Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Beyond Variance: Selecting Low-Rank Temporal Networks That Preserve Causality and Spreading

Madjid Eshaghi Gordji
Semnan University
Mohamadali Berahman
mohamadali_berahman@semnan.ac.ir
Semnan University

Research Article

Keywords: temporal networks, low-rank approximation, proper orthogonal decomposition, time-respecting paths, causal reachability, epidemic spreading, reproducible network science

Posted Date: September 8th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10842709/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Beyond Variance: Selecting Low-Rank Temporal Networks That Preserve Causality and Spreading

Madjid Eshaghi Gordji[1] Mohamadali Berahman[1,*]

[1]Faculty of Mathematics, Statistics and Computer Science, Semnan University, Semnan, Iran
[*]Correspondence: mohamadali_berahman@semnan.ac.ir

Abstract. A temporal network is a sequence of networks: one adjacency matrix for each time window. A recent proper-orthogonal-decomposition approach showed that this sequence can be represented by a small number of low-rank modes. The unresolved practical question is how many modes are enough when the compressed network will be used to study paths or spreading, rather than only to reproduce matrix entries. We introduce a process-aware rule. For each rank, we reconstruct a binary network while preserving the observed number of links in every snapshot. We then test active-snapshot similarity, time-respecting reachability, and deterministic susceptible–infected (SI) outbreak size. The dynamics-preserving rank is the smallest integer rank that keeps reachability Jaccard at least $1 - \varepsilon$ and SI size error at most ε in the worst of up to 24 prespecified active windows. A stochastic susceptible–infected–recovered (SIR) process is used as an external check, not as a selection target. A constructed rare-bridge network demonstrates why variance can fail: three modes retain 99.928% of the variance and 99.950% mean edge Jaccard, yet worst-window causal reachability falls to 0.494; one additional mode restores it exactly. Across five public SocioPatterns data sets, the conventional 95%-variance rank retained 36.1–58.0% of the full numerical rank, whereas the five-percent dynamics rank retained 45.9–78.0%. In every 30-minute empirical network, the process-safe rank was larger by 7–39 modes. Across the 15-, 30-, and 60-minute analyses, the same direction holds in 14 of 15 data-set/bin-width combinations, but not universally. The main lesson is therefore precise: matrix-optimal compression is not automatically process-safe, and rank should be selected for the downstream question.

Keywords: temporal networks; low-rank approximation; proper orthogonal decomposition; time-respecting paths; causal reachability; epidemic spreading; reproducible network science

1 Introduction

A static network is a single map of who is connected to whom. A temporal network is a movie: contacts appear, disappear, and occur in a particular order [1, 2]. That order is not decorative. If Alice meets Bob before Bob meets Carol, information can travel from Alice to Carol. Reversing the two meetings blocks that route. The same static aggregate contains both contacts, but it cannot tell these two stories apart.

Recent work makes the importance of timing unusually clear. Temporality can change how degree heterogeneity affects the fixation of cooperation [3]; heterogeneous waiting times reshape diffusion [4]; and non-Markovian timing and agent heterogeneity can be combined in scalable stochastic-network models [5]. Ordered events also expose multiscale patterns in tanker shipping [6]. In a different but related meaning of "temporal," the ordered sequence of bond additions reveals self-similar gap statistics and percolation universality classes [7]. Together, these studies say that a temporal representation must be judged by what its ordering allows a process to do, not only by how similar its matrices look.

This need creates a practical storage and analysis problem. Let $A(t) \in \{0,1\}^{n\times n}$ be the adjacency matrix at snapshot $t = 1, \ldots, T$. Keeping all T matrices is expensive, and the resulting object is hard to inspect. Lacasa recently treated the centered sequence as a trajectory in graph space and applied proper orthogonal decomposition (POD) and dynamic mode decomposition [8]. POD orders orthogonal network modes by variance, gives the best rank-r matrix approximation in Frobenius norm, and can preserve intrinsic temporal fingerprints such as periodicity or autocorrelation in low-dimensional projections. This is a clean mathematical foundation for compressing the adjacency sequence.

However, "best matrix approximation" and "best network for a process" are different statements. A link that appears once has little squared energy. If that link is the only correctly timed bridge between two communities, removing it can destroy many time-respecting paths. No contradiction is involved: Frobenius error counts changed entries, whereas reachability counts changed consequences.

Existing compression methods address neighboring questions. Graph commutators guide the merging of adjacent time intervals and bound errors for spreading operators, so the reduced axis is time [9]. Network hashing reduces the number of nodes and estimates temporal out-components [10]. Embedding-based distance [11] and temporal-network dissimilarity [12] compare whole temporal graphs rather than selecting the truncation rank of an adjacency trajectory. Prasse and Van Mieghem reduce observed node-state dynamics and predict them without requiring the interaction graph, which is a different object from the time-varying graph itself [13]. These works are complementary, but they do not answer the following question:

> *How far should a low-rank decomposition of the adjacency sequence be truncated if the reconstructed network must preserve both its matrix structure and the processes that travel through it?*

We answer this question with a deliberately simple procedure. We retain r POD modes, reconstruct a score for every possible link and time, and keep exactly the observed number of links in every snapshot. We then measure three things: active-snapshot edge overlap, time-respecting reachability, and deterministic SI outbreak

size. We define a dynamics rank, $r_{\mathrm{dyn}}(\varepsilon)$, as the first integer rank that passes reachability and SI tolerances in the worst of up to 24 evenly sampled active windows. Reachability and deterministic SI are two summaries of the same probability-one propagation process, rather than independent outcomes: the first tests the set of reachable ordered pairs, whereas the second tests source-wise final sizes. A stochastic SIR model provides a separate robustness check [14, 15].

Here "causal" has a narrow network meaning: a path must respect time order. It does *not* mean statistical causal inference or an intervention effect. This distinction matters because our method preserves possible routes, not claims about why a contact occurred.

We evaluate three explicit hypotheses. **H1**: a low-energy link can have a large causal effect, so a variance threshold can accept an unsafe rank. **H2**: on empirical contact networks at the prespecified 30-minute resolution, $r_{\mathrm{dyn}}(0.05)$ exceeds the conventional 95%-variance rank r_{95}. **H3**: the selected dynamics rank also keeps paired SIR outcome error below five percentage points, although SIR is not used to choose the rank. The code tests every integer rank until the five-percent condition first holds; every number, table, and figure reported below is generated by that code.

2 Results

A rare link can carry almost no variance and almost all cross-community reachability. Figure 1 shows the constructed test. Forty nodes form two groups of 20. Four recurring within-group patterns create most of the variance. At snapshot 31, one within-group link is replaced by a single between-group bridge. The network therefore has the same number of links at every time. Its first three modes retain 99.9281% of the centered variance. The decoded rank-three sequence has mean edge Jaccard 0.999496, so it appears almost exact when snapshots are averaged. Yet it omits the bridge. In the most affected eight-snapshot window, reachability Jaccard is only 0.4935, and the mean absolute SI outbreak-size error across source nodes is 0.4875. The fourth mode contributes only 0.0719% of variance, but restores the bridge and both process metrics exactly. Thus H1 is satisfied by construction and measurement, not by a qualitative drawing.

The reachability matrices make the mechanism visible (Fig. 2). In the original sequence, 95% of cross-group ordered pairs are reachable in the selected window. At r_{95} the two diagonal blocks remain internally reachable, but cross-group reachability becomes zero. At r_{dyn} the original pattern returns. A low-rank method can therefore preserve communities and density while deleting the one event that connects their futures.

Five public contact networks give the same warning at 30 minutes. We analyzed a hospital ward, the Hypertext 2009 conference, a workplace, a primary school, and a high school. All were downloaded from the official SocioPatterns pages and verified by SHA-256 checksum [16]. The hospital and conference networks were used in full, with 75 and 113 nodes. The workplace also fit below the cap. For the two larger schools, the code selected the 120 nodes with the highest number of incident 20-second records before building snapshots. Table 1 reports the exact analyzed sizes.

