

One-Bit Total Variation Denoising over Networks with Applications to Partially Observed Epidemics

BY CLAIRE DONNAT

*Department of Statistics, University of Chicago,
5747 S Ellis Ave, Chicago, Illinois, 60637, United States of America
cdonnat@uchicago.edu*

OLGA KLOPP

*ESSEC Business School
3 Av. Bernard Hirsch, 95000 Cergy, France
CREST, ENSAE, Institut Polytechnique de Paris,
5 Av. Le Chatelier, 91120 Palaiseau, France
kloppolga@math.cnrs.fr*

AND NICOLAS VERZELEN

*MISTEA, Université Montpellier, INRAE, Institut Agro
2 Pl. Pierre Viala, 34000 Montpellier, France
nicolas.verzelen@inrae.fr*

SUMMARY

This paper introduces a novel approach for epidemic nowcasting and forecasting over networks using total variation (TV) denoising, a method inspired by classical signal processing techniques. Considering a network that models a population as a set of n nodes characterised by their infection statuses Y_i and that represents contacts as edges, we prove the consistency of graph-TV denoising for estimating the underlying infection probabilities $\{p_i\}_{i \in \{1, \dots, n\}}$ in the presence of Bernoulli noise. Our results provide an important extension of existing bounds derived in the Gaussian case to the study of binary variables — an approach hereafter referred to as one-bit total variation denoising. The methodology is further extended to handle incomplete observations, thereby expanding its relevance to various real-world situations where observations over the full graph may not be accessible. Focusing on the context of epidemics, we establish that one-bit total variation denoising enhances both nowcasting and forecasting accuracy in networks, as further evidenced by comprehensive numerical experiments and two real-world examples. The contributions of this paper lie in its theoretical developments, particularly in addressing the incomplete data case, thereby paving the way for more precise epidemic modelling and enhanced surveillance strategies in practical settings.

Some key words: High-dimensional statistics; Graph total-variation; Graph-Trend Filtering; Gamma sparsity; Epidemic Forecasting; Epidemic Nowcasting.

1. INTRODUCTION

We consider the problem of signal denoising over networks and apply our results to the analysis of infectious disease transmission over contact networks. The spread of epidemics over networks is an important area of research with implications for public health (Ottaviano et al., 2018; Sanchez et al., 2022), epidemiology (Danon et al., 2011; Keeling & Eames, 2005a; Moreno et al., 2002; Shirley & Rushton, 2005; Spricer & Britton, 2019), and network analysis (Kim et al., 2021; Wan et al., 2014). While the study of disease propagation through interconnected populations had already gained significant attention over the past decade (Keeling, 2005; Newman, 2002a; Pastor-Satorras et al., 2015a; Pellis et al., 2015), the recent emergence of new infectious diseases, such as COVID-19, further underscores the importance of developing comprehensive models that can capture the intricate interplay between epidemic transmission processes and the underlying network structures through which they spread (Keeling & Eames, 2005a; Manzo, 2020).

Studies of epidemic processes typically rely on compartmental models (Kermack & McKendrick, 1927; Tolles & Luong, 2020), which split a population of size n into different compartments (or groups) based on their disease status, and specify how agents transition between disease states (e.g., Susceptible, Infectious and Recovered). These approaches can be further refined into two categories, depending on how they account for randomness and uncertainty in the disease transmission process: *deterministic*, versus *stochastic* epidemic models. Deterministic models (Bowman et al., 2005; Carcione et al., 2020; Eikenberry et al., 2020; Kumar et al., 2021; Mubarak et al., 2021) assume that the disease transmission process can be precisely described using mathematical equations with fixed parameters. These models typically use differential equations to describe how the numbers of infected, susceptible, and recovered individuals change over time, and do not capture individual-level variability in the disease transmission process. Although deterministic models have provided valuable insights into the dynamics of epidemics, they often fail to capture the inherent randomness and unpredictability associated with real-world outbreaks. By contrast, the consideration of variability recognises that epidemic processes are influenced by random events, chance encounters, and unforeseen interactions among individuals. This acknowledgment prompts the need for stochastic models that can explicitly incorporate randomness into the modeling framework (Ball et al., 2019).

Instead of using differential equations, stochastic models often rely on techniques such as agent-based modelling (Hackl & Dubernet, 2019; Perez & Dragicevic, 2009; Siettos et al., 2015; Teweldemedhin et al., 2004) or stochastic differential equations on compartmental models (Andersson & Britton, 2012; Britton, 2010; Bu et al., 2022; Huang et al., 2022; Morsomme & Xu, 2022) to simulate individual interactions and disease transmission events. In this framework, each interaction or transmission event is subject to chance. The fluxes between compartments therefore follow a probability distribution, rather than a deterministic equation. Although these approaches are designed to better capture the uncertainty inherent in epidemic spreads, most of them only typically incorporate randomness in terms of transmission events: given a contact between an infected and susceptible individual, the transmission is subject to chance, and typically modelled as a Bernoulli variable. Specifically, these models do not attempt to remedy another major limitation of compartmental models: the uniform mixing assumption. This assumption implies that every infected individual in the population can potentially transmit the disease to any susceptible individual in the network, which is not always realistic (Keeling & Eames, 2005a). Such a simplification overlooks critical dynamics in epidemic spread on networks, particularly in large populations where the uniform mixing assumption is known to fail. These approximations can significantly affect the accuracy of models that attempt to capture how diseases spread through complex networks.

A possible solution to address this limitation is to model transmissions within a given contact network, thereby capturing more realistically possible transmission events. The structure of contact networks can in fact known to significantly impact disease spread (Ganesh et al., 2005). Indeed, while the standard epidemic models described in the previous paragraph correspond to fully connected networks, accounting for a more realistic structure of interactions and connections between individuals may lead to more accurate predictions and more effective control strategies (Chakrabarti et al., 2008; Wang et al., 2013). For example, Pastor-Satorras & Vespignani (2001) have shown that the structure of the contact network greatly influences the epidemic threshold and disease prevalence. Research on temporal networks (Holme & Saramäki, 2012) has shown that the timing of interactions plays a crucial role in shaping epidemic dynamics by affecting disease transmissibility. Epidemic models on graphs have also highlighted the importance of super-spreaders, that is, individuals or nodes with a disproportionately high potential for transmission. Finally, Keeling & Eames (2005b) introduced the notion of “assortativity” in networks, where highly connected nodes preferentially interact with each other. This assortative mixing has significant implications for the effective targeting interventions and controlling epidemics.

Another key limitation of traditional deterministic and stochastic differential equation models is the assumption that the current state of the epidemic is known: all agents’ infectious, susceptible, or recovered states are observed. Knowledge of the current state of the epidemic is, in fact, essential for making accurate predictions about its future course, as small inaccuracies can lead to substantial changes in epidemic trajectories over time (Siegenfeld et al., 2020). However, accurately assessing the current situation of an epidemic, a process often referred to as nowcasting, is notoriously challenging (Desai et al., 2019; McGough et al., 2020; Wu et al., 2021), due to the availability of quality data and the delays in reporting of current cases (Chakraborty et al., 2022; Rosenfeld & Tibshirani, 2021). In many situations, complete information about individuals’ health status cannot be disclosed due to privacy concerns. It thus becomes essential to develop methods that are adapted to the practical constraints of real data and respect individuals’ rights while still providing valuable insights into epidemic spread.

