

Compressing the Chronology of a Temporal Network with Graph Commutators

Andrea J. Allen^{1,2}, Christopher Moore³, and Laurent Hébert-Dufresne^{1,4,*}

¹Vermont Complex Systems Center, University of Vermont, Burlington, Vermont 05405, USA

²Applied Clinical Research Center, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania 19104, USA

³Santa Fe Institute, Santa Fe, New Mexico 87501, USA

⁴Department of Computer Science, University of Vermont, Burlington, Vermont 05405, USA

(Received 10 January 2023; revised 20 October 2023; accepted 10 January 2024; published 15 February 2024)

Studies of dynamics on temporal networks often represent the network as a series of “snapshots,” static networks active for short durations of time. We argue that successive snapshots can be aggregated if doing so has little effect on the overlying dynamics. We propose a method to compress network chronologies by progressively combining pairs of snapshots whose matrix commutators have the smallest dynamical effect. We apply this method to epidemic modeling on real contact tracing data and find that it allows for significant compression while remaining faithful to the epidemic dynamics.

DOI: 10.1103/PhysRevLett.132.077402

Introduction.—High resolution temporal interaction data is now simple to obtain and widely available thanks to methods such as radio frequency identification [1] or Bluetooth signals [2]. Temporal interactions have rich dynamics in continuous time, yet we often want to combine intervals of temporal data into a series of simpler, static networks in order to compress the data, reduce analytical complexity, or streamline data collection efforts. For example, digital contact tracing protocols ping devices at fixed intervals to save energy and lighten data requirements. However, it is nontrivial to determine when and how to aggregate temporal data without losing critical information about the dynamics of the interactions.

Many methods currently exist to represent and analyze temporal networks [3], and to find patterns in network structure and dynamics. This includes algorithms for detecting temporal states [4], dynamical approaches for generating synthetic temporal network data [5,6], tools to identify community structure in time-varying networks [7], data-driven approaches to model dynamics on temporal networks by determining change points [8], and methods to represent key temporal features as static networks [9,10]. The dynamics of epidemic spread on temporal networks is well studied [11–14], as are synchronization [15–17] and control dynamics [18].

A continuing challenge is the interplay of dynamics on the network (changes in variables on nodes and edges) with dynamics of the network (where the topological structure changes over time). When the timescales of these two are well separated, we can take one of two limits. If the dynamics on the network are much faster, we can use a static limit where the network structure is essentially constant. When the dynamics on the network are much slower, referred to as the annealed limit, then the dynamical variables effectively experience the average of the network

structure over time. Then we can aggregate many snapshots regardless of their chronological order. In between these two limits, neither type of dynamics can be neglected [19]. Then it is less clear how or if the history of the network's structure can be compressed while remaining faithful to the dynamics on the network.

Here, we quantify the importance of chronology in a sequence of network snapshots by considering its effects on dynamics. Our goal is to compress unimportant structural changes while preserving changes that significantly affect the dynamical process. By compression, we mean a reduction in the number of snapshots used to represent a network's history, rather than a minimization of the description length or memory needed for these snapshots [5]. This task has practical benefits, such as reducing the computational cost of simulation, and we consider it a

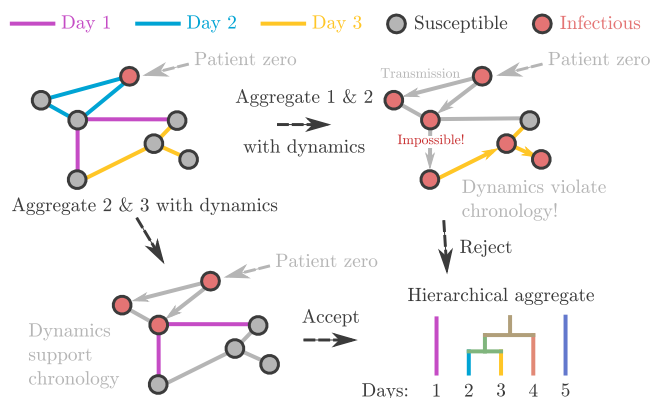


FIG. 1. Schema of our hierarchical aggregation. Given network snapshots, we compare the aggregate spreading dynamics of each adjacent pair of snapshots and combine the pair with the lowest induced error, continuing until we reach a desired number of snapshots.

fundamental scientific question: to what extent do we need to keep track of the chronological changes to a network's structure in order to model dynamical processes on it?