In all five 30-minute networks, $r_{\mathrm{dyn}}(0.05) > r_{95}$ (Table 2 and Fig. 3). The extra requirement ranged from 7 modes in the primary school to 39 in the hospital ward. The conventional ranks retained 36.1–58.0% of the full numerical rank; process-safe ranks retained 45.9–78.0%. Relative to r_{95}, this was a 24.1–46.2% increase in the number of modes. These results support H2 for the prespecified resolution; they are descriptive comparisons across five data sets, not a population-level significance test.

The gap cannot be explained by uniformly poor reconstruction at r_{95}. Mean edge Jaccard over active snapshots was 0.829–0.983 across the empirical networks. Nevertheless, worst-sampled-window reachability Jaccard ranged from 0 to 0.625, with a median of 0.500. The sharpest example was Hypertext 2009: its rank-52 reconstruction had active-snapshot edge Jaccard 0.829, but the original and reconstructed off-diagonal reachability sets did not overlap in the most affected sampled window. At the process-safe rank 76, edge and process metrics were exact for the monitored windows.

The rank curves explain why a single variance cut cannot substitute for direct testing (Fig. 4). Variance increases smoothly because each singular value adds nonnegative energy. Reachability does not. It can remain flat, jump when a bridge returns, or briefly fall because the fixed-density decoder exchanges one selected link for another. SI size error is smoother because it averages source-wise sizes, but it also has small steps. This non-monotonicity is why the code examines consecutive integer ranks until the first passing rank rather than using binary search.

A stochastic process confirms the selected rank, but also shows why averages are not enough. For each data set we ran 12 active-window SIR scenarios with 80 repetitions per scenario, using common random numbers for the original and reconstructed networks. At r_{95}, the paired run-level error in final infected fraction ranged from 0.000147 to 0.006683 in the empirical networks. These averages are small even where reachability fails, because a rare changed path need not be used by a particular stochastic realization. At $r_{\mathrm{dyn}}(0.05)$, the largest paired final-size error was 0.000125 and the largest normalized peak-time error was 0.000586; all were far below the prespecified five-percentage-point bound (Fig. 5), supporting H3. This is an external check of the selected representations, not evidence that the chosen SIR parameters describe a particular disease.

Figure 6 summarizes the practical message. High edge overlap is useful, but points with nearly identical edge Jaccard can have very different reachability. The upper green band is reached only after the causal test passes. The selected diamonds therefore answer a process question; they are not claimed to be universally optimal ranks.

Table 1: Empirical data after deterministic preprocessing. A snapshot is active if it contains at least one retained link. The table is generated directly from the downloaded records.

Data set	Original n	Analysed n	20-s records	Snapshots	Active	Unique links	Days
Hospital ward	75	75	32424	194	162	1139	4.02
Hypertext 2009	113	113	20818	118	81	2196	2.45
Workplace	92	92	9827	549	207	755	11.42
Primary school	242	120	125773	65	36	3369	1.35
High school 2013	327	120	188508	202	82	1227	4.21

Table 2: Variance and dynamics ranks. Reachability is the worst-sampled-window Jaccard; SI error is the worst-sampled-window mean absolute error in source-wise final size. The five-percent rank requires reachability ≥ 0.95 and SI error ≤ 0.05.

Data set	Full k	r95	r-dyn (10%)	r-dyn (5%)	Reach. at r95	SI error at r95	Active edge Jaccard at r-dyn
Rare-bridge model	33	3	4	4	0.494	0.487	1.000
Hospital ward	173	96	135	135	0.091	0.008	0.993
Hypertext 2009	103	52	76	76	0.000	0.053	1.000
Workplace	388	140	178	178	0.500	0.004	1.000
Primary school	50	29	36	36	0.500	0.011	1.000
High school 2013	157	57	78	78	0.625	0.018	1.000

