

# Recovering lost and absent information in temporal networks

James P. Bagrow<sup>1,2,\*</sup> and Sune Lehmann<sup>3,4</sup>

<sup>1</sup>Mathematics & Statistics, University of Vermont, Burlington, VT, United States

<sup>2</sup>Vermont Complex Systems Center, University of Vermont, Burlington, VT, United States

<sup>3</sup>DTU Compute, Technical University of Denmark, Kgs Lyngby, Denmark

<sup>4</sup>Center for Social Data Science, University of Copenhagen, Copenhagen, Denmark

\*Corresponding author. Email: [james.bagrow@uvm.edu](mailto:james.bagrow@uvm.edu), Homepage: [bagrow.com](http://bagrow.com)

July 22, 2021

**Abstract** The full range of activity in a temporal network is captured in its edge activity data—time series encoding the tie strengths or on-off dynamics of each edge in the network. However, in many practical applications, edge-level data are unavailable, and the network analyses must rely instead on node activity data which aggregates the edge-activity data and thus is less informative. This raises the question: Is it possible to use the static network to recover the richer edge activities from the node activities? Here we show that recovery is possible, often with a surprising degree of accuracy given how much information is lost, and that the recovered data are useful for subsequent network analysis tasks. Recovery is more difficult when network density increases, either topologically or dynamically, but exploiting dynamical and topological sparsity enables effective solutions to the recovery problem. We formally characterize the difficulty of the recovery problem both theoretically and empirically, proving the conditions under which recovery errors can be bounded and showing that, even when these conditions are not met, good quality solutions can still be derived. Effective recovery carries both promise and peril, as it enables deeper scientific study of complex systems but in the context of social systems also raises privacy concerns when social information can be aggregated across multiple data sources.

## 1 Introduction

Temporal networks are of increasing importance as models for systems across social, biological and technological domains [1, 2, 3]. Their study has been predicated on data and high-resolution, high-velocity data are becoming more common across fields [4]. In a temporal network, each link turns on and off over time, forming an intricate pattern of dynamic connectivity [5]. In many real-world instances—such as social networks captured from smartphones [6], gene regulatory networks obtained via high-throughput experiments [7], and brain networks measured through neuroimaging [8, 9]—these rich dynamics are unavailable. Instead, we can only know the activation patterns of nodes: the information regarding when a node connects to some other node, but not which specific one. In some cases, this is by design: to preserve user privacy, a smartphone maker may prevent an installed application from accessing user

communication records [10, 11, 12], but allow that application access to, for example, geolocation [13] or accelerometer signals [14]. In other cases, it may be too costly or otherwise infeasible to monitor the activities of every edge. In many biological systems, we simply lack the experimental tools to directly capture *in situ* edge activities [7, 15, 9].

While node activities are generally more available in empirical temporal network data, the loss of information compared to edge activities is currently not understood. This leads us to ask the following question: if a researcher loses, or otherwise cannot ascertain, the edge activity data, but does have access to the node activities, how well can they recover the richer edge activity data? Solely based on nodal activation, at best, this recovery is only possible in a small number of special cases. However, if the network’s static structure is available, e.g. either directly available from a separate measurement or can be inferred via side information, then it may be possible to recover, even if only approximately, the absent edge activity data. Such side information can take the form of additional experiments or joining across datasets. For brain networks, for example, tract-tracing [16] or fluorescence microscopy [17] studies revealing network connectivity can complement functional neuroimaging [9], while for social behavior, data brokers [18] may be able to gather the structure of user social networks by cross-referencing public and private social activity datasets [11, 12].

There is both promise and peril in the recovery of edge activities. In the case of biological systems, this opens up new possibilities for understanding the fundamental dynamics of living cells. In the context of social networks, rich data raises privacy concerns. Data brokers, who may be able to aggregate across both public and nonpublic sources of information, have the potential to circumvent privacy protections and reveal more about individuals than would be possible without such high-resolution data.

In this paper, we make the following contributions. In Sec. 2 we define and motivate the problem of information loss in temporal networks and how to recover it. We apply solution methods to a representative network corpus, including studying how well network analysis tasks can be performed on recovered data (Sec. 3) and better understanding both theoretically and empirically the topological and dynamical sources of recovery error (Sec. 4). We conclude with a discussion providing context for the recovery problem across different scientific domains in Sec. 5, including both the benefits for advancing network measurement and the concerns with protecting privacy raised by this work.

## 2 Information loss and recovery in temporal networks

Suppose a researcher is interested in studying a time-evolving, or temporal, network, fully defined through its edge activations. For many networks, such edge data are not always readily available. It may be impossible, costly or otherwise impractical to monitor the activities of every edge in the network. In some cases, one must even observe

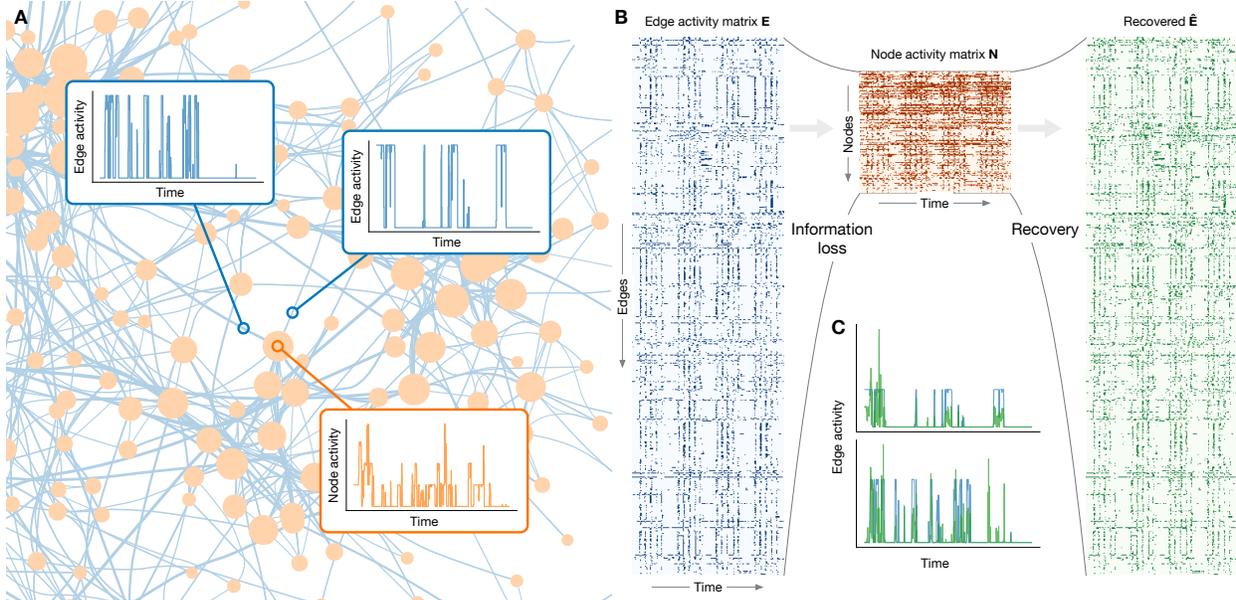


Figure 1: Information loss and recovery in temporal networks. **(A)** Using the Copenhagen Network to illustrate two edge activity time series' and a node's corresponding activity (aggregated for all its incident edges). **(B)** Significant information is lost when only the much smaller node activity matrix is available. (The Copenhagen network has an average degree  $\langle k \rangle \approx 8.8$  meaning there are  $\approx 4.4$  edges per node and thus an over four-fold reduction in information when going from  $\mathbf{E}$  to  $\mathbf{N}$ ) However, as the network topology is learned, the edge activity matrix can be recovered, potentially with high accuracy,  $\mathbf{E} \approx \hat{\mathbf{E}}$ . **(C)** Comparing the original and recovered edge activities for the two edges highlighted in panel A.

every pair of nodes, a costly, quadratic operation. Therefore, in practice one often lacks the ideal edge activity data. We argue that progress on studying the network's time dynamics can still be made, however, by studying the time series of node activities, when available. Node activities, which describe when a node is active but not with whom, are less informative but more readily available, as monitoring a singleton is less resource-intensive than monitoring a pair. Further, there are fewer singletons that need to be monitored than duos, meaning there will be (unless the network is exceptionally sparse) fewer node time series than edge time series. Yet, while node activity data are more practical, the loss of information is typically significant. The amount of information lost depends on network structure, which determines how edge activities are projected down to the node activities. (See Fig. 1A, illustrated using 'Copenhagen', the physical proximity component of the Copenhagen Networks Study dataset [19], a smartphone-derived social network which is part of our corpus of temporal networks; described below).