We propose a method to do so by assessing the sensitivity of the dynamics to aggregating pairs of network snapshots. We take an epidemic spreading model as our main example, but we abstract the dynamics in a way that could represent other dynamical processes on networks (e.g., synchronization of coupled oscillators or cascading failures in power grids). Specifically, we formulate a pairwise error measure using matrix commutators that captures how aggregating snapshots affects the epidemic process, and we aggregate network snapshots as long as this error remains low. Using synthetic networks and real data, we find that this approach is successful at producing a compressed snapshot sequence that still mimics the dynamic behavior of the original sequence.

Analytical framework.—We assume that we have some temporal network data consisting of a large number of snapshots of the network structure. (If we have a series of brief contacts between pairs of nodes in continuous time, this data would consist of a large series of mostly empty graphs.) We consider the dynamical error we would introduce by aggregating a given consecutive pair of snapshots into one. In an epidemic model, this error consists of creating paths that allow the contagion to progress backwards in time, as depicted in Fig. 1. If aggregating two snapshots would create many such chronology-violating paths, we should keep these snapshots separate and respect their chronological order.

In our epidemic example, a contagion spreads along edges connecting infectious nodes to their susceptible neighbors. We measure the effect of pairwise aggregation with a measure of error capturing the difference in the number of infected individuals, with and without aggregation over the duration of the snapshot pair. As we will see, we can approximate this error in terms of the commutator of two matrices, each of which linearizes the dynamics over the time interval of the snapshots.

Linearizing the SI model.—Let $A(t)$ be the time-varying adjacency matrix of a network with N nodes. Let $P(t)$ be a vector of length N where $P_i(t)$ is the probability that a node i is infected at time t . A susceptible-infected-susceptible (SIS) model with infection rate β and recovery rate γ can be modeled by the following system of differential equations:

$$\frac{dP_i}{dt} = \beta[1 - P_i(t)][A(t)P(t)]_i - \gamma P_i, \quad (1)$$

and the total number of infected individuals at time t is $I(t) = |P(t)|$ where $|P| = \sum_{i=1}^N P_i$. In this paper we will focus on the SI model where $\gamma = 0$. We then consider the early time linearization of Eq. (1) around $P = 0$ (without losing much accuracy [20]), and thus have

$$\frac{dP}{dt} = \beta A(t)P(t). \quad (2)$$

The solution to this equation involves a time-ordered exponential of $A(t)$. In the case of a series of snapshots A_1, \dots, A_ℓ , i.e., a sequence of static adjacency matrices which hold over time intervals of duration $\delta t_1, \dots, \delta t_\ell$, we have

$$P(t) = \left[\prod_{i=\ell}^1 \exp(\beta \delta t_i A_i) \right] P(0). \quad (3)$$

[This product is in reverse order since we treat $P(t)$ as a column vector and multiply on the left.] If all the A_i commute, this product reduces to $\exp(\beta \sum_i \delta t_i A_i)$, a single exponential of their weighted average. However, in general their noncommutative nature must be taken into account.

Aggregating snapshots.—Say we have two consecutive snapshots, A and B , which are valid for time intervals $[t_0^A, t_1^A]$ and $[t_0^B, t_1^B]$ of duration $\delta t_A = t_1^A - t_0^A$ and $\delta t_B = t_1^B - t_0^B$ respectively, where $t_0^B = t_1^A$. In the linearized SI dynamics of Eq. (2), aggregating A and B in a single step is equivalent to replacing the time evolution operator $\exp(\beta \delta t_B B) \exp(\beta \delta t_A A)$ with

$$\exp[\beta(\delta t_A + \delta t_B) \overline{(A, B)}],$$

where $\overline{(A, B)}$ is the time-weighted average

$$\overline{(A, B)} = \frac{\delta t_A A + \delta t_B B}{\delta t_A + \delta t_B}. \quad (4)$$

The effect of this aggregation can be quantified through the well-known Baker-Campbell-Hausdorff formula, which expresses the product $\exp X \exp Y$ as a single exponential $\exp Z$:

$$Z = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}([X, [X, Y]] - [Y, [X, Y]]) + \dots \quad (5)$$

where $[X, Y] = XY - YX$ is the matrix commutator. Setting $X = \beta \delta t_B B$ and $Y = \beta \delta t_A A$, we have

$$\exp(\beta \delta t_B B) \exp(\beta \delta t_A A) = \exp Z,$$

where

$$Z = \beta(\delta t_A + \delta t_B) \overline{(A, B)} + \frac{1}{2} \beta^2 \delta t_A \delta t_B [B, A] + O(\beta^3). \quad (6)$$

Up to this point, we have used a known matrix formulation of dynamical processes on networks with or without temporal aggregation. This formulation was previously used to study the impact of temporal networks on the

epidemic threshold [21,22] and to measure autocorrelations in temporal networks [22].