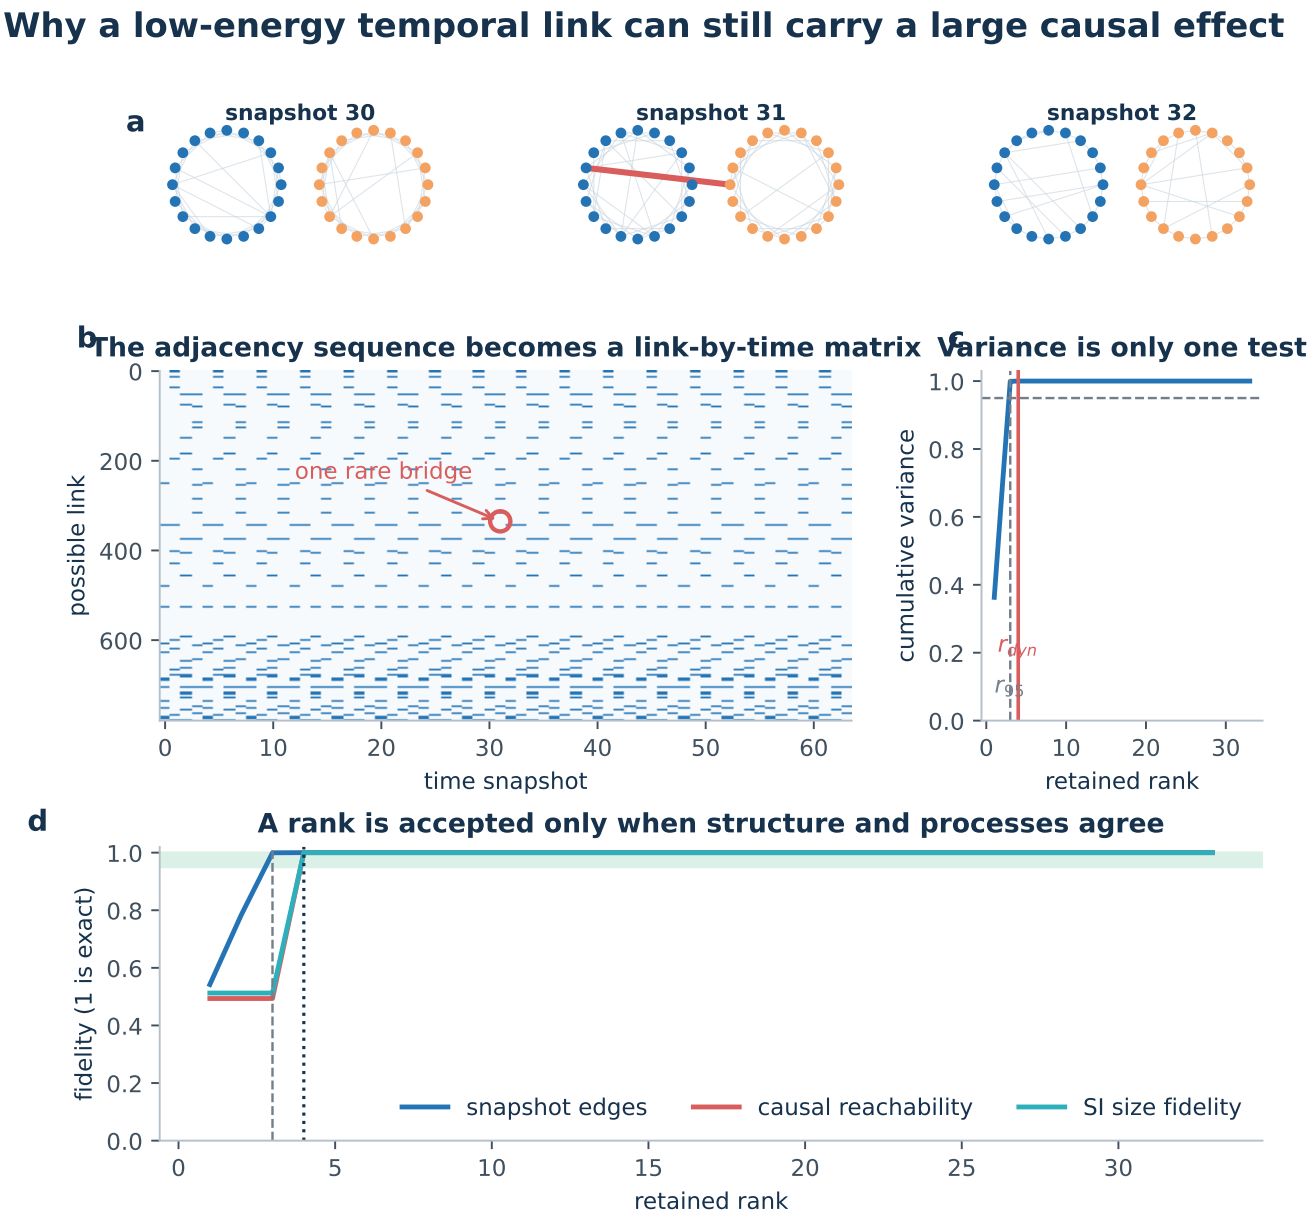


Figure 1: Why variance alone can miss a process-critical event. (a) Three snapshots around the rare bridge; the red link is the only cross-group contact. (b) The same sequence as a link-by-time matrix. (c) Three modes already exceed the conventional 95% variance threshold; the bridge is carried by the next, very low-energy mode. (d) Edge overlap is almost exact at $r_{95} = 3$, but worst-window reachability and SI fidelity remain near one half until $r_{\mathrm{dyn}}(0.05) = 4$. All panels are produced by the supplied Python code.

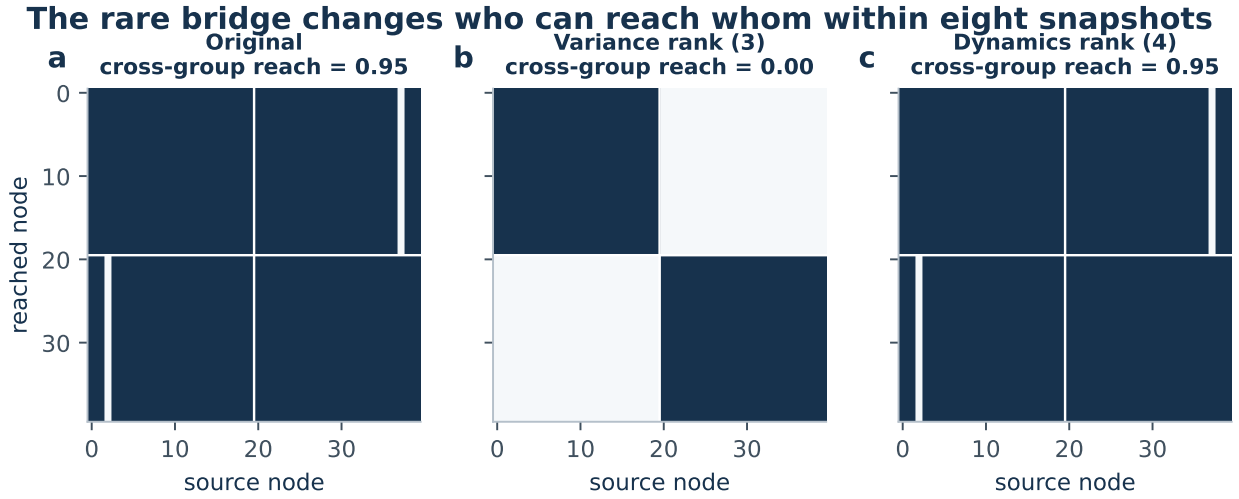


Figure 2: Time-respecting reachability in the rare-bridge window. Dark entries mean that the source on the horizontal axis can reach the node on the vertical axis in at most eight snapshots. White cross-blocks at the variance rank show the causal failure. The dynamics rank restores the original matrix.

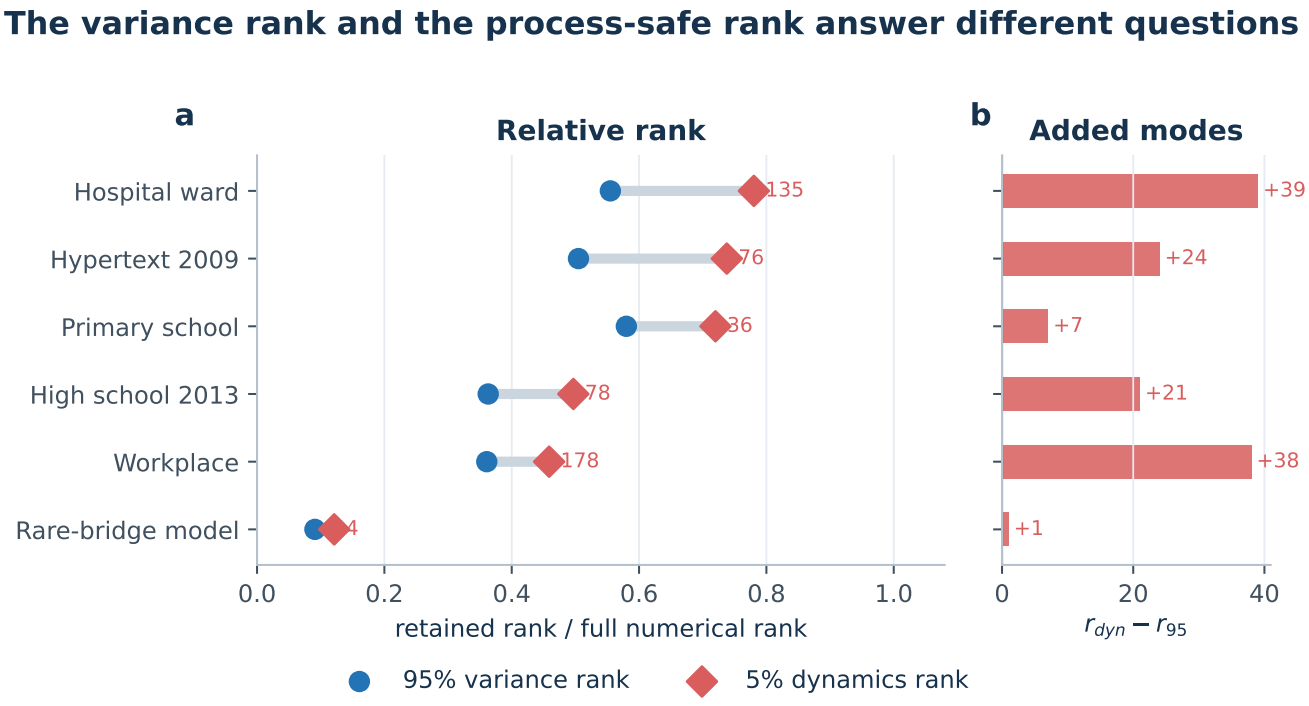


Figure 3: Process safety generally requires additional modes. (a) Blue circles are r_{95}; red diamonds are $r_{\mathrm{dyn}}(0.05)$. Ranks are divided by the full numerical rank so networks of different duration can be compared. Labels beside diamonds give the unnormalized dynamics rank. (b) The exact number of additional modes.