We now formalize the information loss and recovery problem. Consider a network  $G = (V, E)$  of  $n = |V|$  nodes (or vertices) and  $m = |E|$  edges (or links). Let  $\mathbf{E}$  be an  $m \times T$  matrix of edge activity data, where  $E_{ij,t}$  is the number of interactions between  $i$  and  $j$  during timestep  $t$  (or during a time window  $(t, t + \Delta t)$ ). Each row of  $\mathbf{E}$  corresponds to a time series of  $T$  values for an edge in the network. These time series' encode the temporal activities of the network, when edges are active and at what strength. A reduced view of the temporal network can be given by the node activity

matrix  $\mathbf{N}$ . Here  $\mathbf{N}$  is an  $n \times T$  matrix where each row captures the activity of a corresponding node in the network.

The edge and node activity matrices are related to one another via the graph *incidence matrix*  $\mathbf{B}$ . The (unoriented) incidence matrix is an  $n \times m$  binary matrix that details which edges are incident on a node. Using  $\mathbf{B}$ , we then have

$$\mathbf{B}\mathbf{E} = \mathbf{N}. \tag{1}$$

This relationship shows us that the edge activity matrix is more fundamental than the node matrix given the network structure, as  $\mathbf{N}$  can be derived from  $\mathbf{E}$ . Figure 1B illustrates this information loss by the reduction in size as  $\mathbf{E}$  becomes  $\mathbf{N}$  (Eq. (1)).

Now, suppose we do not have the edge activities but we do have the node activities. The question now becomes: can we recover  $\mathbf{E}$  from  $\mathbf{N}$  using the network structure (Fig. 1B,C)? Equation (1) is a system of linear equations in the standard form  $Ax = b$ , which can usually be solved by, for example,  $x = A^{-1}b$ . However,  $\mathbf{B}$  is a rectangular matrix and, except for certain edge cases such as a network with no cycles, Eq. (1) is an underdetermined system. This means there will be either no solutions or an infinite number of solutions  $\hat{\mathbf{E}}$  such that  $\|\mathbf{B}\hat{\mathbf{E}} - \mathbf{N}\|_F^2 = 0$  ( $\|\cdot\|_F$  denotes the Frobenius norm). From this, one can conclude that it is generally not possible to fully recover  $\mathbf{E}$  given  $\mathbf{N}$  and  $\mathbf{B}$ . However, while exact recovery of  $\mathbf{E}$  may not be possible, approximations may still be possible and, if accurate (i.e., if  $\hat{\mathbf{E}} \approx \mathbf{E}$ ), they may be sufficient for subsequent network analysis tasks (see also Sec. 3).

The classical approach to solving the underdetermined system (1) is to find the least-norm solution [20]. While  $\mathbf{B}^{-1}$  in general does not exist, the pseudoinverse  $\mathbf{B}^+$  can be efficiently computed, using, for instance, singular value decomposition (SVD), and then used to find the minimum or least-norm solution:  $\hat{\mathbf{E}} = \mathbf{B}^+\mathbf{N}$ . We show below (Sec. 4) that the accuracy of information recovery using the least-norm solution depends on a combination of topological and dynamical sparsity.

In this context sparsity is also the fundamental drawback with respect to the least-norm solutions—they do not promote sparse solutions, meaning that  $\hat{\mathbf{E}}$  will be a dense matrix, which is unlikely to be a good representation of  $\mathbf{E}$  unless all edges are active at all times. A period of little network activity may correspond to many zeros appearing within the time series. Moreover, sparsity in the time series may be intermixed with sparsity in the network structure.

To address this, sparsity-promoting solution techniques, optimization problems that explicitly reward zero values in solutions, have in recent years become a pillar of statistical learning methods [21, 22, 23, 24]. Thus, to improve on the classical approach, we formulate finding sparse solutions  $\hat{\mathbf{E}}$  to Eq. (1) as an optimization problem whose objective function consists of a least-squares error term and a regularization or penalty term that promotes zero values in discovered solutions (Supporting Information (SI) Eq. (S1)). This optimization problem can be interpreted as a (multi-target) Lasso regression [21] without an intercept term and with an added constraint enforcing nonnegative

regression coefficients, capturing the properties of an edge activity matrix. We use an  $L_1/L_2$  regularization to find sparse solutions for multiple problems (columns of  $\mathbf{N}$  and  $\hat{\mathbf{E}}$ ) jointly by treating each edge as a group over time [25, 26]. Further, this formulation is related to the similar problem of compressed sensing [27, 28]. In our case, solutions were found using coordinate descent; see SI Sec. S2 for full details, including implementation details and Bayesian selection of the regularization hyperparameter<sup>1</sup>.

To understand how well temporal network information can be recovered using the least-norm and sparse methods, we assembled a temporal network corpus consisting of five networks representing a range of different systems. The first, ‘Copenhagen’, used in Fig. 1 to illustrate the information loss and recovery problem, is a social network derived from smartphone bluetooth proximity data that serve as a proxy of face-to-face interactions [19, 29]. The remaining networks are ‘Hospital’, another proximity-based social network derived from wearable sensors carried by healthcare workers [30]; ‘Ant Colony’, a physical interaction network taken from manually-annotated video footage of a *Temnothorax rugatulus* colony [31]; ‘Manufacturing Email’, an interaction network derived from internal emails sent between employees of a mid-size manufacturing firm [32]; and ‘College Message’, a social network derived from messages sent between members of an online community of University of California, Irving college students [33]. These networks cover a span of time scales, from minutes (Ant Colony) to days (Hospital), weeks (Copenhagen) and months (Manufacturing Email, College Message). Full details for all networks, including data processing steps and network statistics, are given in SI Sec. S1.

For each network, information loss (replacing  $\mathbf{E}$  with  $\mathbf{N}$ ) was simulated using Eq. (1). Figure 2 shows heatmaps of the node activity and edge activity matrices for each network. In the figure, rows of the node and edge matrices are drawn to scale to illustrate the extent of information loss across the corpus. For each network, we report  $n$  and  $m$ , the numbers of nodes and edges; their ratio  $m/n$  describes the “aspect ratio” of the incidence matrix  $\mathbf{B}$  which reflects how much information was lost when moving from  $\mathbf{E}$  to  $\mathbf{N}$  and consequently how challenging we may expect the recovery problem to be.

Figure 2 also shows the  $\hat{\mathbf{E}}$  recovered from  $\mathbf{N}$  for both the least-norm and sparse solution methods discussed in Sec. 2. (All matrices in a given row in Fig. 2 use the same colorbar.) Comparing these recovered matrices to  $\mathbf{E}$  shows that both approaches managed to capture the qualitative overall features of  $\mathbf{E}$ . However, we see that the sparse  $\hat{\mathbf{E}}$  solution is far closer in appearance to the true  $\mathbf{E}$ .

In particular, the sparsity pattern of  $\mathbf{E}$  (zero entries in matrices are illustrated with white in Fig. 2) is almost entirely absent in the least-norm solutions while being captured well in the sparse solutions. Qualitatively, visual comparison of  $\hat{\mathbf{E}}$  (sparse) and  $\mathbf{E}$  demonstrates the potential for quite good information recovery. Quantitatively, Table 1 measures

---

<sup>1</sup>An implementation is available at [github.com/bagrow/recovering-information-temporal-networks](https://github.com/bagrow/recovering-information-temporal-networks).