Here, the correction term proportional to $[B, A]$ captures to leading order how aggregation introduces error into the dynamics. As illustrated in Fig. 1, aggregation creates chronology-violating paths whenever A and B do not commute. For example, if A includes the edge (1,2) and B includes the edge (2,3), then a disease could spread from node 1 to 3 in the chronology BA but not in AB . The commutator $[B, A]$ counts paths allowed by BA and subtracts paths allowed by AB , compensating for the fact that the Taylor series of $\exp[\overline{(A, B)}]$ contains terms proportional to both AB and BA .

Thus we use the correction term in Eq. (6) as a measure of the error induced by aggregation, or equivalently a measure of the sensitivity of the dynamics to the temporal ordering of A and B . Specifically, we use the operator norm of this term, i.e., its largest singular value, as an estimate of the error induced from the start to the end of the combined interval $[t_0^A, t_1^B]$:

$$\epsilon_{\text{END}} = \frac{1}{2} \beta^2 \delta t_A \delta t_B \|[B, A]\|_{\text{op}} \quad (7)$$

The operator norm bounds the effect of aggregation on any state $P(t_0^A)$ at the beginning of the interval. Note that ϵ_{END} is small if β is small, i.e., if the epidemic dynamics is slow compared to the dynamics of the network structure. It is also small if either δt_A or δt_B is small, since in that case the effect of A or B is close to the identity.

We are interested in the effect of aggregation on the entire history of the system, not just the final state. We show an explanation of the problem in Fig. 2. Thus we also consider the error induced at the boundary $t_1^A = t_0^B$ between the two intervals. To leading order we have

$$\exp(\beta \delta t_A A) - \exp(\beta \delta t_A \overline{(A, B)}) = \beta \delta t_A (A - \overline{(A, B)}) + O(\beta^2).$$

This term is in fact symmetric in A and B , since

$$\beta \delta t_A (A - \overline{(A, B)}) = \beta \frac{\delta t_A \delta t_B}{\delta t_A + \delta t_B} (A - B).$$

As before we bound the possible error induced at the midpoint as the operator norm of this term,

$$\epsilon_{\text{MID}} = \beta \frac{\delta t_A \delta t_B}{\delta t_A + \delta t_B} \|A - B\|_{\text{op}}. \quad (8)$$

We scale these terms by the cumulative duration of both snapshots—capturing the total effect that ϵ_{MID} and ϵ_{END} have on the epidemic process (as visualized in the top right panel of Fig. 2)—which culminates in our error measure defined as

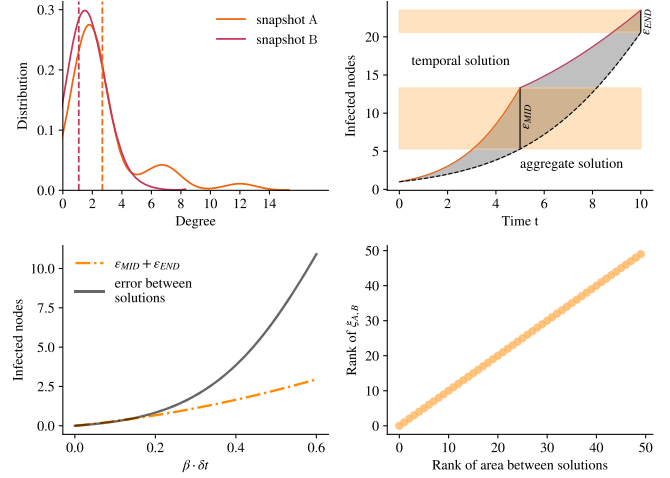


FIG. 2. Top left: degree distributions for two snapshots. Top right: ordinary differential equation solutions of the SI dynamics with $\beta = 0.12$, $\delta t = 5$ on the temporal and aggregate versions of the snapshots with highlighted error terms. Bottom left: ODE solution difference in number of infected nodes under the temporal and aggregate regimes for varying values of βt by varying $t = [0, 5]$. Bottom right: ranking of $\xi_{1,2}$ for snapshots 1 and 2 for increasing values of $\beta \delta t$, compared against the integrated area between solutions.