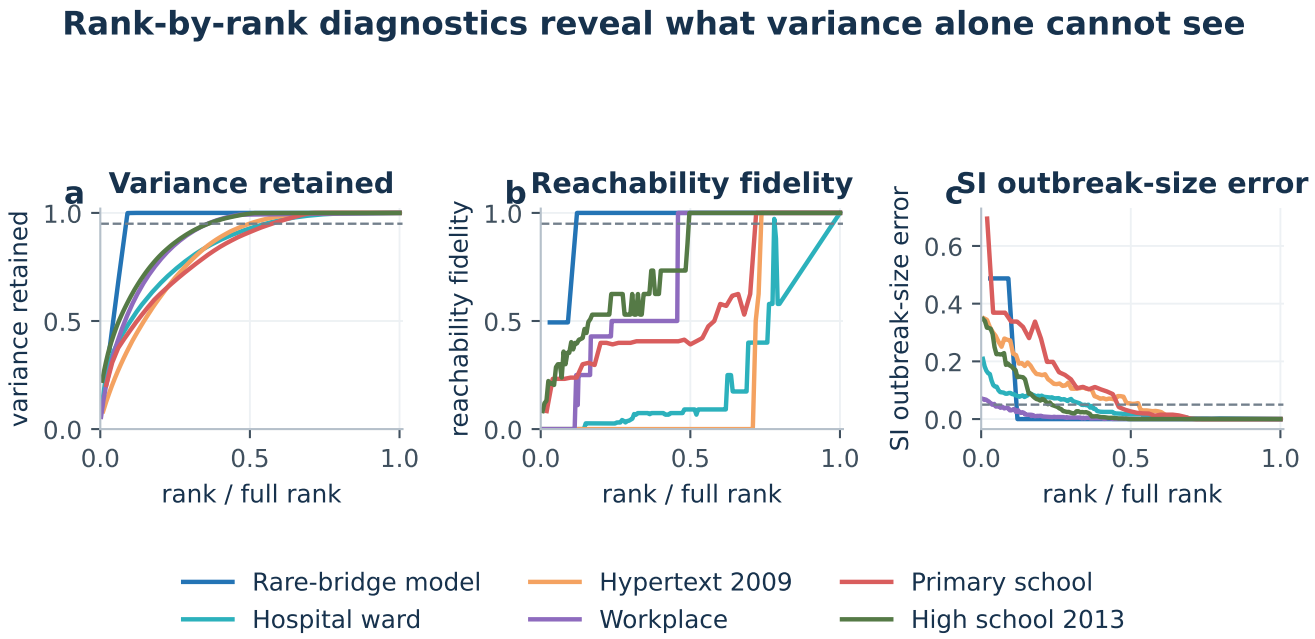


Figure 4: Rank-by-rank diagnostics. (a) Variance retained is smooth and monotone. (b) Worst-sampled-window reachability is discrete and can be non-monotone after binary fixed-density decoding. (c) Worst-sampled-window deterministic-SI size error. Dashed lines are the five-percent acceptance boundaries. Curves include all consecutive ranks tested before acceptance plus reference anchors.

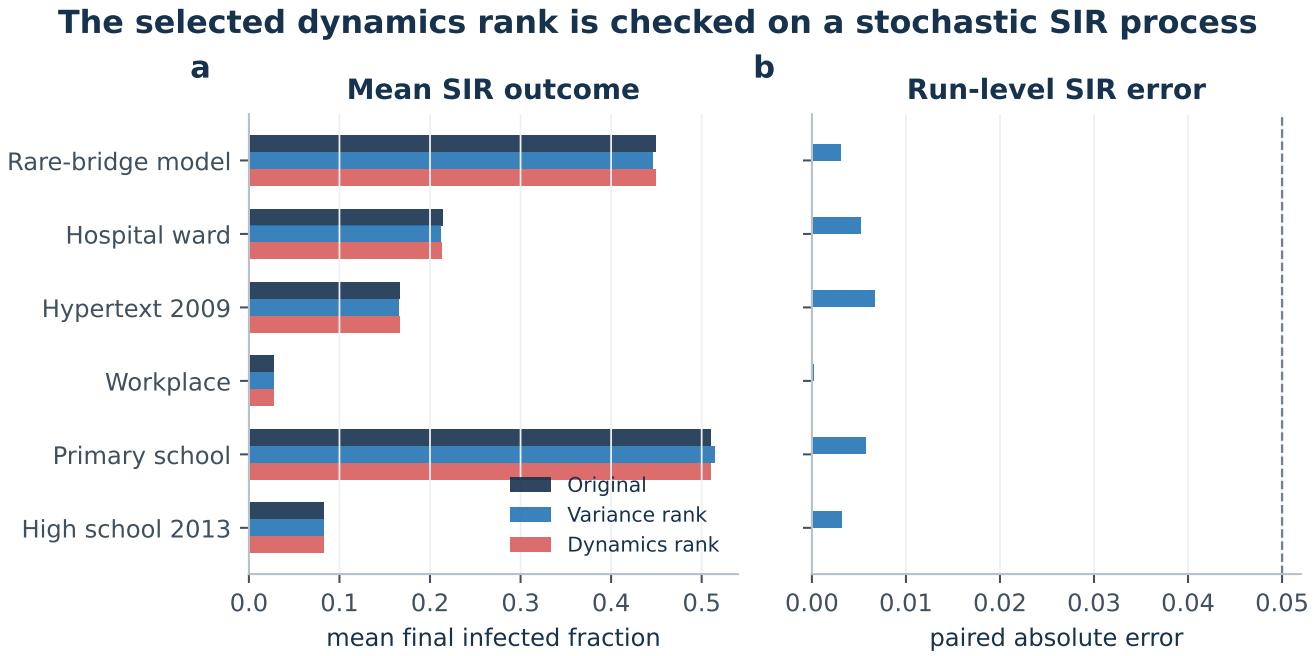


Figure 5: Stochastic SIR validation. (a) Mean final infected fraction for the original, r_{95}, and $r_{\mathrm{dyn}}(0.05)$ sequences. (b) Paired run-level absolute error. Dynamics-rank errors are zero or visually indistinguishable from zero at this scale. The dashed line marks five percentage points.

