quad (\text{D1})$$

$$\mathbf{T}^{(l)} = P(X(t) = X(t-1) + 1 \pmod{K} | \ell(t-1) = l), \quad (\text{D2})$$

where R_{ij} represents the probability of moving from layer i to layer j and $\mathbf{T}^{(l)}$ is the probability of moving rightwards in layer l , i.e., of jumping from node i to node $i + 1 \pmod{K}$. The probability of moving leftwards within layer l is, therefore,

$$1 - \mathbf{T}^{(l)} = 1 - P(X(t) = X(t-1) - 1 \pmod{K} | \ell(t-1) = l).$$

These rightwards-jump probabilities for layers $l = 0, 1, \dots, L-1$ are put together in the vector [99]

$$\mathbf{T}_L = \begin{bmatrix} \mathbf{T}^{(1)} \\ \vdots \\ \mathbf{T}^{(L)} \end{bmatrix}_{L \times 1} \quad (\text{D3})$$

In order to simplify the structure of the stochastic matrix \mathbf{R} , we assume that, at each time t , the walker may stay in the same layer (as in $t-1$) with probability $1-r$ ($r \in (0, 1)$) or may jump to a different layer with probability $[r/(L-1)]$. Hence, the transition matrix \mathbf{R}_L becomes

$$\mathbf{R}_L = \begin{bmatrix} 1-r & \frac{r}{L-1} & \cdots & \frac{r}{L-1} \\ \frac{r}{L-1} & 1-r & \cdots & \frac{r}{L-1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{r}{L-1} & \frac{r}{L-1} & \cdots & 1-r \end{bmatrix},$$

and there is a single unknown parameter, the probability to jump, r .

2. Approximation of the posterior model probabilities

In the main text, we have focused on scenarios where the node visited at time t can be observed exactly; i.e., $Y(t) = X(t)$ and

$$P(Y(t) = j | X(t) = i) = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise.} \end{cases}$$

This constraint simplifies some of the general equations in Appendix C 2. Specifically, the recursions given by Eqs. (C16)–(C18) for the evaluation of the PMF $P(Y(t) | Y(1:t-1), \mathbf{T}_L, r)$ (note that we have reduced \mathbf{R}_L and \mathcal{T}_L to the scalar r and vector \mathbf{T}_L) become

$$P(X(t) | X(1:t-1), \mathbf{T}_L, r) = \sum_{\ell(t-1)} P(X(t) | \ell(t-1), X(t-1), \mathbf{T}_L) \times P(\ell(t-1) | X(1:t-1), \mathbf{T}_L, r), \quad (\text{D4})$$

$$\begin{aligned} P(\ell(t) | X(1:t), \mathbf{T}_L, r) &= P(X(t) | \ell(t), X(1:t-1), \mathbf{T}_L, r) P(\ell(t) | X(1:t-1), \mathbf{T}_L, r) \\ &= P(X(t) | X(1:t-1), \mathbf{T}_L, r) \sum_{\ell(t-1)} P(\ell(t) | \ell(t-1), r) P(\ell(t-1) | X(1:t-1), \mathbf{T}_L, r). \end{aligned} \quad (\text{D5})$$

Equations (D4) and (D5) can be applied recursively, with initial conditions given by the prior PMFs $P_0(X(0))$ and $P_0(\ell(0))$. In particular, note that

$$\begin{aligned} P(X(1) | \mathbf{T}_L, r) &= \sum_{X(0)} \sum_{\ell(0)} P(X(1) | X(0), \ell(0), \mathbf{T}_L) \\ &\quad \times P(\ell(0)) P(X(0)), \end{aligned} \quad (\text{D6})$$

while

$$P(\ell(1) | \mathbf{T}_L, r) = \sum_{\ell(0)} P(\ell(1) | \ell(0), r) P(\ell(0)). \quad (\text{D7})$$

These priors are assumed to be known as part of the model specification (typically, they can be uniform PMFs, as they have been selected in computer experiments).

Given Eq. (D4), we can evaluate the likelihood of the parameters \mathbf{T}_L and r . In particular,

$$\begin{aligned} P(X(1:N) | \mathbf{T}_L, r) \\ = P(X(1) | \mathbf{T}_L, r) \prod_{t=1}^N P(X(t) | X(1:t-1), \mathbf{T}_L, r), \end{aligned} \quad (\text{D8})$$

with $P(X(1) | \mathbf{T}_L, r)$ computed as in Eq. (D6). Numerically, it is convenient to work with the log-likelihood

$$\begin{aligned} \log P(X(1:N)|\mathbf{T}_L, r) \\ = \log P(X(1)|\mathbf{T}_L, r) + \sum_{t=1}^N \log P(X(t)|X(1:t-1), \mathbf{T}_L, r) \end{aligned} \quad (\text{D9})$$

and transform it into natural units only when strictly needed.

The likelihood of the RV, L , given by the observation record $X(1:N)$ is, therefore, given by the integral

$$P(X(1:N)|L) = \int P(X(1:N)|\mathbf{T}_L, r) p_0(\mathbf{T}_L) p_0(r) d\mathbf{T}_L dr, \quad (\text{D10})$$

where we have assumed prior densities w.r.t. the Lebesgue measure. Indeed, in the absence of any prior knowledge, it is natural to choose $p_0(r)$ to be uniform over the interval $(0,1)$, i.e., $p_0(r) = 1$. The prior $p_0(\mathbf{T}_L)$ is important because it is used to penalize system configurations with two or more identical layers. Note that a multiplex with $L = 2$ and identical transition matrices $\mathbf{T}^{(1)} = \mathbf{T}^{(2)} = \tilde{\mathbf{T}}$ is equivalent to a monoplex with transition matrix $\tilde{\mathbf{T}}$. Similarly, any multiplex with L layers, out of which $L' \leq L$ have identical transition matrices, is fully equivalent to a reduced system with $L - L' + 1$ layers [100]. For our computer simulations, we have chosen

$$p_0(\mathbf{T}_L) \propto \min_{l \neq l'} |\mathbf{T}^{(l)} - \mathbf{T}^{(l')}|; \quad (\text{D11})$$

i.e., we penalize systems with at least two layers that have similar transition probabilities.

3. Deterministic integration

One conceptually simple way to approximate the integral in Eq. (D8) is

$$\begin{aligned} P(X(1:N)|L) \\ \approx \sum_{\mathbf{T}^{(0)} \in \mathcal{G}} \cdots \sum_{\mathbf{T}^{(L-1)} \in \mathcal{G}} \sum_{r \in \mathcal{G}} P(X(1:N)|\mathbf{T}_L, r) p_0(\mathbf{T}_L), \end{aligned} \quad (\text{D12})$$

where $\mathcal{G} = \{g_1, \dots, g_H\}$ is a grid of H points over the interval $(0,1)$, i.e., $0 < g_1 < g_2 < \dots < g_H < 1$. The proportionality constant of the prior $p_0(\mathbf{T}_L)$ can be approximated numerically as well (over the grid \mathcal{G}^L). We have selected a uniform prior on L ; i.e., assuming that $L \in \{1, 2, \dots, L_+\}$, we choose $P_0(L) = (1/L_+)$. In this way, the MAP selection criterion reduces to choosing the value of $L \in \{1, 2, \dots, L_+\}$ that maximizes the model likelihood $P(X(1:N)|L)$.

4. NPMC algorithm for Monte Carlo integration

Let us specify a practical NPMC algorithm for the multiplex composed of ring-shaped layers. The prior densities are $p_0(r) = U(0, 1)$ and $p_0(\mathbf{T}_L) \propto \min_{l \neq l'} |\mathbf{T}^{(l)} - \mathbf{T}^{(l')}|$. We approximate the normalization constant \hat{C} of the latter PDF by standard Monte Carlo integration [simulating the $\mathbf{T}^{(l)}$'s uniformly over $(0,1)$].

The sampling kernels $K(\cdot, \cdot | \mathbf{T}_L, r)$ are truncated Gaussian PDFs. Specifically, if $TN(x|\mu, \sigma^2, a, b)$ denotes the truncated Gaussian density

$$TN(x|\mu, \sigma^2, a, b) = \frac{\exp\{-\frac{1}{2\sigma^2}(x-\mu)^2\}}{\int_a^b \exp\{-\frac{1}{2\sigma^2}(u-\mu)^2\} du}, \quad \text{for } x \in (a, b),$$

then

$$K(\mathbf{T}_L, r | \check{\mathbf{T}}_L, \check{r}) = TN(r|\check{r}, \sigma^2, 0, 1) \prod_{l=1}^L TN(\mathbf{T}^{(l)} | \check{\mathbf{T}}^{(l)}, \sigma^2, 0, 1),$$

where $\sigma^2 = 0.005$.