$$\xi_{A,B} = (\epsilon_{\text{END}} + \epsilon_{\text{MID}})(\delta t_A + \delta t_B). \quad (9)$$

This measure includes both the error incurred by the overall difference between A and B , and the extent to which their chronological order matters. If our main goal is to compute the final state, we can place a smaller weight on ϵ_{MID} or focus entirely on ϵ_{END} . Note that while ϵ_{END} is second order in β and ϵ_{MID} is first order, either can dominate, e.g., if A and B have average degree d then $\|[A, B]\|_{\text{op}}$ can grow as d^2 . This new quantification of compression error captures structure and dynamics, allowing us to introduce a dynamic-preserving algorithm for the compression of temporal networks.

Compression algorithm.—Given a temporal network dataset as a sequence of M snapshots, we can use the framework to compress the snapshots into $M - j$ snapshots via a greedy algorithm. First, the number of desired iterations j is set. For steps from 1 to j , (i) compute the error $\xi_{A,B}$ from Eq. (9) for each ordered pair A, B of consecutive snapshots, (ii) identify the pair $A^*, B^* = \text{argmin}_{A,B}(\xi_{A,B})$ to compress, (iii) replace A and B with their aggregate, $\overline{(A, B)}$ over the union $[t_0^A, t_1^B]$ of their time intervals. Alternately, we can compress until $\min_{A,B} \xi_{A,B}$ reaches some threshold.

This algorithm produces a hierarchical aggregation of snapshots as portrayed in Fig. 1. Periods where the network structure changes slowly or where snapshots commute become single snapshots with a long duration, while the periods with rapidly changing and noncommuting structures are preserved at a higher temporal resolution.

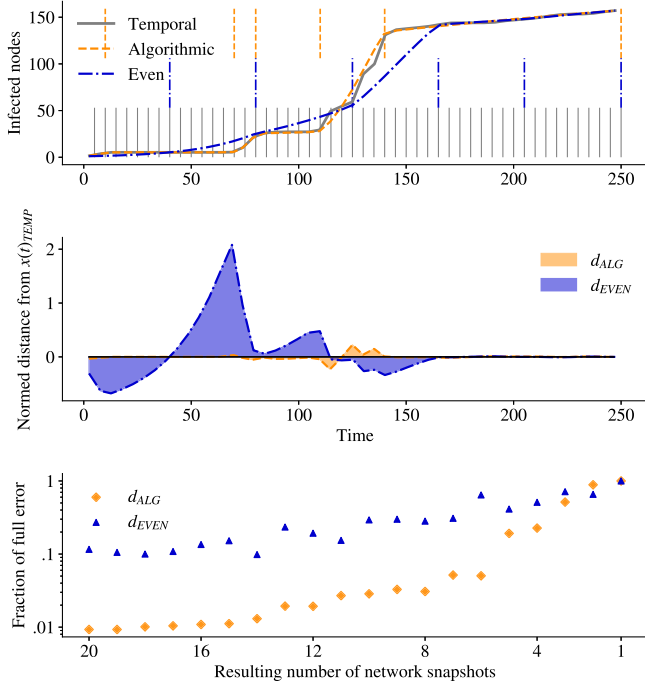


FIG. 3. Compression of a series of 50 synthetic network snapshots (detailed in appendix [20]), showing the SI dynamics with $\beta = 0.0017$. Top: we use our algorithm to compress the network history to 6 snapshots of varying lengths, with boundaries shown by the orange dashed lines. We compare with the SI dynamics using even-width aggregation into 6 windows of fixed length (blue dashed lines). The snapshots produced by our algorithm give SI dynamics closer to that on the full temporal network (gray). Middle: normalized distance from the temporal curve over time for each solution, d_{ALG} for the algorithmic solution, and d_{EVEN} for the even-width solution. The vertical axis shows the shaded area from the middle panel as a function of number of aggregated snapshots, normalized by the error induced by aggregating the entire history into a single static network, measuring the error induced by aggregation as a fraction of the worst-case scenario. Generally, our algorithm results in almost ten times less error than even temporal split.