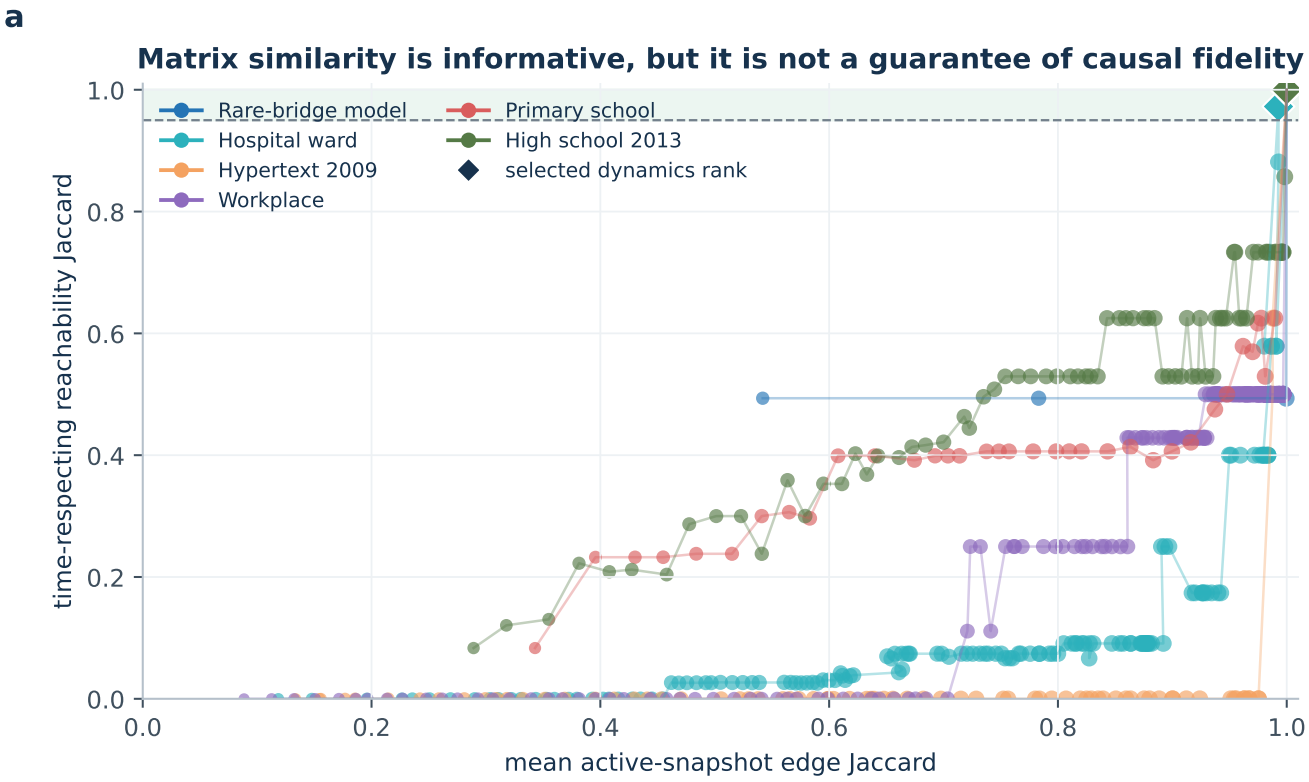


Figure 6: Structural similarity does not guarantee causal fidelity. Each point is a tested rank; point size grows with retained variance. The shaded band denotes at least 0.95 worst-sampled-window reachability Jaccard. Diamonds mark the selected dynamics rank.

3 Methods

Data acquisition and snapshot construction. The five event lists were downloaded from SocioPatterns, a public infrastructure for high-resolution proximity data [16, 17]. The hospital data record 20-second proximity events between patients and health-care workers [18]; Hypertext 2009 records conference contacts [19]; the workplace data were collected in an office building [20]; and the school data describe face-to-face contacts measured by wearable sensors [21, 22]. The exact official URLs, licenses, filenames, and checksums are stored in the code and result manifest. The code stops if a checksum differs.

For the main analysis, time was divided into consecutive 1,800-second bins anchored at the first retained event and continued through the last retained event, including empty overnight bins. A link (i, j) was set to one in bin t if at least one 20-second record for that pair occurred in the bin. Repeated records inside the same bin were collapsed. Self-loops were removed. If a data set had more than 120 nodes, nodes were ranked by their number of incident raw records and the induced temporal subgraph of the first 120 was used. This deterministic cap affects only the two school networks and is reported rather than hidden.

Link-by-time matrix and POD. For an undirected network without self-loops, let $p = n(n-1)/2$ and let $x(t) \in \{0,1\}^p$ be the upper triangle of $A(t)$. Define the temporal mean

$$\bar{x} = \frac{1}{T}\sum_{t=1}^{T} x(t) \tag{1}$$

and stack the centered snapshots as

$$S = [x(1) - \bar{x}, \ldots, x(T) - \bar{x}] \in \mathbb{R}^{p\times T}. \tag{2}$$

The singular value decomposition is

$$S = U\Sigma V^{\mathsf{T}} = \sum_{k=1}^{K} \sigma_k u_k v_k^{\mathsf{T}}, \qquad \sigma_1 \geq \cdots \geq \sigma_K > 0, \tag{3}$$

where $K = \operatorname{rank}(S)$. This is the snapshot form of POD [8, 23]. We computed the exact temporal covariance $S^{\mathsf{T}}S$, diagonalized it, and recovered U by $u_k = Sv_k/\sigma_k$. This avoids forming the larger $p \times p$ covariance.

The rank-r score matrix is

$$X_r = \bar{x}\mathbf{1}^{\mathsf{T}} + \sum_{k=1}^{r} \sigma_k u_k v_k^{\mathsf{T}}. \tag{4}$$

The cumulative variance is

$$E_r = \frac{\sum_{k=1}^{r}\sigma_k^2}{\sum_{k=1}^{K}\sigma_k^2}, \qquad r_{95} = \min\{r : E_r \geq 0.95\}. \tag{5}$$

By the Eckart–Young–Mirsky theorem, the centered component $X_r - \bar{x}\mathbf{1}^{\mathsf{T}}$ is matrix-optimal in Frobenius norm among rank-r approximations to S [24, 25]. A self-contained proof is given in Appendix A.

Fixed-density binary decoder. POD scores are real-valued, but paths need binary links. Let $m_t = \sum_e x_e(t)$ be the observed number of links in snapshot t. The decoder selects the m_t largest entries of $X_r(:, t)$. Exact ties at the selection boundary are resolved by ascending link index, without perturbing distinct scores. Thus

$$\sum_e \widehat{x}_{r,e}(t) = m_t \quad \text{for every } t. \tag{6}$$

Any difference in a process is therefore not caused by changing snapshot density. Appendix A proves that this decoder is the closest binary vector to the score vector among all vectors with m_t ones.

Structural and process metrics. Snapshot fidelity is the mean edge Jaccard over the set $\mathcal{T}_+$ of snapshots containing at least one observed link,

$$J_E(r) = \frac{1}{|\mathcal{T}_+|}\sum_{t\in\mathcal{T}_+} \frac{|E_t \cap \widehat{E}_{r,t}|}{|E_t \cup \widehat{E}_{r,t}|}, \tag{7}$$

so that trivially reconstructed empty overnight snapshots do not inflate the structural score.

For process tests, we considered up to 24 start times sampled evenly from windows containing at least one original contact. Each window spans $h = 8$ snapshots, or four hours at the main resolution. Propagation is synchronous: a node reached during snapshot t can transmit from snapshot $t+1$, not earlier. Let R_w and $\widehat{R}_{r,w}$ be the binary source-to-target reachability matrices in window w, including the diagonal. We remove the diagonal when computing Jaccard and use the conservative fidelity

$$J_R(r) = \min_w \frac{|R_w \cap \widehat{R}_{r,w}|}{|R_w \cup \widehat{R}_{r,w}|}. \tag{8}$$

The deterministic SI size from source i is $s_{i,w} = n^{-1}\sum_j R_w(j, i)$. Its error is

$$\Delta_{\mathrm{SI}}(r) = \max_w \frac{1}{n}\sum_{i=1}^{n} |s_{i,w} - \widehat{s}_{r,i,w}|. \tag{9}$$

Time-respecting reachability and deterministic SI with transmission probability one are exactly equivalent; the induction proof is in Appendix A.

For tolerance ε, the dynamics-preserving rank is

$$r_{\mathrm{dyn}}(\varepsilon) = \min\{r : J_R(r) \geq 1 - \varepsilon \text{ and } \Delta_{\mathrm{SI}}(r) \leq \varepsilon\}. \quad (10)$$

We use $\varepsilon = 0.05$ as the main criterion and 0.10 as a sensitivity criterion. The set in Eq. (10) is nonempty because full numerical rank recovers the observed binary sequence; Appendix A proves this statement.

Stochastic SIR check. The SIR analysis was intentionally generic. In each time step, a susceptible node with k infected neighbors became infected with probability

$$1 - (1 - \beta)^k, \quad (11)$$

with $\beta = 0.18$; each infected node recovered with probability $\gamma = 0.08$. We used 16 snapshots, 12 active-window start times, and the most active node in each window as its seed. Each scenario had 80 repetitions, giving 960 runs per network and representation. Identical node-level uniform random numbers were used for each original/compressed pair. We report paired mean absolute errors in final infected fraction and normalized peak time. These parameters test sensitivity to a stochastic process; they are not fitted to influenza, COVID-19, or any other disease.

Computation, sensitivity, and literature boundary. All analyses used Python 3 with NumPy 2.3.5, pandas 2.2.3, SciPy 1.17.0, and Matplotlib 3.10.8. Synthetic and SIR seeds were 17 and 314159. Unit tests verify full-rank recovery, exact preservation of per-snapshot link count, and sensitivity to reversed contact order. Table 4 in the Appendix is generated from the same parameter source as the analysis manifest.