In the present paper, we propose a new approach to nowcasting and forecasting of epidemic spreads on a network inspired by total-variation signal denoising. Signal denoising based on the total variation (TV) penalty is a widely used technique in signal processing and image reconstruction that aims to remove noise from an observed signal while preserving important spatial or structural features. This approach uses the TV penalty, which measures the variation in the signal over a network. Letting $p_i, i \in \{1, \dots, n\}$ be a signal over a network $\mathcal{G}$ on n nodes and E edges, the total variation penalty can be written as: $\sum_{(i,j) \in E} |p_i - p_j|$. By penalising differences in the signal over neighbouring vertices, this term acts as a regularization penalty that enforces smoothness in the reconstructed signal. TV regularization is typically formulated as an optimisation problem, where the goal is to minimise the TV of the signal subject to a fidelity constraint in the form of a data-fitting loss, ensuring the reconstructed signal remains close to the original data. A more extensive review of Total-variation denoising is provided in Section 7.

In this paper, we study TV denoising in the setting of Bernoulli observations and show that this approach improves nowcasting estimates and prediction accuracy of dynamic behavior of epidemic processes over simple compartmental approaches. To address the challenges of epidemic nowcasting, we also extend our approach to the case of incomplete observations which allows a better assessment of the disease’s spread and can help implement more targeted surveillance and control measures.

Contributions

Our contributions are threefold and summarised as follows. First, we prove that the total variation denoiser that we propose – thereafter referred to as one-bit TV denoiser — is consistent in the case of Bernoulli noise (see Theorem 1). These results extend existing bounds that have been obtained in the literature in the case of Gaussian noise (Hütter & Rigollet, 2016). In particular, we obtain new bounds on the risk of TV denoiser measured in terms of the ℓ_1 -norm. Second, we extend our analysis to the challenging case of missing observations and demonstrate good performance of the one-bit TV denoiser (see Theorem 2). To the best of our knowledge, such results are new in the literature. Finally, we show how the one-bit total variation denoiser can be applied to improve the prediction of virus spread dynamics on networks (see Section 3). Our results are supported by extensive numerical studies that confirm the advantages of the proposed method, as well as two real-data experiments.

Notations

For any k , $\mathbf{1}_k$ denotes the k -dimensional vector of ones and $\mathbf{0}_k$ denotes the k -dimensional vector of zeros. For an integer $k \in \mathbb{N}$, we write $[k] = \{1, \dots, k\}$. We denote by $A^\dagger$ the Moore-Penrose pseudo-inverse of a matrix A and by $\|A\|$ its operator norm. Let $S \subset [n]$ be a subset of the nodes in $[n]$, we will use the notation $d_{\max}(S)$ to denote the maximum degree of nodes i in subset S .

2. SIGNAL DENOISING

Consider a network where each node $i \in V$ has an underlying state $y_i \in \{0, 1\}$ where 0 and 1 may respectively represent, for example, healthy and infected states. The probability that person i is in the infected state (or the proportion of the infected subpopulation at node i) is indicated by p_i^* , that is, $p_i^* = \text{pr}(y_i = 1)$. Let $p^* = (p_i^*)_{i=1, \dots, n}$ denote the vector that captures the infected states of an interconnected population consisting of n nodes. The state of the network can thus be represented by the following “signal + noise” model:

$$Y = p^* + \xi \quad (1)$$

where $Y = (y_i)_{i \in [n]}$ are our observations and ξ is a vector of centred Bernoulli noise. Given one vector of observations Y , we consider the problem of estimating the underlying $(p_i^*)_{i \in [n]}$.

Let $\mathcal{G} = (V, E)$ be the underlying undirected connected graph with vertex set V and edge set E . We set the cardinalities of both sets as $|V| = n$ and $|E| = m$. We denote by $A = (a_{ij})_{(i,j) \in [n] \times [n]}$ the graph’s adjacency matrix. It will be convenient for us to represent a graph by its edge-vertex incidence matrix $D \in \{-1, 0, 1\}^{m \times n}$. To each edge $e = (i, j) \in E$ corresponds a row D_e of D where the k th entry $D_{e,k}$ of D_e is given by

$$D_{e,k} = \begin{cases} 1 & \text{if } k = \min(i, j) \\ -1 & \text{if } k = \max(i, j) \\ 0 & \text{otherwise.} \end{cases}$$

Note that that $L = D^T D = \text{diag}(A\mathbf{1}_n) - A$ is the *unnormalized Laplacian* of the graph $\mathcal{G}$. Here $\text{diag}(A\mathbf{1}_n)$ is the diagonal matrix with j th diagonal element given by the degree of vertex j .

The TV denoiser $\hat{p}$ associated to G is any solution of the following minimization problem:

$$\hat{p} \in \arg \min_{p \in [0,1]^n} \left\{ \frac{1}{n} \sum_{i=1}^n (y_i - p_i)^2 + \lambda \|Dp\|_1 \right\}, \quad (2)$$

where $\lambda > 0$ is a regularization parameter and the TV penalty $\|Dp\|_1$ is the convex relaxation of $\|Dp\|_0$, which corresponds to the number of times p changes values along the edges of the graph G . Note that (2) is a convex problem that may be solved efficiently (Arnold & Tibshirani, 2016). We start by defining the *compatibility factor* and the *inverse scaling factor* that will be crucial in characterising the performance of the TV denoiser:

DEFINITION 1 (COMPATIBILITY FACTOR). Let $D \in \{-1, 0, 1\}^{m \times n}$ be the incidence matrix. The compatibility factor of a set $T \subset [m]$ is defined as

$$\kappa_T = \kappa_T(D) = \inf_{p \in [0, 1]^n} \frac{\sqrt{|T|} \|p\|_2}{\|(Dp)_T\|_1} \quad \text{for } T \neq \emptyset.$$

DEFINITION 2 (INVERSE SCALING FACTOR). Let $S = D^\dagger = [s_1, \dots, s_m]$. The inverse scaling factor of D is defined as

$$\rho = \rho(D) = \max_{j \in [m]} \|s_j\|_2.$$

The following result shown in Hütter & Rigollet (2016) provides a bound on both ρ and κ_T :

PROPOSITION 1. Let D be the incidence matrix of a connected graph G with maximal degree $d_{\max}$ and $\emptyset \neq T \subseteq E$. Let $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of the Laplacian L . Then,

$$\kappa_T \geq \frac{1}{2 \min\{\sqrt{d_{\max}}, \sqrt{|T|}\}} \quad \text{and} \quad \rho \leq \frac{\sqrt{2}}{\lambda_2}.$$

Equipped with these notations, we establish the following bound on the risk of TV denoiser (2):

THEOREM 1 (RISK BOUND FOR ONE-BIT TV DENOISING). Fix $\delta \in [0, 1]$, $T \subset [m]$ and assume that G is connected. Define the regularization parameter

$$\lambda = \frac{\sqrt{2}\rho}{n} \log \left(\frac{4n^2}{\delta} \right). \quad (3)$$

Then, the TV denoiser $\hat{p}$ defined in (2) satisfies

$$\|\hat{p} - p^*\|_2^2 \leq \frac{16\rho^2 |T| \log^2 \left(\frac{4n^2}{\delta} \right)}{\kappa_T^2} + 4\sqrt{2}\rho \|(Dp^*)_{T^c}\|_1 \log \left(\frac{4n^2}{\delta} \right) + \frac{4\|p^*\|_0 \log \left(\frac{4}{\delta} \right)}{n} \quad (4)$$

with probability at least $1 - \delta$.