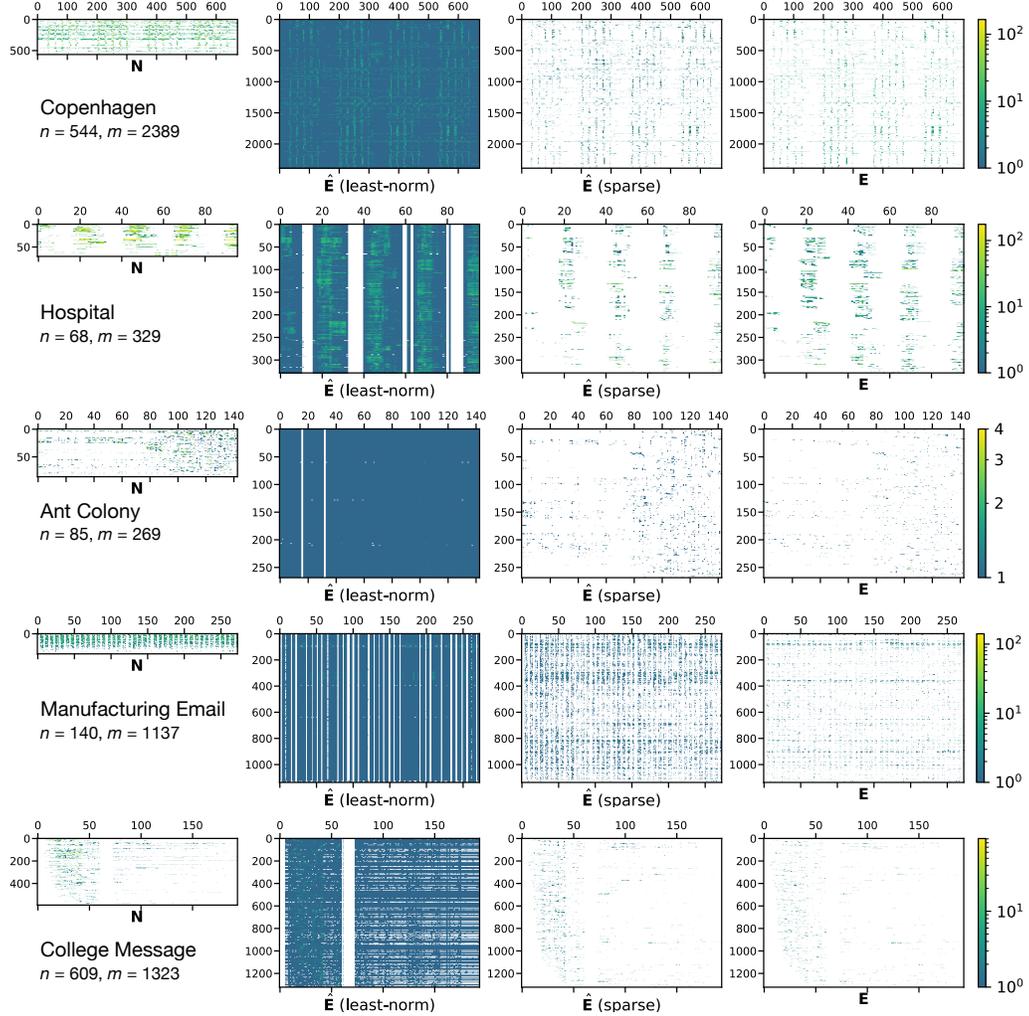


Figure 2: Good quality  $\hat{\mathbf{E}}$  matrices are found across most networks in a representative corpus of temporal networks. One notable exception is Mfg. Email. Each row corresponds to one network with the (logarithmic) color scale, indicating the magnitude of matrix elements, shared across the row. Solution sparsity is illustrated with white color and is well recovered, except for Mfg. Email, when comparing  $\hat{\mathbf{E}}$  (sparse) and  $\mathbf{E}$ . Node and edge activity matrices are drawn to scale—the small size of the  $\mathbf{N}$  compared to  $\mathbf{E}$  underscores the recovery challenge. Least-norm solutions are dense except for time periods of zero edge activity, manifesting as vertical bands in all networks except Copenhagen; see SI for details.

the recovery accuracy by reporting the correlations between matrix entries. We report both Pearson  $r$  and Spearman  $\rho$  correlations to capture both linear and nonparametric associations between  $\mathbf{E}$  and  $\hat{\mathbf{E}}$ . Despite the challenge we would expect when attempting to recover the lost temporal information, when examining the table we see that many networks show a high correlation between the original  $\mathbf{E}$  and the recovered  $\hat{\mathbf{E}}$ . One exception is Manufacturing Email, where both solutions achieved correlations lower than 0.5. Even though it is not the largest network in the corpus, Manufacturing Email is the densest, with  $m/n = 8.12$ . This relatively high density leads to increased aggregation of the edge time series' and thus we expect that undoing the subsequent information loss will be more difficult as a result.

	$r$ (ln)	$r$ (sp)	$\rho$ (ln)	$\rho$ (sp)
Copenhagen	0.7660	0.5060	0.4512	0.7308
Hospital	0.5494	0.5752	0.4067	0.6004
Ant Colony	0.5698	0.8432	0.2727	0.6844
Manufacturing Email	0.4275	0.2650	0.3957	0.3879
College Message	0.7104	0.8788	0.2508	0.7560

Table 1: Correlations between original  $\mathbf{E}$  and recovered  $\hat{\mathbf{E}}$ . Here we report the Pearson correlation  $r$  and Spearman correlation  $\rho$  between elements of  $\mathbf{E}$  and  $\hat{\mathbf{E}}$ . In most instances, high correlations were found, indicating successful recovery of  $\mathbf{E}$ , with Mfg. Email being a notable exception. Often the sparse solution outperforms the (dense) least-norm solution. ln: least-norm; sp: sparse.

We further explore the effects of network structure and the intersecting roles of topological and dynamical sparsity in Sec. 4.

### 3 Network tasks after recovery

The recovered edge activity data will only be useful if the data are useful for subsequent tasks. Here we examine how well the recovered data can be used on two exemplar tasks: estimating tie strength from temporal activity and extracting the network’s multi-scale backbone [34].

Tie strengths, typically represented with edge weights  $w$  accounting for the total quantity of interactions between nodes, are captured by summing  $\mathbf{E}$  over time periods:  $w_{ij} = \sum_t E_{ij,t}$ . When only  $\hat{\mathbf{E}}$  is available, how well does  $\hat{w}_{ij} \approx w_{ij}$ ? We computed  $\hat{w}$  using three methods. One, a baseline, uses a linear kernel to bypass the computation of  $\hat{\mathbf{E}}$  and compute  $\hat{w}$  directly from  $\mathbf{N}$ :  $\hat{w}_{ij} = \mathbf{N}_i \mathbf{N}_j^T$ , where  $\mathbf{N}_i$  is the row vector corresponding to node  $i$  in  $\mathbf{N}$ . The second and third methods estimate  $\hat{\mathbf{E}}$  using the least-norm and sparse solution, respectively, and summing  $\hat{\mathbf{E}}$  over time periods to derive  $\hat{w}$ . After computing each  $\hat{w}$  using  $\hat{\mathbf{E}}$  we compare to  $w$  derived directly from  $\mathbf{E}$ .

As shown in Fig. 3, good recovery of tie strength is possible across the network corpus. In particular, the sparse solution more accurately infers the underlying tie strength than either the node-based baseline or the least-norm approach. Three networks show correlations  $r(w, \hat{w}) > 0.8$  and even the most difficult network, Manufacturing Email, shows  $r > 0.5$ , better performance for that network than may be expected from Table 1. While all networks show high recovery performance, the sparse solutions are especially accurate for the Copenhagen, Ant Colony, and College Message networks, achieving a 23.6% (Copenhagen) to 45.4% (Ant Colony) improvement in correlation  $r$  compared to the least-norm approach.