To test dependence on temporal resolution, the entire rank analysis was repeated at 15, 30, and 60 minutes. Fourteen of 15 empirical data-set/bin-width combinations had $r_{\mathrm{dyn}}(0.05) \geq r_{95}$; the exception was the 15-minute workplace network, where the values were 190 and 224. The horizon was held at eight snapshots, so the corresponding clock-time horizons were two, four, and eight hours. This sensitivity analysis therefore tests the combined effect of binning and the process timescale expressed in snapshot units; it does not isolate bin width alone. The main direction is robust in these combinations but is not presented as a theorem about every temporal network.

The novelty statement is bounded by a targeted search completed on 26 August 2026. We searched publisher pages, Crossref metadata, Nature Portfolio, APS, Springer Nature, PNAS/PMC, and combinations of "proper orthogonal decomposition," "low rank," "temporal network," "time-respecting," "reachability," "epidemic," and "dynamics-preserving rank." We found link-space POD with variance and intrinsic-dynamics validation [8], time-axis commutator compression [9], node hashing for out-components [10], temporal graph distances [11, 12], and state-space prediction [13], but no prior study selecting adjacency-sequence POD rank jointly by time-respecting reachability and spreading. We therefore use the qualified phrase "to our knowledge after a targeted search," not an absolute priority claim.

4 Discussion

The central result is simple: the amount of variance retained by a temporal-network approximation does not determine how much causal structure it retains. POD is doing exactly what its objective asks. The mismatch appears because squared matrix error gives every changed link-time entry a local cost, while temporal reachability can amplify one changed entry into many lost source–target relations.

This observation does not make r_{95} useless. It remains a clear summary of matrix energy, and in our experiments it often gave high snapshot overlap and small mean SIR error. It is simply not a certificate for another task. If the downstream question is visualization, periodicity, or mean reconstruction, variance may be sufficient. If the question is whether an alert, infection, resource, or failure can arrive at a node in time, the rank should be checked against that process.

Table 3 places this contribution beside the closest work. Lacasa provides the link-space eigendecomposition that makes the present question possible [8]. Allen and colleagues protect operator dynamics while reducing the number of time intervals [9]. Vaudaine and colleagues reduce nodes and target out-component statistics [10]. Dall'Amico et al. and Zhan et al. provide principled whole-graph comparisons [11, 12]. Our narrower addition is a stopping rule for the number of POD modes when the binary reconstructed adjacency sequence itself will be used for paths and spreading.

The conservative worst-sampled-window rule is appropriate when a rare failure matters. It also raises the selected rank, sometimes substantially. In the hospital, 135 of 173 numerical modes were needed at five-percent tolerance. This is not a failed result: it says that a strong process guarantee permits only modest compression for this data and decoder. A method that always reports dramatic compression would be misleading when the process genuinely depends on small components.

Several limitations define the claim. First, the analysis uses undirected, unweighted snapshots and a fixed-density top-score decoder; weighted or directed networks may require different decoding and path rules. Second, the 120-node cap changes the two school systems, so conclusions apply to the reported induced subnetworks, not all recorded students. Third, reachability is evaluated in up to 24 prespecified sampled active windows, not every possible start and horizon; "worst" therefore means worst among those windows. Fourth, 30-minute aggregation hides ordering inside a bin. Synchronous updates avoid inventing multi-hop paths within one snapshot, but event-level analysis may give different ranks. The 15/60-minute analysis changes both resolution and the clock-time horizon and should be read as a joint sensitivity test. Fifth, deterministic SI and reachability are related views of the same probability-one process, not independent validations. Sixth, the SIR model is a process probe, not epidemiological calibration. Finally,

Table 3: Scope relative to closely related work. The structured comparison is stored as CSV and converted to LaTeX by the supplied code.

Study	Compressed object	Axis reduced	Validation	Main distinction
Prasse & Van Mieghem (2022)	Node states	State dimension	Prediction error	Fixed/unknown graph
Allen et al. (2024)	Snapshots	Time	Commutator + epidemic	Merges adjacent times
Vaudaine et al. (2024)	Nodes/events	Nodes	Out-components	Hashing proxy
Dall'Amico et al. (2024)	Temporal graphs	Embedding	Graph distance	Comparison, not rank selection
Zhan et al. (2025)	Temporal graphs	Representation	Dissimilarity	Comparison, not compression
Lacasa (2026)	Adjacency sequence	Link-space modes	Variance + temporal fingerprint	POD foundation
This work	Adjacency sequence	Link-space modes	Reachability + SI/SIR	Smallest process-safe rank

a targeted literature search cannot prove that no similar unpublished or differently named method exists.

A useful next step is to assign each POD mode both an energy score and a causal salience score. One could then retain low-energy modes that restore many paths before higher-energy modes that only polish dense regions. Such a method would no longer be the Frobenius-optimal truncation and would require a new optimization problem. Other extensions include weighted earliest-arrival error, non-Markovian spreading [5], directed transport networks such as shipping [6], and process-specific ranks for cooperation or diffusion [3, 4]. The present contribution is intended as the transparent baseline against which those more complex methods can be compared.

5 Data availability

All empirical data are public and were obtained from the official SocioPatterns pages [16]. The hospital, Hypertext 2009, primary-school, and high-school event lists are distributed under Creative Commons Attribution–NonCommercial–ShareAlike terms as stated on their source pages; the workplace event list uses a Creative Commons public-domain dedication. The raw files are not duplicated in the code archive. Instead, `run_all.py` downloads the exact official files and verifies the SHA-256 values recorded in `src/temporal_rank/data.py`. This preserves provenance and avoids silently using a modified mirror.

No personal or restricted data were generated. The exact official source pages, direct-download URLs, licences, filenames, and SHA-256 checksums are recorded in the accompanying code and generated data-summary table.

6 Code availability

The complete reproducibility package, including the Python source code, unit tests, environment specification, generated CSV results, LaTeX tables, vector PDF figures, high-resolution PNG figures, and analysis manifest, is openly available in Zenodo at https://doi.org/10.5281/zenodo.22123166. The archived workflow downloads the public data from their official sources, verifies their SHA-256 checksums, and reproduces all quantitative results, tables, and figures reported in this study.

7 Declarations

Author contributions. Both authors contributed to the conceptualization, methodology, interpretation, writing, critical revision, and approval of the manuscript. Mohamadali Berahman is the corresponding author.

Use of artificial intelligence. Generative artificial-intelligence tools were used during manuscript preparation to assist with literature organization, Python-code drafting and checking, figure styling, and language editing. The authors are responsible for the research question, source verification, mathematical arguments, execution and inspection of the analyses, interpretation of every output, and the final text. No AI system is listed as an author.

Competing interests. The authors declare no competing financial or non-financial interests.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

A Mathematical guarantees

This appendix proves the six statements used by the method. They are stated for finite real matrices, which is the setting of every experiment.

Lemma 1 (Trace inequality in the required form)**.** *Let X and Y have singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq 0$ and $\tau_1 \geq \tau_2 \geq \cdots \geq 0$. Then*

$$\langle X, Y \rangle_{\mathrm{F}} = \operatorname{tr}(X^{\mathsf{T}} Y) \leq \sum_i \sigma_i \tau_i. \tag{12}$$

Proof. Write $X = U\Sigma V^{\mathsf{T}}$ and $Y = P\mathcal{T}Q^{\mathsf{T}}$. Expanding the inner product gives terms $\sigma_i \tau_j (u_i^{\mathsf{T}} p_j)(v_i^{\mathsf{T}} q_j)$. Since $ab \leq (a^2 + b^2)/2$, the total is no larger than one half of two sums whose weights are squared overlaps. The matrices $[(u_i^{\mathsf{T}} p_j)^2]$ and $[(v_i^{\mathsf{T}} q_j)^2]$ have nonnegative entries and row and column sums at most one, because both sets of singular vectors are orthonormal. Each is therefore doubly substochastic. A linear functional with decreasing nonnegative coefficients reaches its maximum on a permutation matrix; the rearrangement inequality places largest σ_i with largest τ_i. Each of the two sums is consequently at most $\sum_i \sigma_i \tau_i$, proving Eq. (12). □

Theorem 1 (Optimality of truncated POD/SVD)**.** *Let S have singular values $\sigma_1 \geq \cdots \geq \sigma_K > 0$. Among all matrices Y of rank at most r,*

$$\min_{\operatorname{rank}(Y) \leq r} \|S - Y\|_{\mathrm{F}}^2 = \sum_{k=r+1}^{K} \sigma_k^2, \tag{13}$$

and a minimizer is $S_r = \sum_{k=1}^{r} \sigma_k u_k v_k^{\mathsf{T}}$.