Results.—Our algorithm produces aggregated snapshots that are able to better support the epidemic dynamics on the temporal network than the standard approach of evenly dividing time into windows of some fixed length and aggregating snapshots in each window. We call $x(t)_{\text{TEMP}}$ the SI dynamics of Eq. (1) integrated using the full temporal dataset and $x(t)_{\text{ALG}}$ the dynamics integrated over the snapshots produced by our algorithm. As comparison methods, we use the common approach of evenly dividing the full time window into a certain number of snapshots of equal length [producing $x(t)_{\text{EVEN}}$] and an information-theoretic method that optimally compresses the full sequence as a set of modes [producing $x(t)_{\text{MDL}}$] [23]. This last method is a parameter-free approach designed to compress general multilayer networks using the minimum description length (MDL) principle. In particular, its level

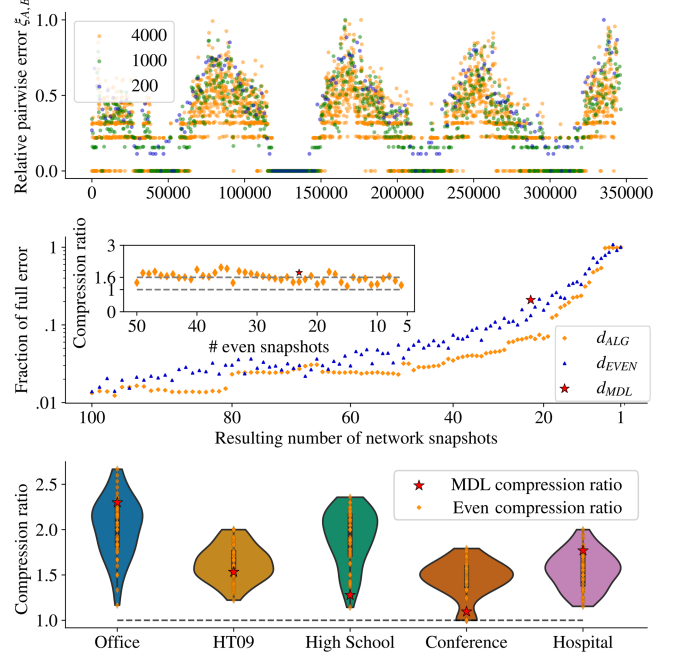


FIG. 4. Application to empirical temporal network data. Top: error measure $\xi_{S(t), S(t+1)}$ computed for consecutive snapshot pairs at three different levels of preaggregation on a hospital contact network [24]. The hospital contact data contains contacts for approximately 9000 unique timestamps. We preaggregate by evenly coarse-graining the data to 4000, 1000, and 200 snapshots. Preaggregation of the data into small static snapshots does not affect the overall quality of our compression [20]. Middle: error of the SI process solutions, relative to the error induced by aggregating the entire history as a single static network (full error), as a function of resulting number of aggregated snapshots. The red star shows the error induced by the MDL optimal compression, which consists of 23 snapshots for this data. The vertical axis of the inset shows the ratio of compression achieved at a given error level by our algorithm versus even-width aggregation in orange markers, and versus MDL in a red star. For example, our algorithm can compress the data to 18 snapshots while maintaining lower error than 30 even snapshots, leading to a 30/18 (~ 1.7) compression ratio. Bottom: summary of the inset of the middle panel across other datasets [26–28] showing the distribution of compression ratios. Our algorithm can be expected to further compress the number of snapshots by 50% to 100%. Importantly, it always outperforms MDL compression, although the two approaches can reach very similar outputs, for example on the conference dataset.

of compression and error cannot be tuned as our algorithm’s or the even division approach can be.

We measure the multiplicative error d_{ALG} as

$$d_{\text{ALG}} = \int \frac{|x(t)_{\text{ALG}} - x(t)_{\text{TEMP}}|}{x(t)_{\text{TEMP}}} dt \quad (10)$$

and similarly for d_{EVEN} and d_{MDL} . We apply the algorithm to synthetic networks in Fig. 3, and to a dataset of four days and nights of contacts in a hospital [24] in Fig. 4, along with several other temporal contact network datasets [25].