While the high performance of the sparse solution naturally follows from its accuracy when inferring  $\mathbf{E}$ , in some ways, it is also surprising. The sparsity-promoting penalty imposed on solutions induces a bias in the solution and we might expect that bias to reduce the accuracy of the aggregated quantity  $\hat{w}$  as opposed to the least-norm solution where

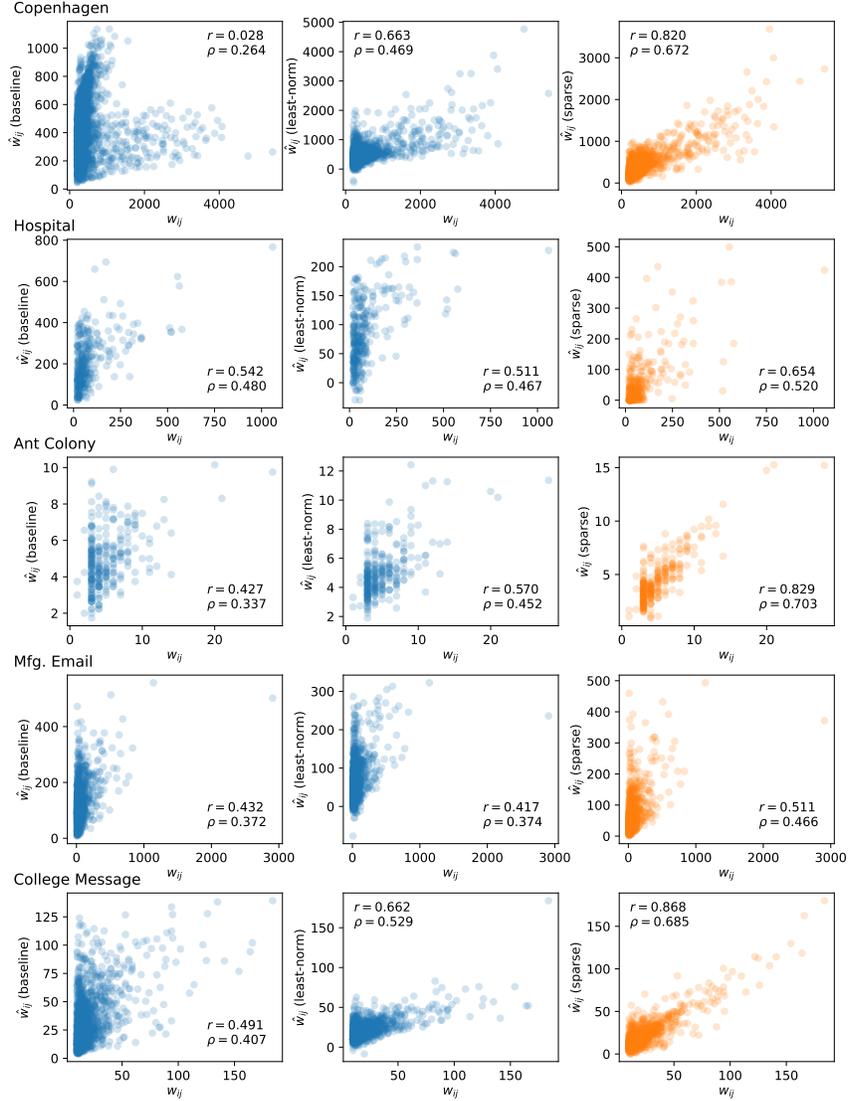


Figure 3: Quality of tie strength inference using recovered  $\hat{\mathbf{E}}$ . Here we compare the edge weights  $w$ , computed from  $\mathbf{E}$ , with  $\hat{w}$  computed directly from  $\mathbf{N}$  (baseline) and from the least-norm and sparse  $\hat{\mathbf{E}}$ . In all cases, the sparse solutions (right column) best recover the edge weights.

solution errors over and under  $\mathbf{E}$  may cancel out during summation. That correlations remain high despite this further underscores the usefulness of sparse estimation.

Our second example of a post-recovery task is to extract a network’s multiscale backbone [34]. The multiscale backbone represents the central or key set of edges undergirding the network’s topology. The multiscale backbone is found by examining the local distribution of edge weights [34]. Specifically, an edge  $(i, j)$  incident to node  $i$  is flagged as belonging to the backbone if  $(1 - p_{ij})^{k_i - 1} < \alpha_B$ , where  $p_{ij} = w_{ij} / \sum_j w_{ij}$ ,  $k_i$  is the degree of  $i$ , and  $\alpha_B$  is the

multiscale backbone strength parameter. (Technically, edge  $(i, j)$  is retained if this inequality is satisfied on either  $(i, j)$  or  $(j, i)$ ; see Serrano *et al.* [34] for details.)

For each network, we compare the backbone inferred from the original data with the backbone found on the recovered data. To do so, we interpret backbone discovery as a binary classification or partitioning of edges: each edge is flagged as either a backbone edge or non-backbone edge and we compare these classifications using false positive and true positive rates. However, there is an additional layer of complexity: as the multiscale backbone method depends on parameter  $\alpha_B$ , the “true” solution (i.e., the solution found on  $\mathbf{E}$ ) will depend on  $\alpha_B$  and an experimenter may not know  $\alpha_B$  when working with  $\hat{\mathbf{E}}$ . We therefore distinguish between two parameters by denoting as  $\alpha_B$  the value used on  $\mathbf{E}$  and  $\alpha'_B$  the value used on  $\hat{\mathbf{E}}$  and we consider both cases  $\alpha_B = \alpha'_B$  and  $\alpha_B \neq \alpha'_B$  when comparing results.

To illustrate the extracted backbones, we begin by drawing the Copenhagen network twice in Fig. 4A, once highlighting the edges of the “oracle” backbone computed using  $\mathbf{E}$  and again highlighting the backbone found using the sparse  $\hat{\mathbf{E}}$ . While there are some visual discrepancies between the two backbones, overall they agree considerably. Moving to a quantitative assessment, in Fig. 4B we plot the Receiver Operating Characteristic (ROC) curve comparing  $\mathbf{E}$ -backbones and  $\hat{\mathbf{E}}$ -backbones on the Copenhagen network. The dashed line illustrates the expected classifier performance if randomly labeling edges as backbone or non-backbone; we see generally far better performance (as typified by low false positive rates and high true positive rates) regardless of  $\alpha_B$  and  $\alpha'_B$ .

Moving beyond the Copenhagen network, Fig. 4C shows classifier performance across the entire corpus, for both the least-norm and sparse recovered  $\hat{\mathbf{E}}$ . We simplify the ROC curve by plotting the AUC (area under the ROC curve) as a function of  $\alpha_B$  (individual ROC curves for all methods and networks are shown in Sec. S3, Fig. S3). Overall, we observe high AUC values, far from the  $\text{AUC} = 1/2$  value for a null classifier. Further, for any given network and value of  $\alpha_B$ , the sparse solution performs better at classifying backbone edges than the least-norm solution. Performance degrades as  $\alpha_B$  increases, when edges were less selectively added to the “true” backbone, and improves considerably at lower  $\alpha_B$ , when edges were more selectively added. This trend in performance indicates that the strongest and thus most important members of the multiscale backbone can be inferred accurately even when the detailed information within  $\mathbf{E}$  is lost or absent.

## 4 Understanding recovery errors

We now explore the errors made when recovering  $\mathbf{E}$  using the classic, least-norm and modern, sparsity-promoting solution methods applied above. Using our network corpus, we also relate how information loss is affected by different aspects of network structure and temporal activity.