Proof. Let the nonzero singular values of a rank-r candidate Y be $\tau_1, \ldots, \tau_r$. By the lemma,

$$\begin{aligned}
\|S - Y\|_{\mathrm{F}}^2 &= \|S\|_{\mathrm{F}}^2 + \|Y\|_{\mathrm{F}}^2 - 2\langle S, Y \rangle_{\mathrm{F}} && (14)\\
&\geq \sum_{k=1}^{K} \sigma_k^2 + \sum_{k=1}^{r} \tau_k^2 - 2\sum_{k=1}^{r} \sigma_k \tau_k && (15)\\
&= \sum_{k=1}^{r} (\sigma_k - \tau_k)^2 + \sum_{k=r+1}^{K} \sigma_k^2 && (16)\\
&\geq \sum_{k=r+1}^{K} \sigma_k^2. && (17)
\end{aligned}$$

Choosing $Y = S_r$ makes $\tau_k = \sigma_k$ for $k \leq r$, aligns the singular vectors, and makes both inequalities equalities. No rank-r matrix can have smaller Frobenius error. □

Proposition 1 (Optimality of the fixed-density decoder)**.** *For scores $z \in \mathbb{R}^p$ and an integer $0 \leq m \leq p$, the binary vector $b \in \{0,1\}^p$ with exactly m ones that minimizes $\|z - b\|_2^2$ is obtained by placing ones at the m largest scores.*

Proof. Because $\sum_e b_e = m$,

$$\|z - b\|_2^2 = \|z\|_2^2 + m - 2\sum_e z_e b_e. \tag{18}$$

The first two terms are fixed, so minimizing distance is equivalent to maximizing the sum of selected scores. If a candidate selects i but not j with $z_j > z_i$, swapping the two increases the sum by $z_j - z_i > 0$. Repeating this exchange removes every inversion and leaves exactly the m largest entries selected. Ties have equal objective value; the code resolves them deterministically. □

Proposition 2 (Small matrix error does not bound temporal reachability error)**.** *There is a family of temporal networks that differs from another family in one link-time entry, so its relative link error tends to zero as n grows, while a nonzero fraction of ordered node pairs loses reachability.*

Proof. Take two groups C_1 and C_2, each with q nodes. Choose gateways $a \in C_1$ and $b \in C_2$. At time 1, connect every other node in C_1 to a. At time 2, add only the bridge (a, b). At time 3, connect b to every other node in C_2. Every source in C_1 can then reach every target in C_2 through a time-respecting three-step path. Now remove only (a, b). All q^2 ordered pairs from C_1 to C_2 lose reachability. The two sequences differ in one undirected link-time event out of $2(q-1)+1$ events, so their eventwise mismatch fraction is $1/(2q-1) \to 0$, while the lost-pair fraction among all $4q^2$ ordered pairs is $q^2/(4q^2) = 1/4$. Therefore no bound that depends only on a vanishing relative entry error can force reachability error to vanish. □

Proposition 3 (Reachability equals deterministic SI)**.** *Under synchronous updates and transmission probability one, node j is infected by the end of a temporal window from seed i if and only if a time-respecting path from i to j exists in that window.*

Proof. At the start, only i is infected and the only length-zero path ends at i, so the statement holds. Assume it holds after snapshot $t-1$. At snapshot t, a susceptible node v becomes infected exactly when it is adjacent to an already infected node u. By the induction hypothesis, a time-respecting path reaches u by $t-1$; appending edge (u, v, t) gives a path to v. Conversely, any path that first reaches v at t ends with an edge from a node reached by $t-1$, so deterministic SI infects v. Induction over all snapshots proves equivalence. □

Proposition 4 (Existence of the dynamics rank)**.** *For every observed binary edge-time matrix and either tolerance used here, the set in Eq. (10) is nonempty.*

Proof. At $r = K$, Eq. (3) reconstructs the centered matrix S exactly. Adding $\bar{x}\mathbf{1}^{\mathsf{T}}$ therefore gives the original binary edge-time matrix. In each snapshot, all observed links have score one and all absent links have score zero. Selecting the observed count m_t returns exactly the observed links. Hence $J_R(K) = 1$ and $\Delta_{\mathrm{SI}}(K) = 0$, which passes every $\varepsilon \geq 0$. □

B Parameters, SIR outputs, and temporal-resolution sensitivity

Table 4: Fixed analysis parameters. Generated by the supplied code.

Parameter	Value	Use
Snapshot width	30 min	All empirical data sets
Node cap	120	Activity-ranked induced subgraph
Reachability horizon	8 snapshots	4 h of clock time
Sampled windows	24	Evenly sampled active windows
Primary tolerance	5%	Worst-window criterion
Sensitivity tolerance	10%	Worst-window criterion
SIR beta / gamma	0.18 / 0.08	Per 30-min step
SIR scenarios	12 x 80 runs	Common random numbers
SIR horizon	16 snapshots	8 h of clock time
Synthetic seed	17	Rare-bridge model
SIR random seed	314159	All data sets

Table 5: SIR validation values. "Paired MAE" is the mean absolute run-level difference in final infected fraction under common random numbers.

Data set	Representation	Rank	Final size	Paired MAE
Rare-bridge model	Original	33	0.449	0.000
Rare-bridge model	Variance rank	3	0.446	0.003
Rare-bridge model	Dynamics rank	4	0.449	0.000
Hospital ward	Original	173	0.213	0.000
Hospital ward	Variance rank	96	0.212	0.005
Hospital ward	Dynamics rank	135	0.213	0.000
Hypertext 2009	Original	103	0.166	0.000
Hypertext 2009	Variance rank	52	0.165	0.007
Hypertext 2009	Dynamics rank	76	0.166	0.000
Workplace	Original	388	0.027	0.000
Workplace	Variance rank	140	0.027	0.000
Workplace	Dynamics rank	178	0.027	0.000
Primary school	Original	50	0.510	0.000
Primary school	Variance rank	29	0.514	0.006
Primary school	Dynamics rank	36	0.510	0.000
High school 2013	Original	157	0.082	0.000
High school 2013	Variance rank	57	0.082	0.003
High school 2013	Dynamics rank	78	0.082	0.000

Table 6: Temporal-resolution sensitivity. The full rank, variance rank, and dynamics rank change with bin width. Fourteen of 15 combinations have $r_{\mathrm{dyn}}(0.05) \geq r_{95}$; the 15-minute workplace case is the transparent exception.