The number of clipped weights is $M_c = \lfloor \sqrt{M} \rfloor$. Recall that the function $\phi(i, \{\tilde{w}_j^m\}_{m=1}^M)$ truncates the biggest M_c weights to be equal (and identical to the M_c th largest weight).

(1) *Initialization.*

- (a) Draw M independent samples $(\tilde{\mathbf{T}}_{L,0}^i, \tilde{r}_0^i)$, $i = 1, \dots, M$, from the prior PDFs $p_0(\mathbf{T}_L)$ and $p_0(r) = 1$ for $r \in (0, 1)$.
- (b) Compute non-normalized IWs, $\tilde{w}_0^i = P(X(1:N) | \tilde{\mathbf{T}}_{L,0}^i, \tilde{r}_0^i)$, $i = 1, \dots, M$.
- (c) Compute non-normalized TIWs, $\tilde{w}_0^i = \phi(i, \{\tilde{w}_0^m\}_{m=1}^M)$, $i = 1, \dots, M$.
- (d) Normalize the TIWs,

$$w_0^i = \frac{\tilde{w}_0^i}{\sum_{m=1}^M \tilde{w}_0^m}, \quad i = 1, \dots, M.$$

- (e) Resample M times the set $\{\tilde{\mathbf{T}}_{L,0}^i, \tilde{r}_0^i\}_{i=1}^M$, with replacement and using the normalized TIWs as probability masses, to yield an unweighted sample set $\{\mathbf{T}_{L,0}^i, r_0^i\}_{i=1}^M$.

(2) *Iteration.* For $j = 1:L$:

- (a) Draw M independent samples

$$\tilde{r}_j^i \sim TN(r|r_{j-1}^i, \sigma^2, 0, 1), \quad i = 1, \dots, M,$$

$$\tilde{\mathbf{T}}_j^{(l),i} \sim TN(\mathbf{T} | \mathbf{T}_{j-1}^{(l),i}, \sigma^2, 0, 1),$$

$$l = 1, \dots, L, \quad i = 1, \dots, M.$$

(b) Compute non-normalized IWs,

$$\tilde{w}_j^i = \frac{P(X(1:N)|\tilde{\mathbf{T}}_{L,j}^i, \tilde{r}_j^i) p_0(\tilde{\mathbf{T}}_{L,j}^i) p_0(\tilde{r}_j^i)}{TN(\tilde{r}_j^i|r_{j-1}^i, \sigma^2, 0, 1) \prod_{l=1}^L TN(\tilde{\mathbf{T}}_j^{(l),i}|\mathbf{T}_{j-1}^{(l),i}, \sigma^2, 0, 1)}, \quad i = 1, \dots, M.$$

(c) Compute non-normalized TIWs, $\tilde{w}_j^i = \phi(i, \{\tilde{w}_j^m\}_{m=1}^M)$, $i = 1, \dots, M$.

(d) Normalize the TIWs,

$$w_j^i = \frac{\tilde{w}_j^i}{\sum_{m=1}^M \tilde{w}_j^m}, \quad i = 1, \dots, M.$$

(e) Resample M times the set $\{\tilde{\mathbf{T}}_{L,j}^i, \tilde{r}_j^i\}_{i=1}^M$, with replacement and using the normalized TIWs as probability masses, to yield an unweighted sample set $\{\mathbf{T}_{L,j}^i, r_j^i\}_{i=1}^M$.

5. Numerical results

From deterministic integration to NPMC: Efficiently validating the full architecture.—We first consider a different example than the one shown in the main text, where now the true model from which the observations are generated has only $L = 2$ layers, and we aim at estimating this hidden model via posterior probability maximization. As for the case considered in the main text, the posterior probabilities are computed via the approximation in Appendix D 2, in particular, using a 19-point grid $\mathcal{G} = \{0.05, 0.10, \dots, 0.95\}$ for each unknown variable, which includes the probability to move rightwards, $\mathbf{T}^{(l)}$, over each layer and the additional parameter r , which is the probability to jump to a different layer at each iteration. For a maximum of $L = 4$ layers, this makes a grid of 19^5 nodes,

plus a 19^4 grid for $L = 3$, plus a 19^3 grid for $L = 2$, plus a 19-point grid for $L = 1$.

In the left panel of Fig. 18, we plot the posterior probability of four possible models with $L = 1, \dots, 4$ layers. It peaks at the correct value, although less emphatically than for the case shown in the main text. The reason for this is that, in this example, these probabilities are obtained using a record of only $N = 10^4$ observations, while we used a larger time series of $N = 2 \times 10^4$ observations in the case reported in the main text. Nevertheless, the method correctly classifies the underlying architecture, which suggests that the estimation protocol is robust against short series size, something extremely relevant for real-world applications. Note also that the posterior probability of $L = 1$ is zero up to the computer degree of accuracy. The observations $X(1:N)$, $N = 10^4$, were generated from a model with $1 - r = 0.89$, $\mathbf{T}^{(1)} = 0.14$, and $\mathbf{T}^{(2)} = 0.56$. This is not a node of the grid \mathcal{G}^3 , and yet the model is identified unequivocally.

In a second step, we explore this particular case with our NPMC algorithm, with $I = 10$ iterations, $M = 900$ Monte Carlo samples per iteration, and clipping parameter $M_c = \sqrt{M} = 30$, to approximate the posterior probabilities $P(\tilde{L}|X(1:N))$ for $\tilde{L} = 1, 2, 3$, and 4. We observe, in the right panel of Fig. 18, that the posterior probability peaks at $\tilde{L} = L = 2$, with similar values of the approximations. The runtime was approximately 68 minutes, while the computations for deterministic grid integration took

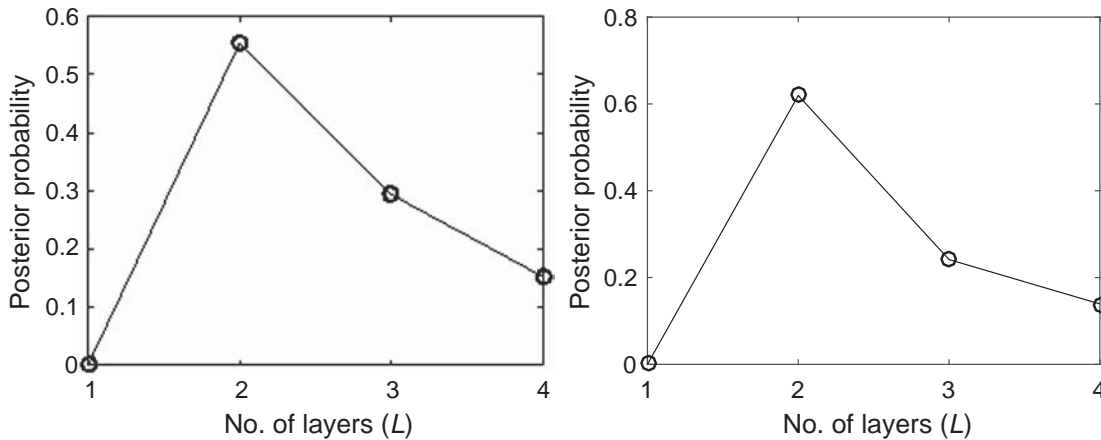


FIG. 18. (Left panel) Posterior probability $P(L|X(1:N))$ as a function of the number of layers L , computed from a trajectory of 10^4 time steps generated using a model with $L = 2$ layers and parameters $1 - r = 0.89$, $\mathbf{T}^{(1)} = 0.14$, and $\mathbf{T}^{(2)} = 0.56$. Computation is performed using deterministic grid integration. The algorithm clearly picks out the correct model $L = 2$. (Right panel) Similar results as for the left panel but using the NPMC algorithm instead of the deterministic integration, which is much faster and allows us to reconstruct the full architecture.




FIG. 19. Scatter plots of 100 Monte Carlo samples of $1 - r$ (top panel) and the interlayer transition probabilities $T^{(1)}$ and $T^{(2)}$ (bottom panel). They are all concentrated in a region close to the ground truth (note that there are two clouds in the bottom panel due to the symmetry $T^{(1)} \leftrightarrow T^{(2)}$).

approximately 6480 minutes (a reduction by a factor of approximately 95) on the same computer and with the same software.