The proof of Theorem 1 is provided in Appendix A. As we now show, inequality (4) allows trading off $|T|$ and $\|(Dp^*)_{T^c}\|_1$. Taking $T = \text{supp}(Dp^*)$, we obtain the following bound on the estimation error:

$$\|\hat{p} - p^*\|_2^2 \leq \frac{16\rho^2 \|Dp^*\|_0 \log^2 \left(\frac{4n^2}{\delta} \right)}{\kappa_T^2} + \frac{4\|p^*\|_0 \log \left(\frac{4}{\delta} \right)}{n}.$$

In particular, in the context of epidemic modelling, we can assume that in the first stages of the epidemic, the number of infected persons is relatively small and that p^* is sparse. Let $S = \text{supp}(p^*)$ and $s = |S| = \|p^*\|_0$. Then, using $|\text{supp}(Dp^*)| \leq s d_{\max}(S)$, the previous inequality

can be simplified to:

$$\|\widehat{p} - p^*\|_2^2 \leq \frac{16\rho^2 s d_{\max}(\mathcal{S}) \log^2\left(\frac{4n^2}{\delta}\right)}{\kappa_T^2} + \frac{4s \log\left(\frac{4}{\delta}\right)}{n}.$$

On the other hand, taking T to be the empty set in (4), we obtain the alternative error bound:

$$\|\widehat{p} - p^*\|_2^2 \leq 4\sqrt{2}\rho \|Dp^*\|_1 \log\left(\frac{4n^2}{\delta}\right) + \frac{4s \log\left(\frac{4}{\delta}\right)}{n}.$$

The next proposition provides a risk bound in ℓ_1 norm.

PROPOSITION 2. *Fix $\delta \in [0, 1]$, $T \in [m]$ and assume that G is connected. Define the regularization parameter*

$$\lambda = \frac{\sqrt{2}\rho}{n} \log\left(\frac{4n^2}{\delta}\right)$$

Then, the TV denoiser $\widehat{p}$ defined in (2) satisfies

$$\begin{aligned} \|\widehat{p} - p^*\|_1 &\leq \frac{4\rho\sqrt{s|T|}}{\kappa_T} \log\left(\frac{4n^2}{\delta}\right) + 2s\sqrt{\frac{\log\left(\frac{4}{\delta}\right)}{n}} \\ &\quad + 3\sqrt{\rho s \|(Dp^*)_{T^c}\|_1 \log\left(\frac{4n^2}{\delta}\right)} + \sqrt{2}\rho s d_{\max}(\mathcal{S}) \log\left(\frac{4n^2}{\delta}\right), \end{aligned} \quad (5)$$

with probability at least $1 - \delta$. Here, we recall that s stands for $\|p^\|_0$.*

The proof of Proposition 2 is provided in Appendix B. In Proposition 2, we can again trade off $|T|$ and $\|(Dp^*)_{T^c}\|_1$. Taking $T = \text{supp}(Dp^*)$ and using the fact that $|\text{supp}(Dp^*)| \leq s d_{\max}(\mathcal{S})$, the previous bound becomes:

$$\|\widehat{p} - p^*\|_1 \leq \rho s \left(\frac{4\sqrt{d_{\max}(\mathcal{S})}}{\kappa_T} + 2d_{\max}(\mathcal{S}) \right) \log\left(\frac{4n^2}{\delta}\right) + 2s\sqrt{\frac{\log\left(\frac{4}{\delta}\right)}{n}}.$$

On the other hand, taking T to be the empty set, we obtain:

$$\|\widehat{p} - p^*\|_1 \leq 3\sqrt{\rho s \|(Dp^*)\|_1 \log\left(\frac{4n^2}{\delta}\right)} + \sqrt{2}s \left(\rho d_{\max}(\mathcal{S}) \log\left(\frac{4n^2}{\delta}\right) + \sqrt{\frac{2 \log\left(\frac{4}{\delta}\right)}{n}} \right).$$

The errors terms given by (4) and (5) crucially depend on the inverse scaling factor ρ and the compatibility factor κ_T . Hütter & Rigollet (2016) provides upper bounds for these quantities for different graphs. Using $\lesssim$ for an inequality where the left-hand side is bounded by the right-hand side up to a numerical constant, we specialize the upper bound on the risk of the TV denoiser to specific graphs. All the following bounds (case 1 through 4) hold with probability at least $1 - 4/n$.

Case 1: 2D-grids. For the 2D grid, $\rho \lesssim \sqrt{\log n}$ (see Hütter & Rigollet (2016)). The ℓ_2 and ℓ_1 bounds on the risk of Proposition 2 and Theorem 1 thus yield that:

$$\|\widehat{p} - p^*\|_2 \lesssim \left\{ \sqrt{\|Dp^*\|_1 + \frac{s}{n \log^{1/2}(n)}} \right\} \log^{3/4}(n) \quad \text{and}$$

$$\|\widehat{p} - p^*\|_1 \lesssim \left\{ \sqrt{\frac{s(\|Dp^*\|_1)}{\log^{3/2}(n)}} + s \right\} \log^{3/2}(n).$$

Case 2: Complete graphs. For a complete graph, $\rho \lesssim 1/n$ and $\kappa_T \gtrsim 1/\sqrt{n}$ (Hütter & Rigollet, 2016). This implies:

$$\|\widehat{p} - p^*\|_2 \lesssim \sqrt{\frac{\|Dp^*\|_1 + s}{n}} \log^{1/2}(n) \quad \text{and}$$

$$\|\widehat{p} - p^*\|_1 \lesssim \left(\sqrt{\frac{s\|Dp^*\|_1}{n \log(n)}} + s \right) \log(n).$$

Case 3: Star graphs. For a star graph, $\rho \lesssim 1$ and $\kappa_T \gtrsim 1/\sqrt{|T|}$, so that:

$$\|\widehat{p} - p^*\|_2 \lesssim \left\{ \sqrt{\|Dp^*\|_1 + \frac{s}{n}} \right\} \log^{1/2}(n) \quad \text{and}$$

$$\|\widehat{p} - p^*\|_1 \lesssim \left(\sqrt{\frac{s\|Dp^*\|_1}{\log(n)}} + s d_{\max}(S) \right) \log(n)$$

Case 4: Graphs with spectral gaps. If the graph Laplacian has a spectral gap, i.e., there exists $\nu > 0$ such that the second smallest eigenvalue of the Laplacian is larger than ν , then $\rho \leq 2/\nu$ (see Proposition 14 in Hütter & Rigollet (2016)). By Proposition 1, we also know that $\frac{1}{\kappa_T^2} \leq 2\min(|T|, d_{\max})$. In particular, for a random d_n -regular graph with $d_n = d_0(\log n)^\beta$ for some $\beta > 0$ and $d_0 > 1$ and for Erdős-Rényi graphs $\mathcal{G}(n, p)$ with $p = d_n/n$, we obtain:

$$\|\widehat{p} - p^*\|_2 \lesssim \sqrt{\frac{d_n\|Dp^*\|_0}{\nu^2} + \frac{s}{n \log^{1/2}(n)}} \log(n) \quad \text{and}$$

$$\|\widehat{p} - p^*\|_1 \lesssim \left(\sqrt{\frac{s d_n\|Dp^*\|_0}{\nu^2}} + \frac{s d_n}{\nu} \right) \log(n)$$

Discussion. These bounds allow to shed more light on the interplay between the size of the support $\text{supp}(Dp^*)$, the size of the epidemic $\|p^*\|_0$, as well as the topology of the graph. At a high level — and as confirmed by our experiments in Section 5 —, we expect our total variation denoiser to be better suited to situations where the cardinality of the signal's support is small compared to the size of the graph ($\|p^*\|_0 = s < n$), and the maximal degree of the graph remains of constant order, so that $|\text{supp}(Dp^*)| \leq s d_{\max}(V)$ is also controlled. This means for instance that the denoiser is likely to provide a better control of the ℓ_1 -error in the 2D-grid (where the ℓ_1 error is of the order of $s \log^{3/2}(n)$) compared to star graph, where the error could be in the order of $sn \log(n)$.