Data set	Bin (min)	Full k	r95	r-dyn (5%)	Reach. at r95
High school 2013	15	288	102	112	0.783
High school 2013	30	157	57	78	0.625
High school 2013	60	72	33	39	0.696
Hospital ward	15	320	172	259	0.250
Hospital ward	30	173	96	135	0.091
Hospital ward	60	94	52	77	0.362
Hypertext 2009	15	190	93	136	0.000
Hypertext 2009	30	103	52	76	0.000
Hypertext 2009	60	51	28	41	0.000
Primary school	15	94	57	62	0.900
Primary school	30	50	29	36	0.500
Primary school	60	29	16	19	0.429
Workplace	15	661	224	190	0.952
Workplace	30	388	140	178	0.500
Workplace	60	204	81	99	0.500

How to read the sensitivity table. Changing the bin width changes both the number of snapshots and the numerical rank, so raw ranks should be compared within a row, not across resolutions. The consistent observation is that a variance threshold and a process threshold remain different quantities. Hospital and Hypertext require substantially more modes for process safety at every tested width. Primary school and high school show the same direction with smaller gaps. Workplace at 15 minutes is the only reversal: $r_{95} = 224$ is 34 modes larger than $r_{\mathrm{dyn}}(0.05) = 190$, and reachability at r_{95} already equals 0.952. This exception prevents an overly broad conclusion. We do not claim that process-safe rank must always exceed variance rank; we claim that the two ranks must be measured separately.

The table also shows why a reported rank is inseparable from its temporal resolution and process horizon. Halving the bin width nearly doubles the number of possible temporal degrees of freedom in several data sets, while a 60-minute bin merges more contacts and lowers the full rank. Because the analysis holds the horizon at eight snapshots, clock time also changes with bin width. The supplied workflow therefore treats bin width, node cap, horizon, and tolerance as part of the result definition. A user applying the method to another system should choose these quantities before inspecting the rank curves and should report them with the final rank.

C Reproduction checklist and searched sources

The archive is self-auditing in the following order:

1. download five official files and verify SHA-256 values;
2. build 15-, 30-, and 60-minute snapshots with documented node selection;
3. run the unit tests;
4. compute the exact temporal covariance decomposition;
5. examine consecutive integer ranks until the five-percent condition first passes;
6. run common-random-number SIR checks;
7. write CSV results, then generate every LaTeX table and figure from those results;
8. compile the manuscript and inspect every rendered PDF page.

The searched source groups were: the six 2026 Nature Portfolio articles listed in the Introduction [3–8]; temporal compression by commutators and hashing [9, 10]; graph comparison and dissimilarity [11, 12]; graph-free state prediction [13]; classic temporal-network, POD/SVD, and epidemic references [1, 2, 14, 15, 23–25]; and the official data descriptions and primary data papers [16, 18–22]. Exact DOI and URL metadata appear in the numbered reference list below.

References

[1] Petter Holme and Jari Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012. doi: 10.1016/j.physrep.2012.03.001.

[2] Naoki Masuda and Renaud Lambiotte. *A Guide to Temporal Networks*. World Scientific, Singapore, second edition, 2020.

[3] Aming Li, Yao Meng, Lei Zhou, Naoki Masuda, and Long Wang. Temporality modulates the effect of network heterogeneity on cooperation fixation. *Nature Communications*, 17:6238, 2026. doi: 10.1038/s41467-026-72717-w.

[4] Cheng Luo, Renaud Lambiotte, and Peng Ji. Temporal heterogeneity shapes diffusion dynamics in complex networks. *Nature Communications*, 17:5601, 2026. doi: 10.1038/s41467-026-72161-w.

[5] Aurélien Pélissier, Miroslav Phan, Didier Le Bail, Niko Beerenwinkel, and María Rodríguez Martínez. Unifying non-Markovian dynamics and agent heterogeneity in scalable stochastic networks. *Nature Communications*, 17:3345, 2026. doi: 10.1038/s41467-026-69817-y.

[6] Kevin Teo, Naomi Arnold, Andrew Hone, Michael Coulon, Martin Ireland, Mauricio Santillana, and István Z. Kiss. Unveiling individual and collective temporal patterns in the tanker shipping network. *Nature Communications*, 17:3300, 2026. doi: 10.1038/s41467-026-70013-1.

[7] Sheng Fang, Jun Meng, Qing Lin, Bingsheng Chen, Gaogao Dong, Chenglin Bai, Jan Nagler, Youjin Deng, Renaud Lambiotte, and Jingfang Fan. Temporal self-similarity reveals percolation universality classes in complex networks. *Nature Communications*, 17:8605, 2026. doi: 10.1038/s41467-026-75436-4.

[8] Lucas Lacasa. Fluid dynamics reduction methods for temporal networks. *Scientific Reports*, 16:18913, 2026. doi: 10.1038/s41598-026-49912-2.

[9] Andrea J. Allen, Cristopher Moore, and Laurent Hébert-Dufresne. Compressing the chronology of a temporal network with graph commutators. *Physical Review Letters*, 132:077402, 2024. doi: 10.1103/PhysRevLett.132.077402.

[10] Rémi Vaudaine, Pierre Borgnat, Paulo Gonçalves, Rémi Gribonval, and Márton Karsai. Temporal network compression via network hashing. *Applied Network Science*, 9:3, 2024. doi: 10.1007/s41109-023-00609-9.

[11] Lorenzo Dall'Amico, Alain Barrat, and Ciro Cattuto. An embedding-based distance for temporal graphs. *Nature Communications*, 15:9954, 2024. doi: 10.1038/s41467-024-54280-4.

[12] Xiu-Xiu Zhan, Chuang Liu, Zhipeng Wang, Huijuan Wang, Petter Holme, and Zi-Ke Zhang. Measuring and utilizing temporal network dissimilarity. *Communications Physics*, 8:40, 2025. doi: 10.1038/s42005-025-01940-6.

[13] Bastian Prasse and Piet Van Mieghem. Predicting network dynamics without requiring the knowledge of the interaction graph. *Proceedings of the National Academy of Sciences*, 119(44):e2205517119, 2022. doi: 10.1073/pnas.2205517119.

[14] William Ogilvy Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society A*, 115 (772):700–721, 1927. doi: 10.1098/rspa.1927.0118.

[15] Eugenio Valdano, Luca Ferreri, Chiara Poletto, and Vittoria Colizza. Analytical computation of the epidemic threshold on temporal networks. *Physical Review X*, 5:021005, 2015. doi: 10.1103/PhysRevX.5.021005.

[16] SocioPatterns Collaboration. Public temporal contact-network data sets. https://sociopatterns.org/datasets/, 2026. Accessed 26 August 2026.

[17] 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(7):e11596, 2010. doi: 10.1371/journal.pone.0011596.

[18] 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(9):e73970, 2013. doi: 10.1371/journal.pone.0073970.

[19] 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. *Journal of Theoretical Biology*, 271(1):166–180, 2011. doi: 10.1016/j.jtbi.2010.11.033.

[20] Mathieu Génois, Christian L. Vestergaard, Julie Fournet, André Panisson, Isabelle Bonmarin, and Alain Barrat. Data on face-to-face contacts in an office building suggest a low-cost vaccination strategy based on community linkers. *Network Science*, 3(3):326–347, 2015. doi: 10.1017/nws.2015.10.

[21] Juliette Stehlé, Nicolas Voirin, Alain Barrat, Ciro Cattuto, Lorenzo Isella, Jean-François Pinton, Marco Quaggiotto, Wouter Van den Broeck, Corinne Régis, Bruno Lina, and Philippe Vanhems. High-resolution measurements of face-to-face contact patterns in a primary school. *PLoS ONE*, 6(8):e23176, 2011. doi: 10.1371/journal.pone.0023176.

[22] Rossana Mastrandrea, Julie Fournet, and Alain Barrat. Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLoS ONE*, 10(9):e0136497, 2015. doi: 10.1371/journal.pone.0136497.

[23] Lawrence Sirovich. Turbulence and the dynamics of coherent structures. part i: Coherent structures. *Quarterly of Applied Mathematics*, 45(3):561–571, 1987. doi: 10.1090/qam/910462.

[24] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. doi: 10.1007/BF02288367.

[25] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *The Quarterly Journal of Mathematics*, 11(1):50–59, 1960. doi: 10.1093/qmath/11.1.50.