Moreover, the NPMC also allows us to efficiently estimate, given $L = 2$, the most probable architecture of the multiplex, i.e., r , $T^{(1)}$, and $T^{(2)}$. The NPMC method generates a grid of the parameter space that is tighter in the regions with high probability measure; therefore, it is guaranteed that all the nodes in this grid yield parameters with high likelihood. If we view the problem as one of optimizing the parameter values for a given number of layers L , one can simply take this grid and choose the point that displays the highest probability density (which can be approximated numerically) or the highest likelihood (which can be computed exactly). The former is a MAP estimate, and the latter is a ML estimate. This ML estimate is indeed global. The method guarantees that the global maximum of

the likelihood is found (when the number of samples is sufficiently large), while other techniques, namely, EM algorithms, are local and can get stuck in local maxima of the likelihood no matter how many times they are iterated. All these computations (including the calculation of the parameter likelihoods) are implicitly carried out by our algorithm. Therefore, parameter estimation of the full architecture comes at no extra cost.

In Fig. 19, we depict the parameters sampled by the NPMC (100 Monte Carlo samples—note that this is not shown in the figure, as a resampling procedure yields many samples that are very similar). The importance sampling algorithm is emphatically concentrated in two small regions of the space, whose coordinates are indeed very close to the true parameters (highlighted in red). Incidentally, note that there are two clouds of points instead of one due to the symmetry between layers. This implies that the posterior PDF of the parameters is multimodal (which would create difficulties with the application of local search algorithms, such as the EM method). The histograms that approximate the (posterior) probability density function of these parameter probabilities are actually shown in Fig. 20. As can be seen, the histograms are concentrated close to the true values (the bimodal nature of the distributions in the case of \mathbf{T} is due to the symmetries of the model). The histograms in the figure can be smoothed out to obtain an approximation of the posterior PDF and, from it, MAP estimates of $T^{(1)}$, $T^{(2)}$, and r . Alternatively, one can obtain ML estimates at no extra cost since the likelihood of each configuration $(T^{(1)}, T^{(2)}, r)$ is computed by the importance sampling scheme.

Estimated architecture is a likelihood maximum.—As discussed above, we have also confirmed that the likelihood of the parameters, $P(X(1:N)|\mathbf{T}_L, r)$ [which is the key ingredient in the approximation of $P(L|X(1:N))$], is maximized at the values corresponding to the true model that generates the observations. In order to illustrate this fact further, we again consider the system under study in the main text (a model with $L = 3$ layers, with $1 - r = 0.84$ and $\mathbf{T}_L = [0.16, 0.76, 0.24]^T$), together with ten slightly perturbed models with $L = 3$ layers and parameters $1 - r' = 1 - r + \Delta$ and $\mathbf{T}'_L = [0.16 + \Delta', 0.76 + \Delta'', 0.24 + \Delta''']$, where $\Delta, \Delta', \Delta'', \Delta'''$ represent independent Gaussian




FIG. 20. Probability density of each parameter obtained from the 100 Monte Carlo samples of Fig. 19.




FIG. 21. Log-likelihood difference between a model with the true values and ten models with slightly perturbed values in the parameters, as a function of the size of the observation sequence. In every case, we find that the log-likelihood difference is larger than zero, meaning that the likelihood of the true model is always larger than the likelihoods of the models with perturbed parameters: The model with maximum *a posteriori* probability is indeed also the model with the correct parameters (not just a model with the correct number of layers), and this is an apparent global maximum. This result is robust as it gets more acute as we increase the size of the observed walker sequence.

perturbations of zero mean and variance 5×10^{-4} . As a function of the time series size [the observation window $X(1:N)$], we have computed the likelihood of each set of parameter values, which, in every case, we compare with the likelihood obtained for the true parameters. In Fig. 21, we plot the logarithm of the ratio of these likelihoods, where a positive value indicates that the ratio is larger than 1, i.e., that the likelihood of the true parameters is larger than the likelihood of the perturbed parameters. We systematically find that this is the case, and interestingly, we find that any small perturbation grows and generates a monotonically decreasing likelihood ratio as the size of the observed series increases.

Scalability.—Finally, we consider the issue of scalability and computational cost. All simulations have been carried out using Matlab R2016a running on an Apple iMac equipped with an Intel Core i7 processor, with 4 cores and 32 GB of RAM. For a first experiment, we have considered the same model and data as in Fig. 18, i.e., a system with $L = 2$ layers, parameters $1 - r = 0.89$, $T^{(1)} = 0.14$, and $T^{(2)} = 0.56$, and a series of $N = 10^4$ observations. As discussed previously, using NPMC instead of deterministic integration reduces the computing time by a factor of approximately 95. Therefore, equipped with the NPMC algorithm, we can tackle more complex systems. For the last computer experiment, we have generated a sequence of $N = 20 \times 10^3$ observations from a model with $L = 5$ layers, and parameters $T^{(1)} = 0.1$,




FIG. 22. Posterior probability $P(L|X(1:N))$ as a function of the number of layers L , computed from a trajectory of $N = 20 \times 10^3$ time steps generated using a model with $L = 5$ and parameters $T^{(1)} = 0.1$, $T^{(2)} = 0.2$, $T^{(3)} = 0.6$, $T^{(4)} = 0.7$, $T^{(5)} = 0.9$, and $1 - r = 0.85$. The algorithm clearly picks out the correct model $L = 5$. These results are obtained using the NPMC algorithm.

$T^{(2)} = 0.2$, $T^{(3)} = 0.6$, $T^{(4)} = 0.7$, $T^{(5)} = 0.9$, and $1 - r = 0.85$. We have applied the NPMC algorithm—with $J = 10$ iterations, $M = 1000$ Monte Carlo samples per iteration, and $M_c = 31$ —to approximate the posterior probabilities $P(\tilde{L}|X(1:N))$ for $\tilde{L} = 1, \dots, 10$. This simulation took around 370 minutes on the same computer. Figure 22 displays the results, where we observe that the correct model is still clearly identified.

APPENDIX E: BIOCHEMICAL NETWORKS: ADDITIONAL DETAILS FOR THE CASE OF RNA POLYMERASE

In Fig. 23, we plot (right panels) the five experimental traces of RNA polymerase I analyzed in this work. To make the inference problem more challenging, we systematically select the initial 10^4 time steps (about 10 seconds, as the signal is sampled at 1 kHz), where it is often not possible to visually distinguish the periods where elongation and backtracking are at play. Because of the inherent noise inside the cell, and as one can see in the figure, only for two cases (traces #2 and #3) are the two operating mechanisms—elongation and backtracking—observable, while the rest are hidden below the noise. Our methodologies (both the layer detection and layer estimation protocols) are successful at detecting and correctly inferring the presence of these two modes, as we find that the more likely number of hidden layers is $L = 2$. Intriguingly, if we select a period where apparently only one mode (backtracking) is at play, then our methodology again predicts multiplexity [$\mathcal{D}(3) > 0$] and $L = 2$ (see Fig. 24). This is a surprising result, which suggests that