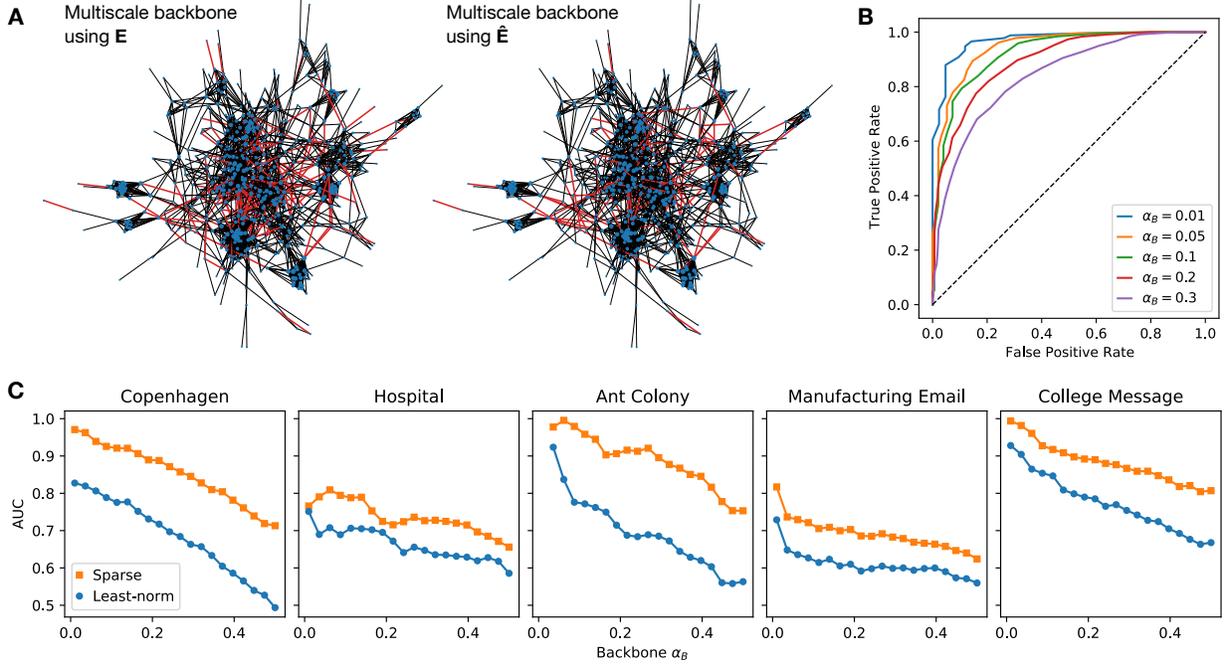


Figure 4: Extracting multiscale network backbones using recovered data. (A) The multiscale backbones ( $\alpha_B = 0.125$ ) using the original and recovered edge activity data for the Copenhagen social network. Red denotes edges flagged as belonging to the backbone. (B) Comparing backbones inferred from the  $\hat{\mathbf{E}}$  (sparse) for the Copenhagen network when  $\alpha_B \neq \alpha'_B$ . Here each ROC curve corresponds to a different choice for  $\alpha_B$  while the curve itself is parameterized by sweeping through values  $\alpha'_B$ . The diagonal line denotes the null case where edges are randomly labeled as backbone or non-backbone. ROC curves for all networks are shown in Fig. S3. (C) Expanding upon panel B to the entire corpus, we report the area under the curve (AUC) of ROC curves.

It is instructive to first consider the classic approach to solving the underdetermined system (Eq. (1)). Using the pseudoinverse  $\mathbf{B}^+$  to compute the least-norm solution will not be perfectly accurate:

$$\hat{\mathbf{E}} = \mathbf{B}^+ \mathbf{N} = \mathbf{B}^+ \mathbf{B} \mathbf{E} \neq \mathbf{E} \quad (2)$$

because  $\mathbf{B}^+ \mathbf{B} \neq \mathbf{I}$  since  $\mathbf{B}^+$  is a right-inverse, not a left-inverse, of  $\mathbf{B}$ . Further, since  $\|\mathbf{B} \hat{\mathbf{E}} - \mathbf{N}\|_F^2 = 0$ , we cannot characterize the solution using the residual. Instead, we consider in theory the difference between  $\mathbf{E}$  and  $\hat{\mathbf{E}}$ . We show (Sec. S4) that

$$\|\hat{\mathbf{E}} - \mathbf{E}\|_F \leq \sqrt{m-n} \|\mathbf{E}\|_F, \quad (3)$$

where  $\|\cdot\|_F$  is the Frobenius norm. In other words, the error bound Eq. (3) tells us how much our recovered edge activity matrix  $\hat{\mathbf{E}}$  will be near to or far from the actual edge activity matrix  $\mathbf{E}$  depends on the product of two terms, the norm of the true solution  $\|\mathbf{E}\|_F$  and the density of the network, specifically  $\sqrt{m-n}$ .

Both terms on the right of Eq. (3) relate the difficulty of the recovery problem to the sparsity of the data. Denser networks will have larger  $m$  relative to  $n$ , leading to a wider  $\mathbf{B}$ . Likewise, networks with denser temporal activity will

have more edges active at any given time, leading to fewer zero elements in  $\mathbf{E}$  and potentially larger values for non-zero elements, both of which contribute to a larger  $\|\mathbf{E}\|_F$ . As either form of sparsity decreases, the gap between  $\hat{\mathbf{E}}$  and  $\mathbf{E}$  is likely to grow, leading to larger recovery errors.

Turning from the least-norm solution, we now focus on the sparse solution method. Sparse estimation techniques have a rich history of theoretical study [35, 22, 26, 24]. Here it is key to understand the shape of the optimization problem’s loss function. If the loss function is strongly convex, then a single minimum exists and, for optimizations such as ours, theoretical bounds can be put on the estimation error. If, however, the loss function is not strongly convex but only convex, then there will exist directions for which the loss is flat, and the optimization solution can become arbitrarily far from the minimum. For our optimization, at timestep  $t$ , i.e., column  $t$  of  $\mathbf{E}$  and  $\mathbf{N}$ , the least-squares loss  $f(\mathbf{E}_t) = \frac{1}{2n} \|\mathbf{N}_t - \mathbf{B}\mathbf{E}_t\|_2^2$  is always convex. It is strongly convex when the  $m \times m$  Hessian matrix for this loss,  $\nabla^2 f = \mathbf{B}^T \mathbf{B}/n$ , has eigenvalues bounded away from zero. Unfortunately, a Gram matrix of the form  $\mathbf{B}^T \mathbf{B}$  where  $\mathbf{B}$  is  $n \times m$  will have rank at most  $\min\{n, m\}$  and so will be rank-deficient in the high-dimensional setting ( $n < m$ ) [24]. Standard practice then is to seek weaker requirements than strong convexity for the loss function, known as *restricted strong convexity* [36], where strong convexity is required only for a subset of estimation error vectors. For a linear model, this means that we seek to lower bound the *restricted eigenvalues* (RE) of  $\mathbf{B}^T \mathbf{B}$  [37]. If zero is an eigenvalue of  $\mathbf{B}^T \mathbf{B}$ , we will be unable to guarantee a bound on the estimation error, although in practice it is still possible to find good solutions [22].

By combining results from the literature of sparse statistical learning with results from spectral graph theory, we prove (Sec. S5) that RE holds when the smallest eigenvalue of the adjacency matrix of the network’s *line graph* is greater than  $-2$ . This occurs only when every connected component of the network is a tree or has one cycle only, of odd length [38]. Intuitively, this makes sense: the incidence matrix of a tree for instance will be narrower than it is wide ( $m/n < 1$ ) and the system (Eq. (1)) will not be underdetermined. Yet this is also a strong requirement—no networks in the corpus are close to being trees. However, what matters for RE to hold is the network’s structure at each time step, not the overall structure (see SI), and some networks, Ant Colony and College Message in particular, often meet this requirement (Fig. S4). And, as in many practical situations, even without the theoretical guarantee, we have observed here successful use of recovered data in practice (Figs. 3 and 4 above).

Beyond the theoretical analysis, let us examine the estimation error between  $\hat{\mathbf{E}}$  and  $\mathbf{E}$  empirically, using our network corpus. In Fig. 5A,B we plot the (relative) estimation error at time  $t$  as a function of the active node  $n_t/n$  and active edge fractions  $s_t/m$ , respectively. Both active fractions strongly correlate with estimation errors across the network corpus. Somewhat surprisingly, we find comparable correlations with error for both  $n_t/n$  and  $s_t/m$ :  $r = 0.831$  and  $r = 0.825$ , respectively. We expand further on this correlation in SI Table S2. Being able to use  $n_t$  to judge the