3. VIRUS SPREAD DYNAMICS ON NETWORK

We now apply our one-bit total variation denoiser to the analysis of epidemic spreads over networks. As highlighted in the introduction, numerous compartmental models have been suggested to depict the spread of epidemic processes over networks. We refer the reader to the recent surveys by Draief & Massoulié (2009); Pastor-Satorras et al. (2015b) or Paré et al. (2020a) for a more extensive review of these methods. We first propose deploying our model to the simple Susceptible - Infectious - Susceptible (SIS), before showing how our approach can be used in a wider variety of epidemic models.

SIS Model

We start by considering the networked version of the susceptible - infectious - susceptible (SIS) compartmental model, originally introduced in Kermack & McKendrick (1927). This widely used model has been indeed adapted for network settings in various subsequent studies (Ahn & Hassibi, 2013; Paré et al., 2020b), as detailed in the subsequent paragraph. We take the vector $(p_i)_{i \in [n]}$ to represent the probabilities of each node being infected, with each entry corresponding to an individual node's infection probability. In the networked setting (Paré et al., 2020a), the evolution of the p_i s is governed by the following differential equation:

$$\dot{p}_i(t) = (1 - p_i(t))\beta_i \sum_{j=1}^n \omega_{ij} p_j(t) - \gamma_i p_i(t). \quad (6)$$

Here ω_{ij} quantifies the strength of the connection from node j to node i , $\beta_i > 0$ is the infection (or susceptibility) rate and γ_i is the healing rate. The *infection rate* β_i quantifies the rate at which a susceptible individual (node i) in the population is infected through contacts with infected individuals (denoted here by the sum $\sum_j \omega_{ij} p_j$). In many cases, the susceptibility to a virus may not be uniform across the population. Susceptibility to a virus often varies across a population due to factors such as age, pre-existing health conditions, or immunity from previous exposure or vaccination. To account for this heterogeneity in infection rates, one solution consists in identifying relevant subgroups within the population and determining the infection rate for each subgroup.

The *healing rate* γ_i models the rate at which an infected individual i transitions between infected and recovered states. When γ_i is assumed to be identical across all individuals ($\gamma_i = \gamma$, $\forall i = 1, \dots, n$), it can be understood as the proportion of infected individuals who recover from the virus in a given time window Δt . In this case, the healing rate can be estimated by tracking the number of individuals who transition from infected to recovered status over time, and calculated as the number of recoveries divided by the number of confirmed cases during that time window Δt . Usually, the time window for calculating the recovery rate is determined by the average duration of infection. In general, to estimate the healing rate, we need information on the typical progression of the disease, including the duration of the infection, the severity of symptoms, and the factors that contribute to recovery. Note that the healing rate can vary depending on factors such as age, underlying health conditions, and the availability and effectiveness of treatments.

The SIS model introduced in (6) can be derived from a subpopulation perspective or through a mean-field approximation of a 2^n state Markov chain model (Paré et al., 2020b). As the epidemics dynamic is usually sampled at discrete time steps, we will consider the discrete version of model (6):

$$p_i^{k+1} = p_i^k + (1 - p_i^k)\beta_i \sum_{j=1}^n \omega_{ij} p_j^k - \gamma_i p_i^k \quad (7)$$

where $k > 0$ is the time index. In matrix form, the model can be re-written as:

$$p^{k+1} = p^k + [(1 - P^k)B\Omega - \Gamma] p^k \quad (8)$$

where $P^k = \text{diag } p^k$, $B = \text{diag } \beta_i$, $\Omega = (\omega_{ij})$ and $\Gamma = \text{diag } \gamma_i$. We denote by

$$O_k = 1 + (1 - P^k)B\Omega - \Gamma.$$

For this model to be well defined, we need the following assumption:

Assumption 1. For all $i \in [n]$, we have $\gamma_i < 1$ and $\sum_{j=1}^n \beta_i \omega_{ij} < 1$.

In this case, we have the following result:

LEMMA 1 ((PARÉ ET AL., 2020B)). *Consider the model (8) under Assumption 1. Suppose $p_i^0 \in [0, 1]$ for all $i \in [n]$. Then, for all $k > 0$ and $i \in [n]$, $p_i^k \in [0, 1]$.*

Using data on healthy / infected individuals at time point k^0 , we can estimate p^{k^0} using the TV estimator given by (2). Let

$$\widehat{O}_k = 1 + (1 - \widehat{P}^k)B\widehat{\Omega} - \Gamma \quad (9)$$

be the evolution operator, where $\widehat{P}^k = \text{diag}(\widehat{p}^k)$ and $\widehat{\Omega}$ is an estimator of Ω from noisy observations of the network (e.g. using variational maximum likelihood from some (partial) observations of the network (Gaucher & Klopp, 2021), as discussed in Remark 3 below). Then, by recursively computing $\widehat{p}^k = \widehat{O}_{k-1}\widehat{p}^{k-1}$ for $k > k^0$, we can predict the evolution of the epidemic. In particular, the expected number of infected individuals at time k , $E(I^k) = \sum p_i^k = \|\widehat{p}^k\|_{l_1}$ can be estimated using Proposition 2.

PROPOSITION 3. *Consider the discrete time SIS given by (7). Assume that we observe $Y^{k^0} \in \{0, 1\}^n$ which encodes the data on the healthy / infected status of a population of n connected individuals at time point k^0 . Fix $\delta \in [0, 1]$, $T \subset [m]$ and assume that the contact graph G is connected. Then, with probability at least $1 - \delta$, we have*

$$\begin{aligned} \left| \|\widehat{p}^k\|_{l_1} - E(I^k) \right| &\leq (1 - \gamma_{\min} + \beta_{\max} \|\Omega\|_{1,\infty})^{k-k^0} \left(\frac{4\rho\sqrt{s|T|}}{\kappa_T} \log \left(\frac{4n^2}{\delta} \right) + \right. \\ &\quad \left. + 2s \sqrt{\frac{\log \left(\frac{4}{\delta} \right)}{n}} + 3 \sqrt{\rho s \|(Dp^{k^0})_{T^c}\|_1 \log \left(\frac{4n^2}{\delta} \right)} + \sqrt{2\rho s d_{\max}(S) \log \left(\frac{4n^2}{\delta} \right)} \right). \end{aligned}$$

where $\gamma_{\min} = \min_i \gamma_i$ and $\beta_{\max} = \max_i \beta_i$

Proof. We have $E(I^k) = \|p^k\|_{l_1}$ and

$$\left| \|\widehat{p}^k\|_{l_1} - \|p^k\|_{l_1} \right| \leq \|\widehat{p}^k - p^k\|_{l_1} \leq (1 - \gamma_{\min} + \beta_{\max} \|\Omega\|_{1,\infty})^{k-k^0} \|\widehat{p}^{k^0} - p^{k^0}\|_{l_1}.$$

The result follows by bounding $\|\widehat{p}^{k^0} - p^{k^0}\|_{l_1}$ using the results of Proposition 2. $\square$

Note that Assumption 1 implies $\beta\|\Omega\|_{1,\infty} \leq 1$.

Remark 1. Beyond SIS model, one can also consider Susceptible - Infected - Recovered (SIR) and Susceptible - Exposed - Infected - Recovered (SEIR) models (Paré et al., 2020a). For a SIR process, for example, let r_i^k denotes the probability of node i being recovered at time k . Then, the

following equations govern the evolution of p^k and r^k :

$$p_i^{k+1} = p_i^k + (1 - p_i^k - r_i^k)\beta_i \sum_{j=1}^n \omega_{ij} p_j^k - \gamma_i p_i^k$$

$$r_i^{k+1} = r_i^k - \gamma_i p_i^k.$$

Our denoiser can therefore easily be deployed in the context of SIR and SEIR models by simply substituting the evolution operator of Equation 9 with a appropriate version.