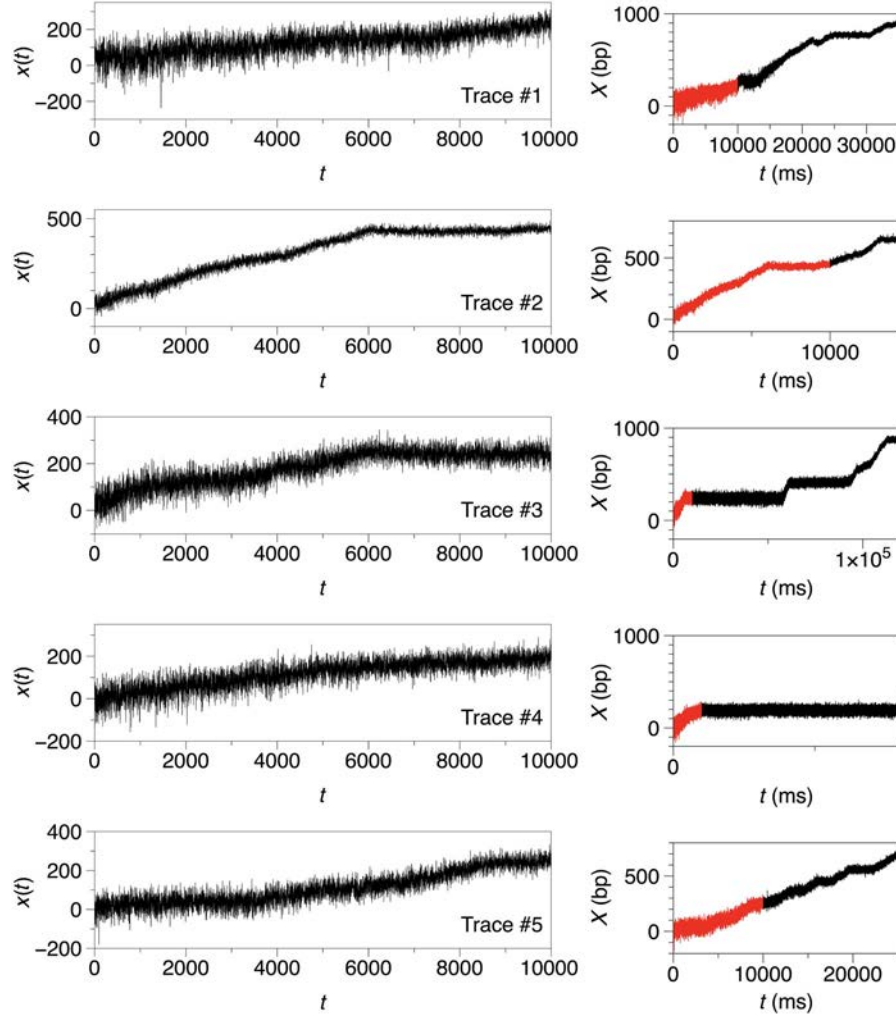


FIG. 23. Experimental traces of RNA polymerase I. In particular, and to make the inference problem more challenging, we only select the first 10^4 steps (about 10 seconds) for each trace. These parts are highlighted in red in the right panels, and plotted in the left panels. Interestingly, only for traces #2 and #3 are the two operating mechanisms—elongation and backtracking—observable and clearly discernible, while the rest are hidden below the noise.

either (i) short elongation times are hidden among a long backtracking pause or (ii), while in backtracking, stochastic fluctuations have a non-Markovian character. Both interpretations are in contrast with the state of the art and deserve further investigation.

APPENDIX F: PROOF OF THEOREM 1

It is well known that a higher-order Markov model ($h > 1$) can be converted into a first-order Markov system ($h = 1$) defined over a higher-dimensional state space (see, e.g., Ref. [101]). In the same vein, we show that every higher-order Markov model can be converted into an equivalent hidden jump-Markov model of the type discussed in Sec. II (which, from now on, we call the multiplex model) with a sufficiently large number of layers. The proof hence consists of two parts:

- (i) We first describe a constructive procedure to obtain a multiplex model from the transition probabilities of a generic Markov model of order h .
- (ii) We then prove that the probability of generating any sequence of length $h + 1$ is the same for the (order h) Markov model and the multiplex model.

1. Construction of the multiplex model

Recall that the Markov sequence takes values on the space $\mathbb{K} = \{1, \dots, K\}$ (this is the node set in the multiplex network) and is defined by the transition probabilities $P_{\mathbb{K}}^h(i_t | i_{t-h:t-1})$, $i_{t-h:t} \in \mathbb{K}^{h+1}$. We aim at constructing a model with $L = K^{h-1}$ layers, which implies that the random layer sequence $\ell(t)$ takes values on the set $\mathcal{L}_{h-1} := 1, \dots, K^{h-1}$. As a first step, we establish a bijection between the set of layers \mathcal{L}_{h-1} and the set of vectors \mathbb{K}^{h-1} .

To be specific, let us introduce the map $b: \mathbb{K}^{h-1} \rightarrow \mathcal{L}_{h-1}$, which associates every vector of states

$$\mathbf{v} = \begin{bmatrix} i_{h-1} \\ i_{h-2} \\ \vdots \\ i_1 \end{bmatrix} \in \mathbb{K}^{h-1}$$

with a unique layer $b(\mathbf{v}) := 1 + \sum_{s=1}^{h-1} (i_s - 1)K^{s-1} \in \mathcal{L}_{h-1}$. Intuitively, we interpret the vector

$$\mathbf{v} = \begin{bmatrix} i_{h-1} \\ i_{h-2} \\ \vdots \\ i_1 \end{bmatrix} \in \mathbb{K}^{h-1}$$

as a sequence of digits in a base K number system, and the function $b(\cdot)$ simply returns the evaluation of this number (plus one, for notational coherence). Since b is bijective, we can readily construct the inverse function $b^{-1}: \mathcal{L}_{h-1} \rightarrow \mathbb{K}^{h-1}$, which takes a layer number $\ell \in \mathcal{L}_{h-1}$ and returns its base- K representation (in vector form) $b^{-1}(\ell)$.

With these ingredients, we can now construct a multiplex model with the form of a Markov chain on the state space $\mathbb{K} \times \mathcal{L}_{h-1}$ described by the transition probabilities

$$P\left(\begin{bmatrix} X(t) \\ \ell(t) \end{bmatrix} = \begin{bmatrix} i_t \\ j_t \end{bmatrix} \middle| \begin{bmatrix} X(t-1) \\ \ell(t-1) \end{bmatrix} = \begin{bmatrix} i_{t-1} \\ j_{t-1} \end{bmatrix}\right) \\ = \mathbb{M}_* \left(\begin{array}{c|c} i_t & i_{t-1} \\ j_t & j_{t-1} \end{array} \right),$$

where $i_t, i_{t-1} \in \mathbb{K}$, $j_t, j_{t-1} \in \mathcal{L}_{h-1}$ and

$$\mathbb{M}_* \left(\begin{array}{c|c} i_t & i_{t-1} \\ j_t & j_{t-1} \end{array} \right) = \begin{cases} P_{\mathbb{K}}^h(i_t | i_{t-1:t-h}) & \text{if } b^{-1}(j_{t-1}) = \begin{pmatrix} i_{t-2} \\ \vdots \\ i_{t-h} \end{pmatrix} \text{ and } b^{-1}(j_t) = \begin{pmatrix} i_{t-1} \\ \vdots \\ i_{t-h+1} \end{pmatrix} \\ 0 & \text{otherwise.} \end{cases}$$

Note that we have simply encoded the history of states $i_{t-2:t-h}$ into the layer index j_{t-1} . For the system to behave “correctly,” we need to check that the transition from layer j_{t-1} to layer j_t is compatible with the sequence of states $i_{t-1:t-h+1}$ that has to be “recorded” for the next element of the chain. Unfortunately, the Markov chain described by the transition probabilities $\mathbb{M}_*(\cdot|\cdot)$ is semidegenerate. To be specific, the layer value $\ell(t) = j_t$ is deterministic, given $X(t-1) = i_{t-1}$ and $\ell(t-1) = j_{t-1}$. Indeed,

$$\text{if } j_{t-1} = b^{-1} \begin{pmatrix} i'_{h-1} \\ i'_{h-2} \\ \vdots \\ i'_1 \end{pmatrix} \text{ and } X(t-1) = i_{t-1}, \text{ then } j_t = b^{-1} \begin{pmatrix} i_{t-1} \\ i'_{h-1} \\ \vdots \\ i'_2 \end{pmatrix}$$

with probability 1. This determinism is not coherent with our discussion of multiplex models in Secs. I and II. Fortunately, this difficulty can be easily removed if we further extend the state space of the model.