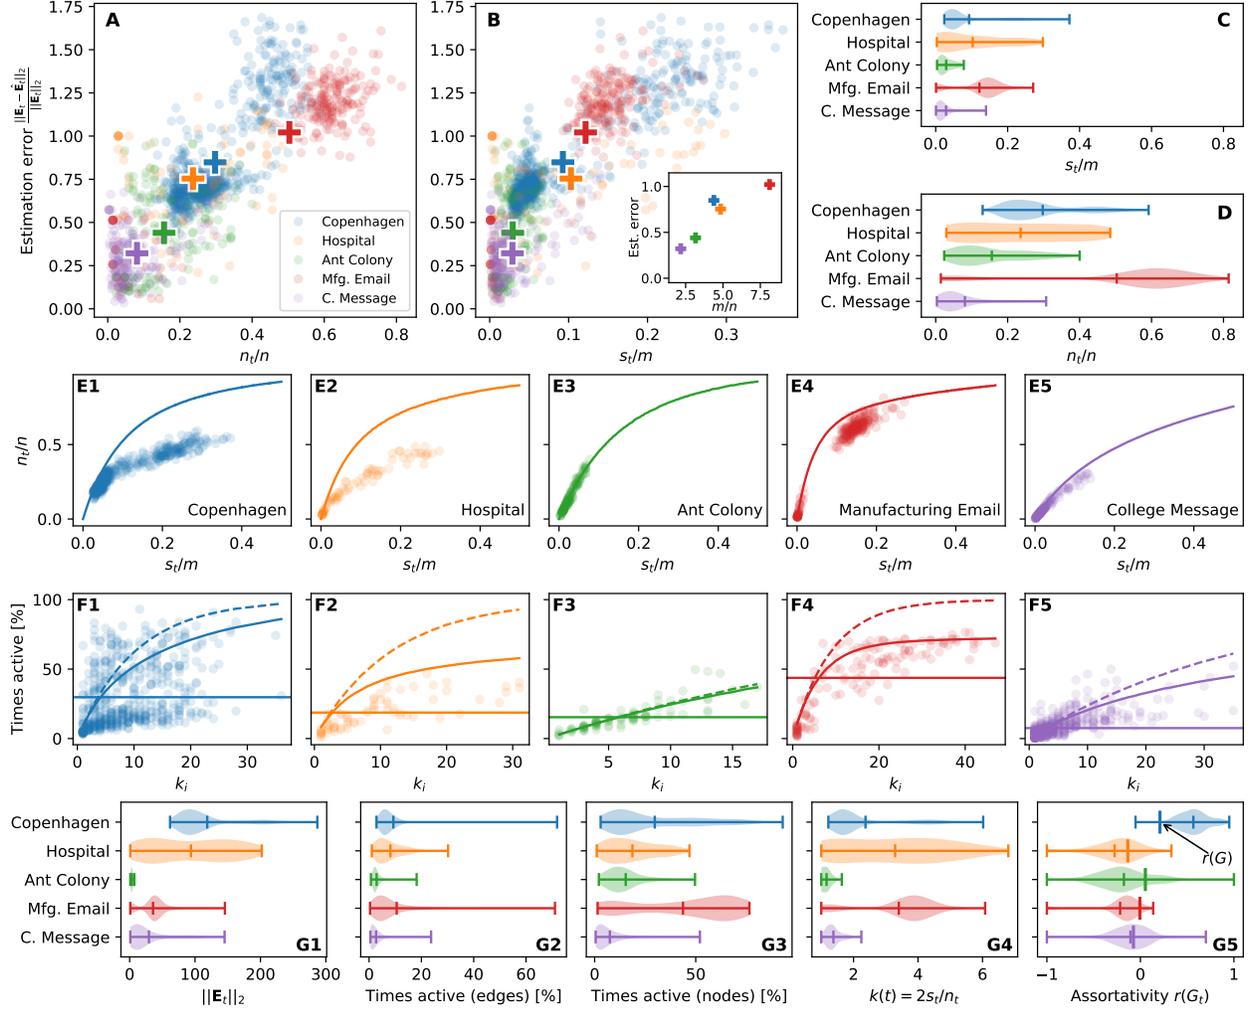


Figure 5: Topological and dynamical characteristics of recovery error. (A) Estimation error at time  $t$  depends on active node fraction  $n_t/n$  and (B) active edge fraction  $s_t/m$ . Crosses in A and B denote the overall average per network. (inset) the average estimation error is well approximated by the network density  $m/n$ . (C, D) Distributions of active edges and nodes. Most networks are denser in node activity than in edge activity. (E) Observed relation between active nodes and edges. Lines denote the expected  $n_t/n$  if  $s_t$  distinct edges, chosen uniformly at random, are active. Ant Colony is well explained by this independence null, while other networks, particularly Copenhagen and Hospital deviate at high  $s_t$  from the null. (F) The proportion of times a node  $i$  is active (has at least one active incident edge) as a function of its degree  $k_i$ . Curved lines denote expectations for null models for edge activity that preserve the overall sparsity of  $E$  (dashed) and the sparsity of  $E_t$  at all times  $t$  (solid). Horizontal lines denote the mean fraction of times active. (G) Distributions of dynamical and network features across the network corpus.

magnitude of errors just as well as  $s_t$  is interesting because in practice the actual value of  $s_t$  will not be known, but  $n_t$  will be. A strong correlation here would allow a researcher to anticipate the magnitude of errors made during recovery, an important step to support their investigations in practice, when true solutions are unknown. In that regard, even a simple measure of overall density,  $m/n$ , which is also the “aspect ratio” of B, is a good predictor for the average estimation error (Fig. 5B inset).

Further examining the scatter plots in Fig. 5A,B we see that the two sparsest networks, Ant Colony and College Message, have lower active fractions and lower recovery errors. Conversely, the densest network, Manufacturing Email, has a very high active node fraction and high estimation errors. Copenhagen occupies an interesting middle ground, with a bimodality in active fraction and estimation error, meaning there are time periods where the network is more strongly activated and recovery is difficult and other periods where the network’s activity is sparse and recovery is less error-prone. Examining the distributions of active fractions over time (Fig. 5C,D) we see that only a minority of edges are active at a given time (C) but those edges will land upon and activate a larger fraction of the nodes (D). This is most extreme for Manufacturing Email, where typically less than 20% of active edges will give rise to 60% active nodes.

While the distributions of node and edge activity fractions are useful to study independently, the interaction of the two also matters. In particular, do active edges follow a pattern in their distribution over the network, or are they effectively randomly distributed? In Fig. 5E we plot for each network  $n_t/n$  vs.  $s_t/m$  over  $t$ . Along with these scatter plots we include a randomized null model where we selected the same number  $s_t$  of distinct edges uniformly at random and then determined  $n_t$  from how many unique nodes were incident upon that set of edges. For small  $s_t$  relative to  $m$ , we expect  $n_t = 2s_t$  as it will be unlikely for two edges chosen at random to fall on a common node (cf. graph matchings [39]). Indeed, this appears to hold for Ant Colony and to a lesser extent College Message; both are well described by the randomized null model. In comparison, Copenhagen and Hospital, both proximity-based social networks, stay close to the null model only for the smallest values of  $s_t$ . At higher values of  $s_t$ , both deviate significantly, indicating that edge activations cluster around a comparatively smaller set of nodes. This clustering of active edges will make the recovery problem more challenging. Manufacturing Email shows a similar clustering but to a lesser extent, while the high number of nodes active at the same time (cf. Fig. 5D) makes recovery a challenge. Notice that transforming from  $n_t$  and  $s_t$  to  $n_t/n$  and  $s_t/m$  introduces a scale-dependency related to the overall density of the network:  $n_t = 2s_t$  becomes  $n_t/n = \langle k \rangle s_t/m$  where  $\langle k \rangle = 2m/n$  is the average degree of the network and, further, is twice the aspect ratio of **B**. In many ways, the steepness of the curves in panel E at small  $s_t$  summarizes the topological difficulty of the recovery problem.

Another view for how edge activations distribute over nodes is given in Fig. 5F, where we study the fraction of times a node is active as a function of its degree. As only a single incident active edge is needed to activate a node, we anticipate that high degree hub nodes will be disproportionately activated compared to low degree nodes. In fact, all networks in the corpus exhibit this trend. We also observe that the steepness of this trend correlates well with the difficulty of the recovery problem. For comparison, each scatter plot in Fig. 5F also shows two null models (curved lines) capturing the expected proportion of times active if edges are activated at random. The first null model (dashed lines) preserves the overall number of active edges (the number of nonzero elements in **E**) while the second, more

strict model (solid lines) preserves the active number of edges at each time  $t$  (the number of nonzero elements in each column  $\mathbf{E}_t$ ). Compared to the actual times active, some networks are well explained by one null model (Manufacturing Email, College Message) or even both (Ant Colony). Other networks (Copenhagen, Hospital) are not explained by either null, indicating a non-uniformly random distribution in how active edges appear across high- and low-degree nodes. Copenhagen is especially interesting to point out as the bimodality seen above (Fig. 5A,B) exhibits here as two regimes of nodes, one active far less often than expected, and another active as much or more often than expected from the null models.