We show in the top panel of Fig. 4 how the sensitivity of certain temporal ranges is maintained over a large range of resolution, which allows for preaggregation of data to improve the speed of the algorithm. The error metric also allows us to identify the daily patterns of the contact data at a glance. Once integrated in the compression algorithm, the middle panel shows how we can aggregate over nights at the hospital, and capture the daily activity in one or two snapshots. As seen in Figs. 3 (bottom) and 4 (middle), our algorithm compresses more effectively than even-width time windows and the MDL method. In fact, for a given number of even-width aggregation steps, the algorithm can attain the same level of error after significantly more aggregation steps ($\sim 50\%$ – 200%). We call this additional aggregation the *compression ratio* and use it to summarize the results on other temporal data sets in Fig. 4 (bottom).

Discussion.—The measure $\xi_{A,B}$ of the error induced by aggregating adjacent snapshots, and our hierarchical compression algorithm based on it, offer at least four interesting applications.

First, the approach can directly provide bounds of accuracy when studying dynamics on temporal networks with tools developed for spreading on static networks. Second, our algorithm can help compress large sequences of temporal networks by aggregating consecutive pairs of networks to reach a desired level of simplification. Our algorithm could also be modified to aggregate snapshots until the expected error caused by this aggregation is smaller than some threshold. As shown in Fig. 4 using real temporal interaction data, this approach allowed us to consistently meet a certain level of error while decreasing the number of required network snapshots almost by half compared to aggregating into even-width windows.

Third, the error can be used to estimate the accuracy of data collection in the first place by testing how compressible the data might be. This could help focus data collection efforts by identifying places and times with fast temporal variations, as in the top panel of Fig. 4. Fourth, the error can be used on nontemporal data to compare the structure of any two networks that share some of the same nodes and support the same dynamics. At its core, our approach is a network comparison tool: How different are dynamics on two networks compared to dynamics on their average?

One important limitation is the greediness of our algorithm, as it can get stuck in suboptimal compression sequences. Future work should also explore how to predict the optimal stopping point of temporal compression. We hope that our work will inspire more tools to compress temporal network data while preserving the dynamical processes they support, which is an area rich in possible applications.

Code to support the preparation of this letter is available at [29]. The compression algorithm was implemented in Python and is available at [30].

A. J. A. and L. H. D. acknowledge support from the National Institutes of Health 1P20 GM125498-01 Centers of Biomedical Research Excellence Award and the National Science Foundation Grant No. DMS-1829826. C. M. is supported by NSF Grant BIGDATA-1838251.

*Corresponding author: laurent.hebert-dufresne@uvm.edu

- [1] Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, Vittoria Colizza, Jean-François Pinton, and Alessandro Vespignani, Dynamics of person-to-person interactions from distributed RFID sensor networks, *PLoS One* **5**, e11596 (2010).
- [2] Piotr Sapiezynski, Arkadiusz Stopczynski, David Dreyer Lassen, and Sune Lehmann, Interaction data from the copenhagen networks study, *Sci. Data* **6**, 1 (2019).
- [3] Petter Holme and Jari Saramäki, Temporal networks, *Phys. Rep.* **519**, 97 (2012).
- [4] Naoki Masuda and Petter Holme, Detecting sequences of system states in temporal networks, *Sci. Rep.* **9**, 795 (2019).
- [5] Tiago P Peixoto and Martin Rosvall, Modelling sequences and temporal networks with dynamic community structures, *Nat. Commun.* **8**, 582 (2017).
- [6] Xiao Zhang, Christopher Moore, and Mark E. J. Newman, Random graph models for dynamic networks, *Eur. Phys. J. B* **90**, 200 (2017).
- [7] Amir Ghasemian, Pan Zhang, Aaron Clauset, Christopher Moore, and Leto Peel, Detectability thresholds and optimal algorithms for community structure in dynamic networks, *Phys. Rev. X* **6**, 031005 (2016).
- [8] Tiago P. Peixoto and Laetitia Gauvin, Change points, memory and epidemic spreading in temporal networks, *Sci. Rep.* **8**, 15511 (2018).
- [9] Petter Holme, Epidemiologically optimal static networks from temporal network data, *PLoS Comput. Biol.* **9**, e1003142 (2013).
- [10] Ingo Scholtes, Nicolas Wider, and Antonios Garas, Higher-order aggregate networks in the analysis of temporal networks: Path structures and centralities, *Eur. Phys. J. B* **89**, 61 (2016).
- [11] Luis E. C. Rocha, Fredrik Liljeros, and Petter Holme, Simulated epidemics in an empirical spatiotemporal network of 50,185 sexual contacts, *PLoS Comput. Biol.* **7**, e1001109 (2011).
- [12] Mathieu Géniois, Christian L. Vestergaard, Ciro Cattuto, and Alain Barrat, Compensating for population sampling in simulations of epidemic spread on temporal contact networks, *Nat. Commun.* **6**, 8860 (2015).
- [13] Guangming Ren and Xingyuan Wang, Epidemic spreading in time-varying community networks, *Chaos* **24**, 023116 (2014).
- [14] Eugenio Valdano, Luca Ferreri, Chiara Poletto, and Vittoria Colizza, Analytical computation of the epidemic threshold on temporal networks, *Phys. Rev. X* **5**, 021005 (2015).
- [15] Stefano Boccaletti, D.-U. Hwang, Mario Chavez, Andreas Amann, Jürgen Kurths, and Louis M. Pecora, Synchronization in dynamical networks: Evolution along commutative graphs, *Phys. Rev. E* **74**, 016102 (2006).