Remark 2 (Taking into account false positives). In the case of epidemic models on networks, it is also necessary to account for errors that may occur in our observations Y^{k_0} due to false positives. Such false positives might be due to errors due to imperfect test specificity for instance: while gold standard medical tests can have specificity up to 99%, reporting based on symptom data is usually less accurate, and can thus lead to false positives. Here we assume that we know the false positive rate α , and that our observations y_i^k follow Bernoulli distribution with probability $\rho_i^k = (1 - \alpha)p_i^k + \alpha$. In this case, to estimate p^k we add a post-processing step thresholding to zero all the coordinates of $\hat{p}^k$ obtained from (2) that are smaller or equal to α . We test the performance of this method on graphs with various topologies, using the experimental setup described in Section 5. As shown in Figure 1 for the case of the k-NN graph, our procedure can significantly improve the accuracy of the nowcasted estimates $\hat{p}$ over the naive estimates $\hat{p}^{\text{naive}} = y_{\text{observed}}$, particularly as the epidemic size increases. Plots for other graph topologies (small-world, preferential attachment, etc) can be found in Appendix D.




Fig. 1: ℓ_1 error for the k-NN graph, as a function of the noise level (false positive rate) $\alpha \in \{0, 1\%, 5\%\}$, k_0 and β . The healing rate is here fixed at $\gamma = 0.1$. The benchmark values (in red) have been slightly shifted on the x-axis to improve legibility. The results are averaged over 100 independent experiments. Error bars indicate interquartile ranges.

Remark 3 (Estimating the strength of the connections in the network). One of the key quantities that determines the evolution of the epidemic on the network in (7) is the matrix $\Omega = (\omega_{ij})_{(ij) \in [n] \times [n]}$, which quantifies the strength of the connections between nodes in the network. We can estimate Ω from the observation of a network using the inhomogeneous random graph model assuming that for $i < j$, the elements A_{ij} of the adjacency matrix A are independent Bernoulli random variables with the success probability $\omega_{ij} \in [0, 1]$. In this case, each entry ω_{ij} can be interpreted as the probability that the edge (i, j) is present in the graph $\mathcal{G}$.

The problem of estimating Ω from a single observation of the network has been considered in several papers (Gao et al., 2015; Klopp et al., 2017; Klopp & Verzelen, 2019; Gaucher & Klopp, 2021) where strong theoretical guarantees have been obtained for this problem. One of the most popular methods is based on the variational approximation of the likelihood function (Celisse et al., 2012; Bickel et al., 2013). Loosely speaking, this method attaches to the latent variables a distribution with free parameters. These parameters are fitted in a way to obtain a distribution close in Kullback-Leibler divergence to the true posterior. The crux of this approach lies in the realisation that the variational distribution is simpler than the true posterior, which allows to solve the optimization problem approximately, see, for example, Gaucher & Klopp (2021).

Remark 4 (Learning parameters γ and β). Assume that we know $p_k, p_{k+1}, \dots, p_K$. We can write the following equation:

$$\begin{bmatrix} p^{k+1} - p^k \\ p^{k+2} - p^{k+1} \\ \vdots \\ p^{K-1} - p^K \end{bmatrix} = \Phi \begin{bmatrix} \beta \\ \gamma \end{bmatrix}$$

where

$$\Phi = \begin{bmatrix} (I - P^k)\Omega p^k & -p^k \\ \vdots & \vdots \\ (I - P^{K-1})\Omega p^{K-1} & -p^{K-1} \end{bmatrix}$$

Vrabac et al. (2020) propose to estimate the vector of parameters $[\beta, \gamma]^T$ using the inverse (or pseudoinverse) of matrix Φ . However, in practice, we do not have access to p^j but only to its noisy counterpart Y_i , as given by (1). In such situations, one can replace p^j by its estimator $\hat{p}^j$ given by (2) or (11) and compute the inverse (or pseudoinverse) of matrix Φ using $\hat{p}^j$. We illustrate the efficiency of this approach in numerical studies in Section 5.3.

4. PARTIALLY OBSERVED EPIDEMICS

In (2), we assume that we observe the state $y_i \in \{0, 1\}$ for any node in the network. However, in many real-world situations, it is challenging to obtain complete information about each individual's disease status. People may not be aware of their infection or may choose not to disclose it. Therefore, methods that can work with partial information become indispensable. In this section, we extend our analysis to the case of partial observations of Y^k .

Let $\mathfrak{M} = (m_i)$ be a mask vector. That is, $m_i = 1$ if we observe the information about node i , and $m_i = 0$ if not. We assume that m_i are independent Bernoulli random variables with parameter π_i : $m_i \sim \mathcal{B}(\pi_i)$ and $\pi_i > 0$. In the simplest setting, each coefficient is observed with the same probability, i.e., for every $i \in [n]$, $\pi_i = \nu$. In the case of partial observations, we observe two vectors, the mask vector $\mathfrak{M}$ and the vector $\tilde{Y} = \mathfrak{M} \circ Y$ where $\circ$ denotes the entry-wise (Hadamard)

product. For any $v = (v_i)_{i \in [n]} \in \mathbb{R}^n$ we define the weighted l_2 -norm of v as:

$$\|v\|_{l_2(\pi)}^2 = \sum_i \pi_i v_i^2.$$

Let

$$\kappa_\pi = \left(\min_i \pi_i \right)^{-1}$$

and $\pi^{-1} = (\pi_i^{-1})_{i \in [n]}$. It is easy to see that for any $v \in [0, 1]^n$

$$\|v\|_2 \leq \kappa_\pi \|v\|_{l_2(\pi)}. \quad (10)$$

We use the partial observations $(\mathfrak{M}, \tilde{Y})$ to construct the TV denoiser $\hat{p}_{miss}$ associated to G as any solution of the following minimization problem:

$$\hat{p}_{miss} \in \arg \min_{p \in [0, 1]^n} \left\{ \frac{1}{n} \sum_{i=1}^n m_i (\tilde{y}_i - p_i)^2 + \lambda \|Dp\|_1 \right\}. \quad (11)$$

We prove the following bound on the risk of the TV denoiser estimated on partially observed data:

THEOREM 2 (RISK BOUND FOR TV DENOISING WITH PARTIAL OBSERVATIONS). *Assume that G is connected and define the regularization parameter*

$$\lambda = \frac{9\sqrt{2}\rho \log(n)}{n}. \quad (12)$$

For $T = \text{supp}(Dp^)$, the TV denoiser from partial observations, $\hat{p}_{miss}$, defined in (11) satisfies*

$$\|\hat{p}_{miss} - p^*\|_{l_2(\pi)}^2 \leq C \left\{ \frac{\rho^2 \kappa_\pi |T|}{\kappa_T^2} + \frac{\|\pi\|_1 \|\pi^{-1}\|_1}{n^2} \right\} \log^2(n) \quad (13)$$

with probability at least $1 - c/n$ where C, c are fixed numerical constants.

The proof of Theorem 2 is provided in Appendix C. These results therefore allow us to provide theoretical guarantees on the recovery of p^* despite missing observations. In particular, we consider the setting where $\pi_i = v$, in which case, v captures the fraction of observed nodes within the network. In this case, the dependency between the number of missing observations and the quality of the bound becomes clearer: as expected, the smaller the fraction v , the larger the error.

5. NUMERICAL EXPERIMENTS

In this section, we validate our method through a series of synthetic and semi-synthetic experiments. These controlled experiments allow us to test a variety of data regimes by varying the topology of the contact graph - a dependency highlighted by the presence of the scaling factor ρ in Proposition 3 - as well as varying the parameters of the epidemic model (β and γ).