Lastly, we examine a number of other dynamical and network properties in Fig. 5G. Figure 5G1 shows distributions of  $\|\mathbf{E}_t\|_2$ , describing the (unnormalized) magnitudes of dynamical activity we can expect on edges in the network; Copenhagen, the largest network, exhibits a few periods of very high activity but otherwise exhibits a typical scale similar to Hospital. Hospital, however, exhibits a wide, nearly uniform spread in low and high activity. Other networks, especially the sparse Ant Colony, show less activities per timestep. Next, supplementing the distributions of  $s_t/m$  and  $n_t/n$  (Fig. 5C,D), Fig. 5G2 and Fig. 5G3 show the distributions of times active for both nodes and links. These follow the active fractions closely but two networks, Copenhagen and Manufacturing Email, have subsets of frequently active edges not apparent in the distributions of  $s_t/m$  shown in Fig. 5C. Another measure of dynamical network density, the temporal mean degree  $k(t) = 2s_t/n_t$  is shown in Fig. 5G4 (excluding times  $t$  where  $n_t = 0$ ). Here we can see that Hospital and Manufacturing Email have similar high densities, but Manufacturing Email is consistently dense while Hospital has periods of high and low density, in line with the spread of estimation errors seen in Fig. 5A,B. Finally, we report in Fig. 5G5 the degree assortativities [40]  $r(G_t)$  of each network over time, where  $G_t$  is the network of active nodes and edges at time  $t$ . Most networks tend to be degree dissortative ( $r < 0$ ) at any given time (with Copenhagen being a notable exception), meaning that links tend to form between high and low degree nodes. Often the time-dependent network  $G_t$  is more dissortative than the static, cumulative network  $G$ , with Copenhagen again being an exception. See SI Sec. S5.2, Fig. S6 for further comparison between network features and recovery errors. In general, the interplay between both static and time-varying network properties (Fig. 5G), coupled with non-random patterns of edge activations (Fig. 5E,F) will influence the challenge of the recovery problem, as seen here in the per-timestep estimation errors (Fig. 5A,B).

## 5 Discussion

We have studied the problem of information loss in temporal network data. We asked the commonly occurring question: when detailed edge activity data are unavailable but node activities are, can edge activities be recovered from

node activities and the static network itself? This recovery problem (generally) maps into an underdetermined linear system governed by the network structure, and we examined how to recover information loss using both classical and modern approaches to finding solutions to underdetermined systems. Using both theoretical analyses and empirical investigations, we showed that the difficulty of information recovery is governed by a decomposition of the topological and dynamical properties of the network being investigated. We found that temporal networks will often have a high degree of dynamic recoverability, with surprisingly good recovery performance for multiple networks, underscoring the importance of knowing the network structure: a good picture of network structure makes it easier to determine when recovery will be accurate and when it will be inaccurate. In particular, the density of the network structure challenges recovery, with denser networks making recovery more error-prone. Conversely, dynamical sparsity, where less network activity tends to occur, can make recovery more accurate. These competing forms of density and sparsity can be exploited, and we found that sparsity-promoting methods applied to the recovery problem find good solutions, sometimes surprisingly so, for both recovery itself and for informing subsequent network analysis tasks that depend on the recovered information.

Our results here point towards a number of questions that warrant further study. The recovery problem is an inference problem, not prediction, but it would be interesting to see how well past values of  $\mathbf{N}$  can predict future values of  $\mathbf{E}$ . As recovery errors already occur without prediction, we expect compounding errors when trying to predict  $\mathbf{E}$  unless the system under study happens to exhibit stationarity. Likewise, what are the effects of measurement error? In other words, how much more difficult is recovery when  $\mathbf{B}$  is not known perfectly? This question should be addressed and is an important direction for future work. Errors-in-variables methods are the natural starting point. Although this problem may be especially challenging as sparse estimators are known to be unstable under uncertainty, the use of specialized estimators may be fruitful [41]. And local recovery, where only subgraph dynamics are considered, are worth investigating, as local problems are often sufficient around network regions of interest and those subgraphs may be more tractable to analyze than the global network, both from a data collection perspective and from the perspective of the recovery problem.

This work underscores how important it is, across a variety of problem domains, to know the underlying network structure, as the incidence matrix feeds into the solution methods we consider. But this raises an important question: what are the scientific consequences and broader impacts, both positive and negative, that may arise as a result of gathering static network data and then recovering, with perhaps unexpected accuracy, temporality?

We discuss two scenarios. First, consider the example of an application (app) maker distributing their application on smartphones. Smartphones contain many sensors and collect much sensitive data [42, 14], such as location information [13], but also have privacy safeguards in place [43]. While a potentially rich source of social network

and human behavioral data, outside their own application, the application maker will not have access to sensitive information such as communication records. Now suppose that app maker requests access to user address books. These address books constitute local, egocentric snapshots of the users' static social network, and are especially informative after disambiguation and record linkage across datasets [44]. When combined with sensor data or other records, therefore, these static networks snapshots can be used to infer not just the intrinsic activities of users (the node activities) but potentially even the activities of specific social ties (the edge activities). While inferences along these lines will certainly require significant data, and are especially interesting in certain use cases, such as inferring the illicit communications of criminals, this scenario underscores the privacy concerns here, as it means an entity such as a social media platform provider can begin to infer more information than expected about individuals by distributing applications and cross-referencing multiple datasets.

For the second scenario, we consider the implications for network neuroscience [2]. Imaging studies reveal spatiotemporal dynamics, generating time series that act as  $\mathbf{N}$  for an aggregated network where nodes represent brain regions-of-interest [45]. More useful would be to capture the edge dynamics governing when different ROIs “talk” to one another;  $\mathbf{E}$  is not accessible. However, our results here imply that, as network structure is revealed, for instance with tract- or fiber-tracing methods such as diffusion spectrum imaging [16, 46] or fluorescence microscopy [17], then  $\mathbf{B}$  can be inferred and combined with  $\mathbf{N}$  to approximate the missing  $\mathbf{E}$ . Accurate inference of these hidden dynamics can inform studies that move beyond the structural connectome and explore the brain's dynamical “dynome” [47].

## Acknowledgments

We are grateful to M. Almassalkhi and H. Ossareh for useful discussions. JPB acknowledges support by Google Open Source under the Open-Source Complex Ecosystems And Networks (OCEAN) project, U.S. Department of Energy's Advanced Research Projects Agency - Energy (ARPA-E) award DE-AR0000694, NASA under grant 80NSSC20M0213, and by the Broad Agency Announcement Program and Cold Regions Research and Engineering Laboratory (ERDC-CRREL) under Contract No. W913E521C0003. SL thanks the Villum Foundation (Nation Scale Social Networks) for support.

## References

- [1] P. Holme and J. Saramäki, “Temporal networks,” *Physics reports*, vol. 519, no. 3, pp. 97–125, 2012.
- [2] D. S. Bassett and O. Sporns, “Network neuroscience,” *Nature neuroscience*, vol. 20, no. 3, pp. 353–364, 2017.
- [3] A. Li, S. P. Cornelius, Y.-Y. Liu, L. Wang, and A.-L. Barabási, “The fundamental advantages of temporal networks,” *Science*, vol. 358, no. 6366, pp. 1042–1046, 2017.
- [4] P. Holme, “Modern temporal network theory: a colloquium,” *The European Physical Journal B*, vol. 88, no. 9, pp. 1–30, 2015.
- [5] N. Masuda and R. Lambiotte, *A guide to temporal networks*. World Scientific, 2016.