Let us construct a (new but similar) Markov chain on the space $\mathbb{K} \times \mathcal{L}_h$, defined by the transition probabilities

$$P\left(\begin{bmatrix} X(t) \\ \ell(t) \end{bmatrix} = \begin{bmatrix} i_t \\ j_t \end{bmatrix} \middle| \begin{bmatrix} X(t-1) \\ \ell(t-1) \end{bmatrix} = \begin{bmatrix} i_{t-1} \\ j_{t-1} \end{bmatrix}\right) = \mathbb{M} \left(\begin{array}{c|c} i_t & i_{t-1} \\ j_t & j_{t-1} \end{array} \right), \quad (\text{F1})$$

where

$$\mathbb{M} \left(\begin{array}{c|c} i_t & i_{t-1} \\ j_t & j_{t-1} \end{array} \right) := P_{\mathbb{K}}^h(i_t | i_{t-1:t-h}) R_{j_{t-1}j_t} \quad (\text{F2})$$

and

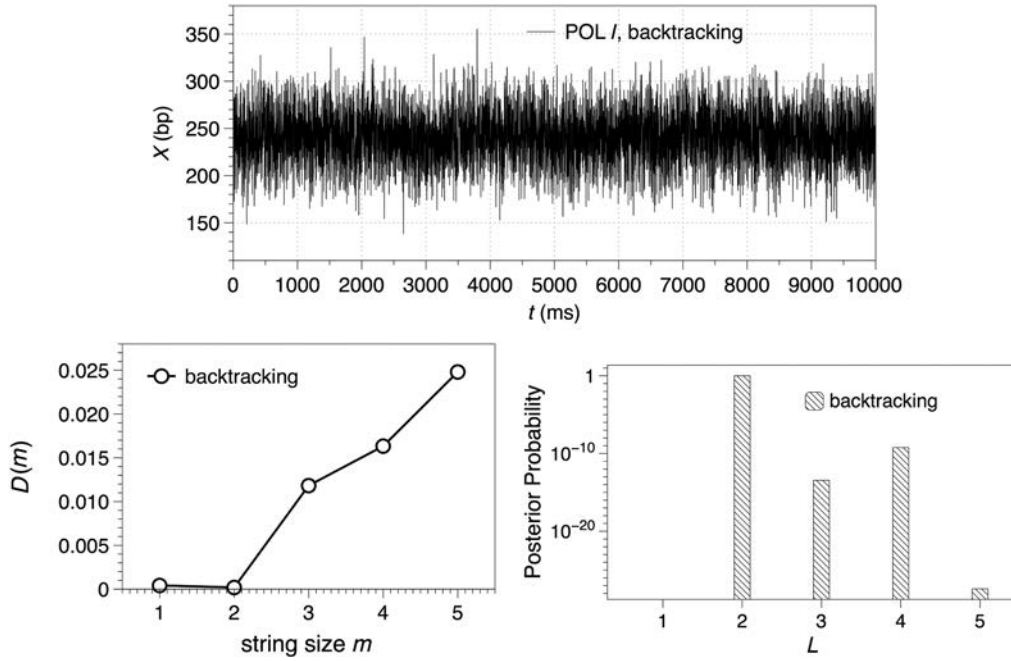


FIG. 24. (Top panel) Sample of Pol I trajectory of 10^4 steps in a backtracking pause. (Bottom-left panel) $\mathcal{D}(m)$ as a function of m . We find $\mathcal{D}(3) > 0$, suggesting that this series is non-Markovian. (Bottom-right panel) Posterior probability $P(L|X(1:N))$ as a function of the number of layers L , peaking at $L = 2$. The probability for $L = 1$ is below the ϵ of the machine, meaning it is virtually null.

$$R_{j_{t-1}j_t} := \begin{cases} \frac{1}{K} & \text{if } b^{-1}(j_t) = \begin{pmatrix} r \\ i_{t-1} \\ \vdots \\ i_{t-h+1} \end{pmatrix} \text{ and } b^{-1}(j_{t-1}) = \begin{pmatrix} s \\ i_{t-2} \\ \vdots \\ i_{t-h} \end{pmatrix} \text{ for some } s, r \in \mathbb{K} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{F3})$$

(Note that, in this case, the bijection b is a function $\mathbb{K}^h \rightarrow \mathcal{L}_h$, with the same interpretation as before.) Intuitively, we have now encoded the system memory (the sequence of states $i_{t-h:t-2}$) into a subset of K layers, instead of a single one. From each layer in this subset, there is probability $1/K$ to jump to a new layer belonging to the subset that encodes the sequence $i_{t-h+1:t-1}$. Comparing Eq. (F2) with our notation for multiplex models in Secs. I and II, we can readily identify the matrix \mathbf{R}_L of transition probabilities between layers, while $T_{i_{t-1}i_t}^{(j_{t-1})} = P_{\mathbb{K}}^h(i_t|i_{t-1:t-h})$ (note that j_{t-1} uniquely identifies the sequence $i_{t-2:t-h}$).

2. Identity of transition probabilities

What remains to be shown is that the conditional probabilities

$$P(X(t) = i_t | X(t-h:t-1) = i_{t-h:t-1})$$

generated by the probability function $P_{\mathbb{K}}^h(\cdot|\cdot)$ (for the order- h Markov model) and the Markov kernel $\mathbb{M}(\cdot|\cdot)$ in Eq. (F2) (for the multiplex model) coincide.

On one hand, if $\{X(t)\}_{t \geq 0}$ is an order- h Markov sequence, then

$$P(X(t) = i_t | X(t-h:t-1) = i_{t-h:t-1}) = P_{\mathbb{K}}^h(i_t | i_{t-h:t-1}) \quad (\text{F4})$$

by definition. On the other hand, if the sequence $\{X(t)\}_{t \geq 0}$ is generated by the multiplex model on $\mathbb{K} \times \mathcal{L}_h$ defined by \mathbb{M} , we have

$$\begin{aligned}
P(X(t) = i_t | X(t-h:t-1) = i_{t-h:t-1}) \\
&= \sum_{j_t \in \mathcal{L}_h} \sum_{j_{t-1} \in \mathcal{L}_h} P(X(t) = i_t, \ell(t) = j_t, \ell(t-1) = j_{t-1} | X(t-h:t-1) = i_{t-h:t-1}) \\
&= \sum_{j_t \in \mathcal{L}_h} \sum_{j_{t-1} \in \mathcal{L}_h} P(X(t) = i_t, \ell(t) = j_t | \ell(t-1) = j_{t-1}, X(t-h:t-1) = i_{t-h:t-1}) \\
&\quad \times P(\ell(t-1) = j_{t-1} | X(t-h:t-1) = i_{t-h:t-1}) \\
&= \sum_{j_t \in \mathcal{L}_h} \sum_{j_{t-1} \in \mathcal{L}_h} \mathbb{M} \left(\begin{matrix} i_t \\ j_t \end{matrix} \middle| \begin{matrix} i_{t-1} \\ j_{t-1} \end{matrix} \right) P(\ell(t-1) = j_{t-1} | X(t-h:t-1) = i_{t-h:t-1}), \tag{F5}
\end{aligned}$$

where the first equality follows from the theorem of total probabilities, the second equality is obtained from the definition of conditional probability, and the third identity follows from the definition of the multiplex model on $\mathbb{K} \times \mathcal{L}_h$. Moreover, for the multiplex model constructed in Appendix F 1,

$$P(\ell(t-1) = j_{t-1} | X(t-h:t-1) = i_{t-h:t-1}) = \begin{cases} \frac{1}{K} & \text{if } b^{-1}(j_{t-1}) = \begin{pmatrix} s \\ i_{t-2} \\ \vdots \\ i_{t-h} \end{pmatrix}, \quad s \in \mathbb{K} \\ 0 & \text{otherwise.} \end{cases} \tag{F6}$$

Hence, substituting Eqs. (F6) and (F2) into Eq. (F5) yields

$$P(X(t) = i_t | X(t-1:t-h) = i_{t-1:t-h}) = \frac{1}{K} \sum_{j_t \in \mathcal{L}_h} \sum_{j_{t-1} \in S_{b^{-1}}(i_{t-h:t-2})} P_{\mathbb{K}}^h(i_t | i_{t-1:t-h}) R_{j_{t-1}j_t}, \tag{F7}$$

where

$$S_{b^{-1}}(i_{t-h:t-2}) := \left\{ j \in \mathcal{L}_h : b^{-1}(j) = \begin{bmatrix} s \\ i_{t-2} \\ \vdots \\ i_{t-h} \end{bmatrix}, s \in \mathbb{K} \right\}.$$