To simulate an epidemic process, we first sample a graph on n nodes, where n is fixed to 1,000 unless otherwise specified. In the results presented here, we consider a variety of random graph models, including Erdős-Rényi graphs (where each edge is independently sampled with probability p), Stochastic Block Models, Small World Networks (parametrised by their average degree m and a rewiring parameter p), Power-Law graphs (parametrised by a parameter, m , indicating the number of edges to attach from a new node to existing nodes as the graph

is generated), as well as k -Nearest Neighbour graphs (generated by uniformly sampling 2D coordinates between 0 and 1, and connecting nodes to their k -nearest neighbours in Euclidean norm). To further improve the realism of our experiments and consider graphs reflective of real-world contact network characteristics, we also generate an epidemic process on the Berkeley graph (Rossi & Ahmed, 2015). The latter is a large social friendship network extracted from Facebook consisting of users (nodes) with edges representing friendship ties. This also allows us to test the performance of our algorithm on a larger-size graph, as this social network has indeed a total of 22,900 nodes and 852,419 edges.

For each experiment, we randomly sample a “patient 0”, and propagate the epidemic on the graph $\mathcal{G}$ as per Equation (8) for k_0 steps, using a specified value of the infection rate β and healing rate γ . The probability of infection for each node here is chosen to be identical: $\beta_i = \beta$. We select the interaction strength between each node i, j as $w_{ij} = \frac{1}{d_i \vee d_j}$. This models a situation in which high degree nodes (e.g. influencers) have looser connections to their neighbours, while peripheral nodes with fewer connections might have tighter bounds to their neighbours. The rationale for this choice of prior is reinforced by the observation that the rate of epidemic transmission increases over time. Individuals with a high number of connections, denoted by high degrees, are here presumed to have, in average, brief interactions with each contact. In contrast, individuals with fewer connections tend to represent closer-knit groups, such as family units, where extended interactions are more common. Note that Assumption 1 imposes additional conditions on the topology of the graph. For instance, in Erdős-Rényi (ER) graphs, Assumption 1 implies that $np_{ER}\beta \leq 1$ so that $p_{ER} = O(\frac{1}{n})$ — in other words, the ER graph cannot be too dense. This is by no means a limitation of our model, but rather, a way of establishing graph topologies in which our sparsity assumption can be verified.

In our experiments, unless specified otherwise, γ is fixed to 0.1. This corresponds to an average recovery time of 9 days, assuming that the recovery is distributed as a negative binomial with a probability γ of recovery at each time step. The result of the simulation is an underlying probability $(p^*)_{i \in [n]}$ that each node is infected after k_0 epidemic diffusion steps. In this setting, the lower the value of k_0 , the more localized the epidemic. Finally, we generate a vector of observations, sampling each entry (i.e. the infection status of each node) as $y_i \sim \text{Bernoulli}(p_i^*)$. An example of this diffusion process is presented in Figure 2.




Fig. 2: Visualisation of a SIS epidemic process on a 5-nearest neighbour graph at time 0 (left: one initial infected node), and after $k_0 = 10$ (middle) and $k_0 = 20$ (right) diffusion steps, for $\beta = 0.9$.

5.1. Algorithm.

To efficiently solve this problem, we use the semi smooth-Newton augmented Lagrangian method (SSNAL) of Sun et al (Sun et al., 2021). SSNAL is a scalable and efficient algorithm that was originally designed to solve the following optimization problem:

$$\min_{X \in \mathbb{R}^{d \times n}} \frac{1}{2} \sum_{i=1}^n \|x_i - a_i\|^2 + \lambda \sum_{i < j} W_{ij} \|x_i - x_j\|_2 \quad (14)$$

where $W_{ij} \in [0, 1]$ is a matrix of weights such that $W_{ij} = 0$ if nodes i and j are disconnected in the graph, and reflects the proximity between node i and j otherwise. This formulation appears in several settings, particularly in convex clustering problems, where it is used to define a hierarchical clustering of the data. Note that in our case, since our vector $(p_i)_{i \in [n]}$ is of dimension 1, so $\|p_i - p_j\|_2 = \sqrt{(p_i - p_j)^2} = |p_i - p_j|$, and our problem shares the same objective function as Equation (14). In SSNAL, the optimisation problem (14) is solved using an augmented Lagrangian method on equation (14), and solving the corresponding sub-problems using semi-smooth Newton conjugate gradient updates. In their work, Sun et al. (2021) show that as long as the dimension of the feature vector is small (i.e., in our case, equal to 1), this method has the potential to considerably improve existing solvers. We provide a comparison of this solver against alternative solvers (e.g. ADMM (Boyd et al., 2011) and CVX (Diamond & Boyd, 2016)) in Appendix D. In our experiments, an appropriate value for the parameter λ is chosen by cross-validation¹.

5.2. Results.

Figures 3 and 4 show the results for the k-NN graph, as the values of k , k_0 and β vary. In these plots, the ℓ_1 and ℓ_2 errors (Figure 3 and 4, respectively) of our estimator are presented as a function of the probability of transmission β . Each column represents a number of diffusion of steps (k_0): as emphasized earlier, the higher the value, the more advanced — and therefore difficult to denoise, from an inference perspective — the epidemic spread. Each row corresponds to a different graph topology: in this case, the higher the value of k , the denser the contact network. Table 1 presents the results for the Berkeley graph, and plots corresponding to other topologies are displayed in Figures 11 and 12, Appendix D. Overall, we observe that for graphs where the epidemic has not propagated, the denoising method exhibits only marginal improvement over the naive estimator $\hat{p}^{\text{naive}} = y_{\text{observed}}$: in the k-NN and Power-Law graphs for instance, when the number of steps is small, the average number of infections is low (less than 10— see Figure 10), and the one-bit TV denoising method only improve the raw observations by 7% ($\|\hat{p} - p^*\|_1 = 2.03$ for the 2-NN denoised estimate vs 2.17 for the raw observations at $k_0 = 10$). However, as the size of the epidemic increases, the effect of the denoising estimator grows considerably. For the 5-NN graph, with $\beta = 0.5$ and $k_0 = 30$, our denoising leads to an almost 50% decrease of the ℓ_1 error for the estimated current state of the epidemic.

Figures 5 shows the ℓ_2 error in the predicted state of the epidemic in 2 days. These estimates are obtained by entering the estimated value of $\hat{p}$ in Equation (7), and rolling out the updates for two time steps. This yields potentially substantial differences in the predictions of the state of the epidemic, as depicted in Figure 5.

5.3. Estimating parameters γ and β .

To illustrate the efficiency of 1-bit denoiser in the case of estimation of parameters γ and β we use a kNN-nearest neighbour graph and propagate the epidemic for 20 steps. We vary the

¹ The code for all of our simulations and analyses is available at the following link: <https://github.com/donnate/epidemics>.




Fig. 3: Effect of the denoising (with λ chosen through cross validation), as a function of epidemic time (columns) and graph topology (row), measured by the difference in ℓ_1 -norm between the ground-truth. Points denote the mean $\ell_1 = \|p^* - \hat{p}\|_1$ error over 200 simulations for $k = 2$, and 100 simulations for the other topologies. Error bars indicate interquartile ranges. A plot showing the size of the corresponding epidemic can be found in figure 10, Appendix D.

values of β and γ , using $k_0 = 30$. We compare the values of the estimated values of β and γ as reported by our estimator with those using the same method, but fitting instead the observed values of the state of the epidemic (the naive estimator). The results for the estimates of β , γ and the reproductive number $R_0 = \frac{\beta}{\gamma}$ - an important characteristic of the disease -, averaged over 100 simulations, are reported in Table 2. In most cases, our TV denoiser improves the parameter estimates, achieving almost perfect reconstruction for $\beta = 0.8$ and $m = 5$. In Figure 6, we evaluate the effect of the number of observed time steps k_0 , the graph parameter k , and the transmission probability β on the estimation error. Overall, we observe that, for this method to be accurate, the epidemic must have a size that is neither too big nor too small: if the epidemic is too big (e.g. as it tends to be the case when $\beta = 0.9$), the sparsity assumption ceases to hold, while if the epidemic is too small (e.g. as it tends to be the case when $\beta = 0.2$), there are too few observations to reliably estimate the parameters.