- [16] Daniel J. Stilwell, Erik M. Bollt, and D. Gray Roberson, Sufficient conditions for fast switching synchronization in time-varying network topologies, *SIAM J. Appl. Dyn. Syst.* **5**, 140 (2006).
- [17] Yuanzhao Zhang and Steven H. Strogatz, Designing temporal networks that synchronize under resource constraints, *Nat. Commun.* **12**, 3273 (2021).
- [18] Aming Li, Sean P. Cornelius, Y.-Y. Liu, Long Wang, and A.-L. Barabási, The fundamental advantages of temporal networks, *Science* **358**, 1042 (2017).
- [19] Guillaume St-Onge, Jean-Gabriel Young, Edward Laurence, Charles Murphy, and Louis J. Dubé, Phase transition of the susceptible-infected-susceptible dynamics on time-varying configuration model networks, *Phys. Rev. E* **97**, 022305 (2018).
- [20] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.132.077402> for further detail on our derivation of the error measure, on our algorithmic implementation, and on our experiments.
- [21] Leo Speidel, Konstantin Klemm, Víctor M Eguíluz, and Naoki Masuda, Temporal interactions facilitate endemicity in the susceptible-infected-susceptible epidemic model, *New J. Phys.* **18**, 073013 (2016).
- [22] Eugenio Valdano, Michele Re Fiorentin, Chiara Poletto, and Vittoria Colizza, Epidemic threshold in continuous-time evolving networks, *Phys. Rev. Lett.* **120**, 068302 (2018).
- [23] Alec Kirkley, Alexis Rojas, Martin Rosvall, and Jean-Gabriel Young, Compressing network populations with modal networks reveal structural diversity, *Commun. Phys.* **6**, 148 (2023).
- [24] Philippe Vanhems, Alain Barrat, Ciro Cattuto, Jean-François Pinton, Nagham Khanafer, Corinne Régis, Byeul-a Kim, Brigitte Comte, and Nicolas Voirin, Estimating potential infection transmission routes in hospital wards using wearable proximity sensors, *PLoS One* **8**, e73970 (2013).
- [25] Ciro Cattuto *et al.*, SocioPatterns is an interdisciplinary research collaboration formed in 2008 that adopts a data-driven methodology to study social dynamics and human activity, <http://www.sociopatterns.org>.
- [26] Mathieu Génois and Alain Barrat, Can co-location be used as a proxy for face-to-face contacts?, *EPJ Data Sci.* **7**, 11 (2018).
- [27] Julie Fournet and Alain Barrat, Contact patterns among high school students, *PLoS One* **9**, e107878 (2014).
- [28] Lorenzo Isella, Juliette Stehlé, Alain Barrat, Ciro Cattuto, Jean-François Pinton, and Wouter Van den Broeck, What's in a crowd? Analysis of face-to-face behavioral networks, *J. Theor. Biol.* **271**, 166 (2011).
- [29] <https://github.com/andrea-allen/network-compression>.
- [30] A. Allen, Compressnets (2024), 10.5281/zenodo.10564812.