However, from the definition of $R_{j_{t-1}j_t}$ in Eq. (F3), we can further reduce Eq. (F7) and obtain

$$P(X(t) = i_t | X(t-h:t-1) = i_{t-h:t-1}) = \frac{1}{K} P_{\mathbb{K}}^h(i_t | i_{t-1:t-h}) \sum_{j_t \in S_{b^{-1}}(i_{t-h+1:t-1})} \sum_{j_{t-1} \in S_{b^{-1}}(i_{t-h:t-2})} R_{j_{t-1}j_t}, \tag{F8}$$

and, for each term in the sum of Eq. (F8), we have $R_{j_{t-1}j_t} = (1/K)$ [from the definition in Eq. (F3)]. Because of the construction of the sets $S_{b^{-1}}(i_{t-h+1:t-1})$ and $S_{b^{-1}}(i_{t-h:t-2})$, there are exactly K^2 terms in the sum of Eq. (F8), which readily yields

$$P(X(t) = i_t | X(t-1:t-h) = i_{t-1:t-h}) = P_{\mathbb{K}}^h(i_t | i_{t-h:t-1})$$

and completes the proof. \square

APPENDIX G: PROOF OF THEOREM 2

Theorem 1 guarantees that any order- h Markov sequence on \mathbb{K} can be represented exactly by a multiplex model with $L = K^h$ layers. Hence, for every random sequence $\{X_h(t)\}_{t \geq 0}$, we have an equivalent sequence $\{X_L(t)\}_{t \geq 0}$ with $L = K^h$ layers and transition probabilities

$$R_{ij} = P(\ell(t) = j | \ell(t-1) = i),$$

$i, j \in \{1, \dots, L\}$ for the layers, and

$$T_{is}^{(j)} = P(X_L(t) = s | X_L(t-1) = i, \ell(t-1) = j),$$

for $i, s \in \{1, \dots, K\}$, $j \in \{1, \dots, L\}$. Therefore, every Markov- ∞ sequence $\{X(t)\}_{t \geq 0}$ can also be represented by a sequence of random processes $\{X_L(t)\}_{t \geq 0}$ described by suitably constructed multiplex models, i.e.,

$$\lim_{L \rightarrow \infty} X_L \stackrel{d}{=} X,$$

where $\stackrel{d}{=}$ denotes equality in distribution. What we need to prove is that the sequence of multiplex models, described by the pairs $(\mathbf{T}_L, \mathbf{R}_L)$, $L = K, K^2, K^3, \dots$, yields the continuous-layer model described by a proper probability measure \mathcal{M}_0 , a Markov kernel \mathcal{M} on $\mathcal{B}([0, 1])$, and a family of transition matrices $\mathbf{T}(y)$, $y \in [0, 1]$, when $L \rightarrow \infty$.

Let us note that every transition probability in the multiplex model can be specified by a simple function [102] defined over subsets of $[0, 1]$. To be specific, we construct

$$\hat{T}_{is}^L(y) := \sum_{\ell=1}^L \mathbb{1}_{[\frac{\ell-1}{L}, \frac{\ell}{L})}(y) T_{is}^{(\ell)}, \quad (\text{G1})$$

where

$$\mathbb{1}_A(y) = \begin{cases} 1 & \text{if } y \in A \\ 0 & \text{otherwise} \end{cases} \quad (\text{G2})$$

is the set-indicator function,

$$\hat{\mathcal{M}}^L(y|y') = \sum_{\ell=1}^L \sum_{j=1}^L \mathbb{1}_{[\frac{\ell-1}{L}, \frac{\ell}{L})}(y') \mathbb{1}_{[\frac{\ell-1}{L}, \frac{\ell}{L})}(y) R_{j\ell} \quad (\text{G3})$$

and

$$\hat{\mathcal{M}}_0^L(y) = \sum_{\ell=1}^L \mathbb{1}_{[\frac{\ell-1}{L}, \frac{\ell}{L})}(y) P(\ell(0) = \ell). \quad (\text{G4})$$

The limits $\lim_{L \rightarrow \infty} \hat{T}_{is}^L(y)$, $\lim_{L \rightarrow \infty} \hat{\mathcal{M}}^L(y|y')$ and $\lim_{L \rightarrow \infty} \hat{\mathcal{M}}_0^L(y)$ are well defined [because $X_L(t)$ yields

the same transition probabilities as $X_h(t)$, for $L = K^h$, and $\lim_{h \rightarrow \infty} X_h \stackrel{d}{=} X$ uniformly over t and \mathbb{K}^{h+1}]. Therefore, from Ref. [102] (see Theorem 5 in Chap. 8), the functions

$$\begin{aligned} T_{is}(y) &:= \lim_{L \rightarrow \infty} \hat{T}_{is}^L(y), \\ \bar{\mathcal{M}}(y|y') &:= \lim_{L \rightarrow \infty} \hat{\mathcal{M}}^L(y|y'), \quad \text{and} \\ \bar{\mathcal{M}}_0(y) &:= \lim_{L \rightarrow \infty} \hat{\mathcal{M}}_0^L(y) \end{aligned} \quad (\text{G5})$$

exist, and they are measurable w.r.t. the Lebesgue measure. In particular, $\bar{\mathcal{M}}(y|y')$ and $\bar{\mathcal{M}}_0(y)$ are probability density functions w.r.t. the Lebesgue measure, and we obtain $\mathcal{M}(dy|y') = \bar{\mathcal{M}}(y|y')dy$ and $\mathcal{M}_0(dy) = \bar{\mathcal{M}}_0(y)dy$. \square

-
- [1] M. Newman, *Networks: An Introduction* (Oxford University Press, New York, 2010).
 - [2] A.-L. Barabási, *Network Science* (Cambridge University Press, Cambridge, England, 2016).
 - [3] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, *Epidemic Processes in Complex Networks*, *Rev. Mod. Phys.* **87**, 925 (2015).
 - [4] A. Arenas, A. Díaz-Guilera, J. Kurths, Y. Moreno, and C. Zhou, *Synchronization in Complex Networks*, *Phys. Rep.* **469**, 93 (2008).
 - [5] C. Castellano, S. Fortunato, and V. Loreto, *Statistical Physics of Social Dynamics*, *Rev. Mod. Phys.* **81**, 591 (2009).
 - [6] S. Boccaletti, G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin, *The Structure and Dynamics of Multilayer Networks*, *Phys. Rep.* **544**, 1 (2014).
 - [7] M. Kivela, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, *Multilayer Networks*, *J. Complex Netw.* **2**, 203 (2014).
 - [8] F. Battiston, V. Nicosia, and V. Latora, *The New Challenges of Multiplex Networks: Measures and Models*, *Eur. Phys. J. Spec. Top.* **226**, 401 (2017).
 - [9] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, *Catastrophic Cascade of Failures in Interdependent Networks*, *Nature (London)* **464**, 1025 (2010).
 - [10] J. Gao, S. V. Buldyrev, H. E. Stanley, and S. Havlin, *Networks Formed from Interdependent Networks*, *Nat. Phys.* **8**, 40 (2012).
 - [11] S.-W. Son, G. Bizhani, C. Christensen, P. Grassberger, and M. Paczuski, *Percolation Theory on Interdependent Networks Based on Epidemic Spreading*, *Europhys. Lett.* **97**, 16006 (2012).
 - [12] G. Bianconi and S. N. Dorogovtsev, *Multiple Percolation Transitions in a Configuration Model of a Network of Networks*, *Phys. Rev. E* **89**, 062814 (2014).
 - [13] F. Radicchi and A. Arenas, *Abrupt Transition in the Structural Formation of Interconnected Networks*, *Nat. Phys.* **9**, 717 (2013).
 - [14] F. Radicchi, *Driving Interconnected Networks to Super-criticality*, *Phys. Rev. X* **4**, 021014 (2014).