- [6] Y.-A. de Montjoye, S. Gambs, V. Blondel, G. Canright, N. de Cordes, S. Deletaille, K. Engø-Monsen, M. Garcia-Herranz, J. Kendall, C. Kerry, G. Krings, E. Letouzé, M. Luengo-Oroz, N. Oliver, L. Rocher, A. Rutherford, Z. Smoreda, J. Steele, E. Wetter, A. S. Pentland, and L. Bengtsson, “On the privacy-conscious use of mobile phone data,” *Scientific Data*, vol. 5, no. 1, p. 180286, 2018.
- [7] Z. Bar-Joseph, A. Gitter, and I. Simon, “Studying and modelling dynamic biological processes using time-series gene expression data,” *Nature Reviews Genetics*, vol. 13, no. 8, pp. 552–564, 2012.
- [8] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton, “Dynamic reconfiguration of human brain networks during learning,” *Proceedings of the National Academy of Sciences*, vol. 108, no. 18, pp. 7641–7646, 2011.
- [9] J. Faskowitz, F. Z. Esfahlani, Y. Jo, O. Sporns, and R. F. Betzel, “Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture,” *Nature neuroscience*, vol. 23, no. 12, pp. 1644–1654, 2020.
- [10] E.-Á. Horvát, M. Hanselmann, F. A. Hamprecht, and K. A. Zweig, “One plus one makes three (for social networks),” *PloS one*, vol. 7, no. 4, p. e34740, 2012.
- [11] D. Garcia, “Leaking privacy and shadow profiles in online social networks,” *Science advances*, vol. 3, no. 8, p. e1701172, 2017.
- [12] J. P. Bagrow, X. Liu, and L. Mitchell, “Information flow reveals prediction limits in online social activity,” *Nature human behaviour*, vol. 3, no. 2, pp. 122–128, 2019.
- [13] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, “Unique in the crowd: The privacy bounds of human mobility,” *Scientific reports*, vol. 3, no. 1, pp. 1–5, 2013.
- [14] J. L. Kröger, P. Raschke, and T. R. Bhuiyan, “Privacy implications of accelerometer data: a review of possible inferences,” in *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy*, pp. 81–87, 2019.
- [15] A. F. M. Altelaar, J. Munoz, and A. J. R. Heck, “Next-generation proteomics: towards an integrative view of proteome dynamics,” *Nature Reviews Genetics*, vol. 14, no. 1, pp. 35–48, 2013.
- [16] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff, “Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging,” *Magnetic resonance in medicine*, vol. 54, no. 6, pp. 1377–1386, 2005.
- [17] J. Livet, T. A. Weissman, H. Kang, R. W. Draft, J. Lu, R. A. Bennis, J. R. Sanes, and J. W. Lichtman, “Transgenic strategies for combinatorial expression of fluorescent proteins in the nervous system,” *Nature*, vol. 450, no. 7166, pp. 56–62, 2007.
- [18] G. Anthes, “Data brokers are watching you,” *Commun. ACM*, vol. 58, p. 28–30, Dec. 2014.
- [19] P. Sapiezynski, A. Stopczynski, D. D. Lassen, and S. Lehmann, “Interaction data from the Copenhagen Networks Study,” *Scientific Data*, vol. 6, no. 1, pp. 1–10, 2019.
- [20] R. Penrose, “On best approximate solutions of linear matrix equations,” *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 52, no. 1, pp. 17–19, 1956.
- [21] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [22] E. Candes and T. Tao, “The Dantzig selector: Statistical estimation when  $p$  is much larger than  $n$ ,” *Annals of statistics*, vol. 35, no. 6, pp. 2313–2351, 2007.

- [23] W. Xu, E. Mallada, and A. Tang, “Compressive sensing over graphs,” in *2011 Proceedings IEEE INFOCOM*, pp. 2087–2095, IEEE, 2011.
- [24] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical learning with sparsity: the Lasso and generalizations*. Chapman and Hall/CRC, 2015.
- [25] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [26] K. Lounici, M. Pontil, S. Van De Geer, and A. B. Tsybakov, “Oracle inequalities and optimal inference under group sparsity,” *Annals of statistics*, vol. 39, no. 4, pp. 2164–2204, 2011.
- [27] W. Lu, T. Dai, and S.-T. Xia, “Binary matrices for compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 66, no. 1, pp. 77–85, 2017.
- [28] M. Zhao, M. D. Kaba, R. Vidal, D. P. Robinson, and E. Mallada, “Sparse recovery over graph incidence matrices,” in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 364–371, IEEE, 2018.
- [29] A. Stopczynski, V. Sekara, P. Sapiezynski, A. Cuttone, M. M. Madsen, J. E. Larsen, and S. Lehmann, “Measuring large-scale social networks with high resolution,” *PLOS ONE*, vol. 9, no. 4, p. e95978, 2014.
- [30] P. Vanhems, A. Barrat, C. Cattuto, J.-F. Pinton, N. Khanafer, C. Régis, B.-a. Kim, B. Comte, and N. Voirin, “Estimating potential infection transmission routes in hospital wards using wearable proximity sensors,” *PLOS ONE*, vol. 8, no. 9, pp. 1–9, 2013.
- [31] B. Blonder and A. Dornhaus, “Time-ordered networks reveal limitations to information flow in ant colonies,” *PLOS ONE*, vol. 6, no. 5, pp. 1–8, 2011.
- [32] R. Michalski, S. Palus, and P. Kazienko, “Matching organizational structure and social network extracted from email communication,” in *Business Information Systems (W. Abramowicz, ed.)*, (Berlin, Heidelberg), pp. 197–206, Springer Berlin Heidelberg, 2011.
- [33] P. Panzarasa, T. Opsahl, and K. M. Carley, “Patterns and dynamics of users’ behavior and interaction: Network analysis of an online community,” *Journal of the American Society for Information Science and Technology*, vol. 60, no. 5, pp. 911–932, 2009.
- [34] M. Á. Serrano, M. Boguná, and A. Vespignani, “Extracting the multiscale backbone of complex weighted networks,” *Proceedings of the national academy of sciences*, vol. 106, no. 16, pp. 6483–6488, 2009.
- [35] S. Chen and D. Donoho, “Basis pursuit,” in *Proceedings of 1994 28th Asilomar Conference on Signals, Systems and Computers*, vol. 1, pp. 41–44, IEEE, 1994.
- [36] S. N. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu, “A unified framework for high-dimensional analysis of  $M$ -estimators with decomposable regularizers,” *Statistical science*, vol. 27, no. 4, pp. 538–557, 2012.
- [37] P. J. Bickel, Y. Ritov, and A. B. Tsybakov, “Simultaneous analysis of Lasso and Dantzig selector,” *The Annals of statistics*, vol. 37, no. 4, pp. 1705–1732, 2009.
- [38] M. Doob, “An interrelation between line graphs, eigenvalues, and matroids,” *Journal of Combinatorial Theory, Series B*, vol. 15, no. 1, pp. 40–50, 1973.
- [39] L. Lovász and M. D. Plummer, *Matching theory*, vol. 367. American Mathematical Soc., 2009.
- [40] M. E. J. Newman, “Mixing patterns in networks,” *Physical review E*, vol. 67, no. 2, p. 026126, 2003.
- [41] M. Rosenbaum and A. B. Tsybakov, “Sparse recovery under matrix uncertainty,” *The Annals of Statistics*, vol. 38, no. 5, pp. 2620 – 2651, 2010.

- [42] J. C. Sipiior, B. T. Ward, and L. Volonino, "Privacy concerns associated with smartphone use," *Journal of Internet Commerce*, vol. 13, no. 3-4, pp. 177–193, 2014.
- [43] I. Mohamed and D. Patel, "Android vs ios security: A comparative study," in *2015 12th International Conference on Information Technology - New Generations*, pp. 725–730, 2015.
- [44] K. Shu, S. Wang, J. Tang, R. Zafarani, and H. Liu, "User identity linkage across online social networks: A review," *SIGKDD Explor. Newsl.*, vol. 18, p. 5–17, Mar. 2017.
- [45] D. C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T. E. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S. W. Curtiss, *et al.*, "The human connectome project: a data acquisition perspective," *Neuroimage*, vol. 62, no. 4, pp. 2222–2231, 2012.
- [46] M. I. Menzel, E. T. Tan, K. Khare, J. I. Sperl, K. F. King, X. Tao, C. J. Hardy, and L. Marinelli, "Accelerated diffusion spectrum imaging in the human brain using compressed sensing," *Magnetic Resonance in Medicine*, vol. 66, no. 5, pp. 1226–1233, 2011.
- [47] N. J. Kopell, H. J. Gritton, M. A. Whittington, and M. A. Kramer, "Beyond the connectome: the dynome," *Neuron*, vol. 83, no. 6, pp. 1319–1328, 2014.