Fig. 4: Effect of the denoising (with λ chosen through cross validation), as a function of epidemic time (columns) and graph topology (row), measured by the difference in square ℓ_2 -norm between the ground truth and estimated infection parameter p . Points denote the mean $\ell_2 = \|p^* - \hat{p}\|_2^2$ error over 200 simulations for $k = 2$, and 100 simulations for the other topologies.

5.4. Partially observed epidemic

We also study the performance of the 1-bit TV denoiser in the case of partial observations, that is the goodness-of-fit term of our estimator is only applied to nodes whose infectious statuses are observed. We vary the proportion of missing values, and plot the corresponding estimation error. The results are presented in Figure 7. We observe that our estimator can provide substantial improvement over the naive estimator (where all masked entries are filled in with 0) for even large proportions of missing values. This improvement is particularly substantial in cases where the epidemic has progressed enough so that the trivial estimation $p_1^* = 0$ is no longer accurate, but not enough to violate our sparsity assumption on the size of the support $\|p^*\|_0$.

6. REAL DATA ANALYSIS

In this section, we propose to analyse two real epidemics datasets using our method.

Steps	β	$\ \hat{p} - p^*\ _1$	$\ \hat{p}^{\text{naive}} - p^*\ _1$
5	0.7	0.00014	0.00014
5	0.9	0.00024	0.00026
10	0.3	0.00011	0.00012
10	0.5	0.00026	0.00032
10	0.7	0.00094	0.00124
20	0.3	0.00033	0.00042

Table 1: Accuracy of the estimation (in ℓ_1 norm) for the Berkeley Graph. Note that in both cases, the accuracy is high, since the epidemic in this case remains very localised.




Fig. 5: ℓ_2 error on the state of the epidemic in 2 days, using the predicted $\hat{p}^{(2)}$. Results are compared to the naive estimator that solely uses the observed values: $\hat{p}^{\text{naive}} = y_{\text{observed}}$ to nowcast the epidemic and is then propagated twice as per Equation (9). The benchmark values (in red) have been slightly shifted on the x-axis to improve legibility. The results are averaged over 100 independent experiments. Error bars indicate interquartile ranges.

k-NN	β	γ	R_0	Method	$\hat{\beta}$	$\hat{\gamma}$	$\hat{R}_0$
5	0.35	0.1	3.5	TV denoiser	0.78 (0.58, 1)	0.32 (0.23, 0.44)	2.4 (2.27, 2.55)
				Naive	0.99 (0.72, 1)	0.39 (0.27, 0.52)	2.47 (1.91, 2.7)
	0.5	0.1	5	TV denoiser	0.72 (0.58, 0.84)	0.17 (0.13, 0.22)	4.1 (3.89, 4.38)
				Naive	0.85 (0.69, 1)	0.21 (0.15, 0.27)	4.11 (3.76, 4.39)
	0.8	0.1	8	TV denoiser	0.84 (0.72, 0.98)	0.11 (0.09, 0.14)	7.53 (7.07, 8.2)
				Naive	0.94 (0.8, 1)	0.12 (0.09, 0.14)	7.38 (6.21, 8.02)
10	0.35	0.1	3.5	TV denoiser	0.62 (0.48, 0.78)	0.23 (0.17, 0.31)	2.64 (2.51, 2.81)
				Naive	0.78 (0.6, 0.97)	0.29 (0.21, 0.38)	2.67 (2.53, 2.81)
	0.5	0.1	5	TV denoiser	0.66 (0.55, 0.76)	0.15 (0.12, 0.19)	4.29 (4.09, 4.6)
				Naive	0.77 (0.64, 0.9)	0.18 (0.14, 0.22)	4.26 (4.08, 4.53)
	0.8	0.1	8	TV denoiser	0.83 (0.74, 0.93)	0.11 (0.09, 0.13)	7.75 (7.38, 8.11)
				Naive	0.91 (0.82, 1)	0.12 (0.1, 0.14)	7.62 (7.11, 7.97)

Table 2: Results for the estimation of the parameters γ and β , as well as the reproductive number $R_0 = \frac{\beta}{\gamma}$. Intervals denotes interquartile ranges. Results averaged over 500 simulations, for $K_0 = 30$.

6.1. Analysis of the COVID-19 outbreak in California

To demonstrate the applicability of our method for real-time forecasting (nowcasting), we propose to re-analyse the initial COVID-19 outbreak in California. Our analysis considers the graph formed by the 58 California counties and their corresponding incidence data from January 31, 2020, to April 1, 2020. Incidence data are known to suffer from considerable variations throughout the week and underreporting (Graham, 2020; Lipsitch et al., 2015), which can significantly skew the estimation of the reproductive number and the predicted epidemic trajectory (Siegenfeld et al., 2020). It is, in fact, standard practice to first smooth the raw incidence data through a 7-day rolling average to mitigate irregularities due to weekend reporting (see, for instance, case reporting by the Johns Hopkins Coronavirus Resource Center (Johns Hopkins CRC, 2024), or by the World Health Organization (World Health Organization, 2024)). However, while the importance of accounting for temporal variations seems widely recognized, the extension of this practice to the spatial domain remains unexplored. Yet, due to movement of people between counties, it seems logical to assume that neighbouring counties should expect similar prevalence: any deviations from this pattern could likely be attributed to measurement bias, and should therefore be corrected. We propose here to correct such biases by employing our method, which we believe leads to a more precise estimation of local disease prevalence.

We determine the optimal regularization parameter λ by splitting the data between training and testing. For a given λ , we train a simple epidemic model on the smoothed data from January 31 to March 14, 2020. This first part of the data is used to estimate the basic reproduction number (R_0) and to forecast the following two weeks. Fitting and predictions are performed here using the packages `R0` and `earlyR`, which are designed for the estimation and prediction of early-stage epidemics. These packages use the incidence matrix to estimate R_0 and predict the epidemic trajectory using simple growth models, which only require as input an estimate of the mean and standard deviation of the serial interval. Following Du et al. (2020), we set these values at $\mu = 3.95$ and $\sigma = 4.75$, respectively.

We then choose the regularization parameter (λ) that yields the most accurate forecast on the next three days, from March 15 to March 18, 2020. Having chosen an appropriate λ , we evaluated our model in the period from March 19 to March 31, 2020. The primary reason for this partitioning is that COVID reporting began only in early March for a subset of California counties.




Fig. 6: Error in the estimation of β , as a function of graph density for the kNN graph (parametrized by k , represented by each column) and number of observation steps K . Results are aggregated over 500 experiments. Points denote the mean error in estimating β , and error bars indicate the interquartile range.

However, on March 16, 2020, California implemented a lockdown order, aiming to aggressively reduce R_0 and therefore changing the dynamics of the epidemic. However, we anticipate that the incidence immediately after the lockdown reflects the initial dynamics of the epidemic spread, as such changes are not instantaneous given the incubation period of the virus of around 4 to 5 days.