- [15] S. Gómez, A. Díaz-Guilera, J. Gómez-Gardeñes, C. J. Pérez-Vicente, Y. Moreno, and A. Arenas, *Diffusion Dynamics on Multiplex Networks*, *Phys. Rev. Lett.* **110**, 028701 (2013).
- [16] M. De Domenico, A. Sole-Ribalta, S. Gomez, and A. Arenas, *Navigability of Interconnected Networks under Random Failures*, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 8351 (2014).
- [17] A. Sole-Ribalta, S. Gomez, and A. Arenas, *Congestion Induced by the Structure of Multiplex Networks*, *Phys. Rev. Lett.* **116**, 108701 (2016).
- [18] F. Battiston, V. Nicosia, and V. Latora, *Efficient Exploration of Multiplex Networks*, *New J. Phys.* **18**, 043035 (2016).
- [19] A. Saumell-Mendiola, M. Á. Serrano, and M. Boguñá, *Epidemic Spreading on Interconnected Networks*, *Phys. Rev. E* **86**, 026106 (2012).
- [20] C. Granell, S. Gomez, and A. Arenas, *Dynamical Interplay between Awareness and Epidemic Spreading in Multiplex Networks*, *Phys. Rev. Lett.* **111**, 128701 (2013).
- [21] C. Buono, L. G. Alvarez-Zuzek, P. A. Macri, and L. A. Braunstein, *Epidemics in Partially Overlapped Multiplex Networks*, *PLoS One* **9**, e92200 (2014).
- [22] J. Sanz, Ch.-Y. Xia, S. Meloni, and Y. Moreno, *Dynamics of Interacting Diseases*, *Phys. Rev. X* **4**, 041005 (2014).
- [23] J. Gómez-Gardeñes, I. Reinares, A. Arenas, and L. M. Floría, *Evolution of Cooperation in Multiplex Networks*, *Sci. Rep.* **2**, 620 (2012).
- [24] Z. Wang, A. Szolnoki, and M. Perc, *Rewarding Evolutionary Fitness with Links between Populations Promotes Cooperation*, *J. Theor. Biol.* **349**, 50 (2014).
- [25] J. T. Matamalas, J. Poncela-Casasnovas, S. Gómez, and A. Arenas, *Strategical Incoherence Regulates Cooperation in Social Dilemmas on Multiplex Networks*, *Sci. Rep.* **5**, 9519 (2015).
- [26] C. Del Genio, J. Gómez-Gardeñes, I. Bonamassa, and S. Boccaletti, *Synchronization in Networks with Multiple Interaction Layers*, *Sci. Adv.* (unpublished).
- [27] M. Diakonova, M. San Miguel, and V. M. Eguíluz, *Absorbing and Shattered Fragmentation Transitions in Multilayer Coevolution*, *Phys. Rev. E* **89**, 062818 (2014).
- [28] M. Diakonova, V. Nicosia, V. Latora, and M. San Miguel, *Irreducibility of Multilayer Network Dynamics: The Case of the Voter Model*, *New J. Phys.* **18**, 023010 (2016).
- [29] M. De Domenico, C. Granell, M. Porter, and A. Arenas, *The Physics of Spreading Processes in Multilayer Networks*, *Nat. Phys.* **12**, 901 (2016).
- [30] M. De Domenico, V. Nicosia, A. Arenas, and V. Latora, *Structural Reducibility of Multilayer Networks*, *Nat. Commun.* **6**, 6864 (2015).
- [31] M. Zanin, *Can We Neglect the Multi-layer Structure of Functional Networks?*, *Physica A (Amsterdam)* **430A**, 184 (2015).
- [32] M. Rosvall and C. T. Bergstrom, *Maps of Random Walks on Complex Networks Reveal Community Structure*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 1118 (2008).
- [33] T. P. Peixoto and M. Rosvall, *Modelling Sequences and Temporal Networks with Dynamic Community Detection*, *Nat. Commun.* **8**, 582 (2017).
- [34] P. J. Davis and P. Rabinowitz, *Methods of Numerical Integration* (Academic Press, New York, 1984).
- [35] N. Chopin, P. E. Jacob, and O. Papaspiliopoulos, *SMC2: An Efficient Algorithm for Sequential Analysis of State Space Models*, *J. R. Stat. Soc. Ser. B* **75**, 397 (2013).
- [36] M. Szell, R. Lambiotte, and S. Thurner, *Multirelational Organization of Large-Scale Social Networks in an Online World*, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 13636 (2010).
- [37] M. Barthélemy, *Spatial Networks*, *Phys. Rep.* **499**, 1 (2011).
- [38] A. Cardillo, J. Gómez-Gardeñes, M. Zanin, M. Romance, D. Papo, F. del Pozo, and S. Boccaletti, *Emergence of Network Features from Multiplexity*, *Sci. Rep.* **3**, 1344 (2013).
- [39] A. Garulli, S. Paoletti, and A. Vicino, *A Survey on Switched and Piecewise Affine System Identification*, *IFAC Proc. Vol.* **45**, 344 (2012).
- [40] R. Gallotti and M. Barthelemy, *Anatomy and Efficiency of Urban Multimodal Mobility*, *Sci. Rep.* **4**, 6911 (2015).
- [41] R. Gallotti, M. A. Porter, and M. Barthelemy, *Lost in Transportation: Information Measures and Cognitive Limits in Multilayer Navigation*, *Sci. Adv.* **2**, e1500445 (2016).
- [42] A. Ajdari and J. Prost, *Drift Induced by a Spatially Periodic Potential of Low Symmetry-Pulsed Dielectrophoresis*, *C.R. Acad. Sci. Ser. Gen. Ser. II* **315**, 13 (1992).
- [43] F. Julicher, A. Ajdari, and J. Prost, *Modeling Molecular Motors*, *Rev. Mod. Phys.* **69**, 1269 (1997).
- [44] G. P. Harmer and D. Abbott, *Losing Strategies Can Win by Parrondo's Paradox*, *Nature (London)* **402**, 864 (1999).
- [45] J. M. R. Parrondo and L. Dinis, *Brownian Motion and Gambling: From Ratchets to Paradoxical Games*, *Contemp. Phys.* **45**, 147 (2004).
- [46] M. V. Smoluchowski, *Experimentell nachweisbare, der üblichen thermodynamik widersprechende molekularephänomene*, *Phys. Z.* **13**, 1069 (1912).
- [47] R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics* (Addison-Wesley, Reading, MA, 1963), Vol. 1, Secs. 46.1–46.9.
- [48] In the definition above, time is considered discrete; however, a similar methodology can be applied for continuous time dynamics by sampling the continuous time series.
- [49] É. Roldán and J. M. R. Parrondo, *Entropy Production and Kullback-Leibler Divergence between Stationary Trajectories of Discrete Systems*, *Phys. Rev. E* **85**, 031129 (2012).
- [50] É. Roldán and J. M. R. Parrondo, *Estimating Dissipation from Single Stationary Trajectories*, *Phys. Rev. Lett.* **105**, 150607 (2010).
- [51] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley & Sons, New Jersey, 2006).
- [52] Even if $X(t)$ is Markovian, the matrix \mathbf{Q} cannot be computed exactly, but it can be estimated from the observations. Similarly, one can only accurately estimate the probabilities $P_X(Z_1, \dots, Z_m)$ for moderate values of m . Therefore, in practical scenarios, the decision rule on whether the underlying network is multiplex or not requires that we introduce an error threshold $\epsilon > 0$ such that $\mathcal{D}(3) > \epsilon$ implies multiplexity. Such a discrimination criterion may be formally described as a statistical test, and error bounds may be obtained under mild assumptions.
- [53] M. E. J. Newman and G. Reinert, *Estimating the Number of Communities in a Network*, *Phys. Rev. Lett.* **117**, 078301 (2016).

- [54] L. R. Rabiner, *A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition*, *Proc. IEEE* **77**, 257 (1989).
- [55] J. A. Fessler and A. O. Hero, *Space-Alternating Generalized Expectation-Maximization Algorithm*, *IEEE Trans. Signal Process.* **42**, 2664 (1994).
- [56] Z. Ghahramani, *An Introduction to Hidden Markov Models and Bayesian Networks*, *Int. J. Pattern Recognition Artificial Intelligence* **15**, 9 (2001).
- [57] E. Ozkan, F. Lindsten, C. Fritsche, and F. Gustafson, *Recursive Maximum Likelihood Identification of Jump Markov Nonlinear Systems*, *IEEE Trans. Signal Process.* **63**, 754 (2015).
- [58] S. M. Siddiqi, G. J. Geoffrey, and A. W. Moore, *Fast State Discovery for HMM Model Selection and Learning*, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics* (PMLR, 2007), p. 492, <http://proceedings.mlr.press/v2/siddiqi07a.html>.
- [59] C. A. McGrory and D. M. Titterton, *Variational Bayesian Analysis for Hidden Markov Models*, *Aust. New Zealand J. Stat.* **51**, 227 (2009).
- [60] O. Cappé, R. Douc, A. Guillin, J. M. Marin, and C. P. Robert, *Adaptive Importance Sampling in General Mixture Classes*, *Stat. Comput.* **18**, 447 (2008).
- [61] E. Koblenz and J. Míguez, *A Population Monte Carlo Scheme with Transformed Weights and Its Application to Stochastic Kinetic Models*, *Stat. Comput.* **25**, 407 (2015).
- [62] Note that a multiplex with $L = 2$ and $\mathbf{T}^{(1)} = \mathbf{T}^{(2)} = \mathbf{T}$ is equivalent to a monoplex with transition matrix \mathbf{T} .
- [63] A. Clauset, C. Moore, and M. E. J. Newman, *Hierarchical Structure and the Prediction of Missing Links in Networks*, *Nature (London)* **453**, 98 (2008).
- [64] R. Guimera and M. Sales-Pardo, *Missing and Spurious Interactions and the Reconstruction of Complex Networks*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 22073 (2009).
- [65] N. Koumakis, C. Maggib, and R. Di Leonardo, *Directed Transport of Active Particles over Asymmetric Energy Barriers*, *Soft Matter* **10**, 5695 (2014).
- [66] C. Bechinger, R. Di Leonardo, H. Lowen, C. Reichhardt, G. Volpe, and G. Volpe, *Active Particles in Complex and Crowded Environments*, *Rev. Mod. Phys.* **88**, 045006 (2016).
- [67] M. C. Gonzalez, C. A. Hidalgo, and A. L. Barabasi, *Understanding Individual Human Mobility Patterns*, *Nature (London)* **453**, 779 (2008).
- [68] C. M. Schneider, V. Belik, T. Couronne, Z. Smoreda, and M. C. Gonzalez, *Unravelling Daily Human Mobility Motifs*, *J. R. Soc. Interface* **10**, 84 (2013).
- [69] A. Lima, R. Stanojevic, D. Papagiannaki, P. Rodriguez, and M. C. Gonzalez, *Understanding Individual Routine Behavior*, *J. R. Soc. Interface* **13**, 116 (2016).
- [70] Z.-D. Zhao, S.-M. Cai, and Y. Lu, *Non-Markovian Character in Human Mobility: Online and Offline*, *Chaos* **25**, 063106 (2015).
- [71] D. Boyer, M. C. Crofoot, and P. D. Walsh, *Non-Random Walks in Monkeys and Humans*, *J. R. Soc. Interface* **9**, 70 (2011).
- [72] V. Salnikov, M. T. Schaub, and R. Lambiotte, *Using High-Order Markov Models to Reveal Flow-Based Communities in Networks*, *Sci. Rep.* **6**, 23194 (2016).
- [73] J. K. Jansen, P. L. Boveng, and J. L. Bengtson, *Foraging Modes of Chinstrap Penguins: Contrasts between Day and Night*, *Mar. Ecol. Prog. Ser.* **165**, 161 (1998).
- [74] M. Lanzendorf, *Mobility Styles and Travel Behavior: Application of a Lifestyle Approach to Leisure Travel*, *Transp. Res. Rec.* **1807**, 163 (2002).
- [75] M. Szell, R. Sinatra, G. Petri, S. Thurner, and V. Latora, *Understanding Mobility in a Social Petri Dish*, *Sci. Rep.* **2**, 457 (2012).
- [76] P. Cramer, K.-J. Armache, S. Baumli, S. Benkert, F. Brueckner, C. Buchen, G. E. Damsma, S. Dengl, S. R. Geiger, A. J. Jasiak *et al.*, *Structure of Eukaryotic RNA Polymerases*, *Annu. Rev. Biophys.* **37**, 337 (2008).
- [77] E. A. Galburt, S. W. Grill, A. Wiedmann, L. Lubkowska, J. Choy, E. Nogales, M. Kashlev, and C. J. Bustamante, *Backtracking Determines the Force Sensitivity of RNAP II in a Factor-Dependent Manner*, *Nature (London)* **446**, 820 (2007).
- [78] E. A. Abbondanzieri, W. J. Greenleaf, J. W. Shaevitz, R. Landick, and S. M. Block, *Direct Observation of Base-Pair Stepping by RNA Polymerase*, *Nature (London)* **438**, 460 (2005).
- [79] M. Dangkulwanich, T. Ishibashi, S. Liu, M. L. Kireeva, L. Lubkowska, M. Kashlev, and C. J. Bustamante, *Complete Dissection of Transcription Elongation Reveals Slow Translocation of RNA Polymerase II in a Linear Ratchet Mechanism*, *eLife* **2**, e00971 (2013).
- [80] A. Lisica, C. Engel, M. Jahnel, É. Roldán, E. A. Galburt, P. Cramer, and S. W. Grill, *Mechanisms of Backtrack Recovery by RNA Polymerases I and II*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 2946 (2016).
- [81] P. Gaspard, *Template-Directed Copolymerization, Random Walks along Disordered Tracks, and Fractals*, *Phys. Rev. Lett.* **117**, 238101 (2016).
- [82] T. Valles-Catala, F. A. Massucci, R. Guimera, and M. Sales-Pardo, *Multilayer Stochastic Block Models Reveal the Multilayer Structure of Complex Networks*, *Phys. Rev. X* **6**, 011036 (2016).
- [83] T. P. Peixoto, *Hierarchical Block Structures and High-Resolution Model Selection in Large Networks*, *Phys. Rev. X* **4**, 011047 (2014).
- [84] T. P. Peixoto, *Model Selection and Hypothesis Testing for Large-Scale Network Models with Overlapping Groups*, *Phys. Rev. X* **5**, 011033 (2015).
- [85] T. P. Peixoto and L. Gauvin, *Change Points, Memory and Epidemic Spreading in Temporal Networks*, [arXiv:1712.08948](https://arxiv.org/abs/1712.08948).
- [86] T. P. Peixoto, *Bayesian Stochastic Blockmodeling*, [arXiv:1705.10225](https://arxiv.org/abs/1705.10225).
- [87] C. C. Strelhoff and J. P. Crutchfield, *Bayesian Structural Inference for Hidden Processes*, *Phys. Rev. E* **89**, 042119 (2014).
- [88] A. P. Dempster, N. M. Laird, and D. B. Rubin, *Maximum Likelihood from Incomplete Data Via the EM Algorithm*, *J. R. Stat. Soc. Ser. B* **39**, 1 (1977).
- [89] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer-Verlag, New York, 2006).
- [90] S. Watanabe, Y. Minami, A. Nakamura, and N. Ueda, *Variational Bayesian Estimation and Clustering for Speech Recognition*, *IEEE Trans. Speech Audio Process.* **12**, 365 (2004).

- [91] S. Sobolev and V. Vaskevich, *The Theory of Cubature Formulas* (Springer, Dordrecht, 2013).
- [92] W. R. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Chapman and Hall/CRC, Boca Raton, 1996).
- [93] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, New York, 2004).
- [94] P. Del Moral, A. Doucet, and A. Jasra, *Sequential Monte Carlo Samplers*, *J. R. Stat. Soc. Ser. B* **68**, 411 (2006).
- [95] A. Beskos, D. Crisan, and A. Jasra, *On the Stability of Sequential Monte Carlo Methods in High Dimensions*, *Ann. Appl. Probab.* **24**, 1396 (2014).
- [96] A. Kong, J. S. Liu, and W. H. Wong, *Sequential Imputations and Bayesian Missing Data Problems*, *J. Am. Stat. Assoc.* **89**, 278 (1994).
- [97] O. Cappé, A. Gullin, J. M. Marin, and C. P. Robert, *Population Monte Carlo*, *J. Comput. Graph. Stat.* **13**, 907 (2004).
- [98] E. Koblenz, J. Miguez, M. A. Rodriguez, and A. M. Schmidt, *A Nonlinear Population Monte Carlo Scheme for the Bayesian Estimation of Parameters of α -Stable Distributions*, *Computational Statistics and Data Analysis* **95**, 57 (2016).
- [99] In the main text, we use \mathbf{T}_L to refer to the set of transition matrices $\mathcal{T}_L = \{\mathbf{T}^{(l)}\}_{l=0}^{L-1}$. For the multiplex network described, where each layer is a ring, the probabilities $\mathbf{T}^{(l)}$, $l = 0, \dots, L-1$, contain the same information as \mathcal{T}_L ; hence, \mathcal{T}_L and \mathbf{T}_L are equivalent.
- [100] Along these lines, also note that a multiplex with L ring-shaped layers and $r = \frac{1}{2}$ can be reduced to a monoplex with the probability to move rightwards, $T_{i,i+1} = (1/L) \sum_{l=0}^{L-1} \mathbf{T}^{(l)}$.
- [101] R. Chen, X. Wang, and J. S. Liu, *Adaptive Joint Detection and Decoding in Flat-Fading Channels via Mixture Kalman Filtering*, *IEEE Trans. Inf. Theory* **46**, 2079 (2000).
- [102] A. N. Kolmogorov and S. V. Fomi, *Introductory Real Analysis* (Dover Publications, New York, 1975).