Given the lack of individual-level data and having only access county-level aggregates, we find it easier to model each county as a 'super node' - or equivalently, the entire population as a graph such that each inhabitant is connected to all other inhabitants of the same county, as well as those in neighbouring counties. Letting n_c denote the number of inhabitants in county c , this approach thus requires solving the optimisation problem defined as follows:

$$\arg \min_{p \in \mathbb{R}^C} \sum_{c \in [C]} I_c (1 - p_c)^2 + (n_c - I_c) p_c^2 + \lambda \sum_{c, c'} n_c n_{c'} w_{c, c'} |p_c - p_{c'}| \quad (15)$$




Fig. 7: ℓ_1 error of our one-bit TV denoiser compared to $\hat{p}^{\text{naive}} = y_{\text{observed}}$ in the presence of missing data as a function of the proportion of missing values for the 5-NN graph. Here we vary $k_0 \in \{10, 20, 30\}$ (one value of k_0 per row), while columns corresponds to various values of β .

where $w_{c,c'}$ is a set of appropriately chosen weights. In the previous equation, I_c denotes the total number of infected inhabitants in county c , while $n_c - I_c$ represents the number of susceptible individuals in county c . Here, consistently with our simulation experiments, we pick $w_{c,c'} = \frac{1}{\sum_{c' \sim c} n_{c'}}$ the inverse of the degree of each node.

The results of this additional preprocessing procedure are presented in Table 3. Overall, for 85% of the counties on which the epidemic could be estimated (meaning that these counties had at least one COVID case before March 20th), our estimator provides a significant improvement over the raw data. This improvement is particularly striking in well populated and connected areas, such as the San Francisco Bay, where our estimator systematically improves the naive estimator.

6.2. The ExFLU Dataset of Aiello et al. (Aiello et al., 2016)

In this second dataset, we reinvestigate the epidemic dataset of Aiello et al. (2016). This dataset records a flu epidemic outbreak over a college campus in Winter 2013. The study's initial aim was to examine the influence of social interventions on the transmission of respiratory infections. A total of 590 university students were enrolled in the study, engaging in weekly surveys regarding symptoms resembling influenza-like illness (ILI) and their social interactions. This enables the generation of a weekly contact graph, which helps in mapping potential epidemic propagation phenomena. More specifically here, to evaluate the potential of our method to denoise contact tracing networks, we consider the ego network of each participant declaring experiencing ILI symptoms induced by the contact graph spanning the week of, and the week after their symptom onset. We define as a contagion event any symptomatic or asymptomatic individual having contracted an ILI illness based on the day of or after the contact's reported symptom onset. This creates a total of 97 ego networks. We then evaluate the ability of our one-bit denoising approach

county	MSE Prediction (Raw Data)	MSE Prediction (Smoothed Data)
Contra Costa	60	27
Fresno	0	347
Humboldt	0	0
Los Angeles	652	425
Marin	285	17
Orange	136	73
Placer	0	1
Riverside	32	25
Sacramento	25	18
San Benito	1,382	297
San Diego	2,140	716
San Francisco	63	23
San Joaquin	2,157	351
San Mateo	158	114
Santa Clara	101	78
Solano	2	0
Sonoma	147	3
Stanislaus	784	211
Tulare	4,558	3
Ventura	8,679	3,043
Yolo	3	32

Table 3: Comparison of the results in terms of the MSE between predicted epidemic trajectory and realized trajectory for the period from March 20th to March 30th 2023, using the raw data or the spatially smoothed data from our estimator.

in correctly estimating the ILI statuses of the nodes within this contact network the following week (see Figure 8).

To select the appropriate regularisation parameter, we split the ego networks into two sets, and select lambda on the first, and evaluate on the second. We randomize this procedure and average the results over 100 iterations. As in the previous example, we quantify the error using the l_2 norm. We obtain, in average, an l_2 error of 0.0482, compared to 0.0493 when setting the regularisation parameter to 0. Despite the noisy nature of this dataset, we are therefore able to improve upon the naive estimator ($\lambda = 0$). This suggests that our TV denoiser might be effective for epidemic nowcasting using contact tracing networks.

7. RELATED WORKS

In this section we provide a more in-depth overview of related works.

Total Variation Denoising. Total Variation (TV) denoising (also known as Fused Lasso) has experienced a considerable practical success in the area of image denoising. First introduced in Rudin et al. (1992), it benefits from the strong theoretical guarantee in the case of Gaussian noise. We refer the readers to the works of Mammen & Geer (1997); Dalalyan et al. (2017) for chain graphs and Wang et al. (2016); Hütter & Rigollet (2016); Padilla et al. (2018) for more




Fig. 8: Examples of contact networks induced by taking the 2-hop induced subgraph around participants exhibiting ILI-like symptoms. The contact network here represents all reported contacts between participants on the week of and the week following the symptom onset of the egonet source node. Colours denote different illness statuses. Grey: healthy participant. Yellow: participant sick with ILI symptoms on the week prior to the source node’s symptom onset. Red: participant healthy before the source node’s symptom onset and developing ILI symptoms in the week following the source node’s symptom onset.

general graphs. Earlier works on TV denoising (Mammen & Geer, 1997) marked an important step by obtaining the first statistical guarantees. Another notable development in understanding the behaviour of the TV denoiser was presented in Needell & Ward (2013). In their work, the authors focused on cases where the noise possesses a small l_2 norm but is otherwise arbitrary. This framework is commonly encountered in the literature on noisy compressed sensing. An analysis of the statistical performance of total variation image denoising was performed in Wang et al. (2016) for sub-Gaussian noise and for Gaussian white noise model in Hütter & Rigollet (2016) where, in particular, optimal fast rates for the 2D grid are obtained. By contrast, in this work, we built on the rich literature on total variation denoising and provide an analysis of its statistical guarantees in the case of binary observations (see Section 2).

Epidemic spreads on networks. Works on epidemic spread on networks and, more generally, on graph dynamical systems, can be found in various areas. Common types of network epidemic models include SIR and SEIR models, Agent-Based models, Percolation models and etc. Network-based stochastic models, such as the Susceptible-Infectious-Recovered (SIR) and Susceptible-Exposed-Infectious-Recovered (SEIR) models, are commonly used to simulate epidemic dynamics. Transition rates between compartments are estimated based on network structure and contact patterns (Newman, 2002b; Kenah & Robins, 2007). Researchers have extended these models to incorporate heterogeneous mixing patterns and complex network structures (Miller & Volz, 2013). Agent-based models (ABMs) have gained popularity for simulating realistic human behavior within network contexts. For example, Balcan et al. (2009) used ABMs to investigate the impact of human mobility on epidemic spread, considering factors like travel patterns and contact behaviour. These models provide insights into the effectiveness of public health interventions and behavioral changes during outbreaks.

Stochastic network models, such as the temporal exponential random graph model (TERGM), capture dynamic changes in network structure over time. They are employed to study how changes in contact patterns impact epidemic dynamics (see, for example, Volz & Meyers (2007)). These models provide a statistical framework for exploring network evolution during epidemics.

Discrete-time Markov Processes (DTMP) and Continuous-time Markov Processes (CTMP) are commonly used in network epidemic models to represent the dynamics of disease spread or information diffusion within a population (Gómez et al., 2010; Zhang & Moura, 2014). In DTMP, time is divided into discrete intervals or time steps. For example, in a SIR (Susceptible-Infectious-Recovered) model, susceptible individuals may become infected with a certain probability during each time step. In Continuous-time Markov Processes (CTMP), time is treated as a continuous variable, allowing for modeling events that occur at any moment in time. State transitions can occur at any time, not just at discrete time steps. Events such as contacts, infections, recoveries, or interventions are represented as stochastic processes. The rates at which these events occur depend on factors such as contact frequencies, transmission probabilities, and recovery rates.

Bayesian inference methods have gained popularity for parameter estimation in epidemic models. Markov Chain Monte Carlo (MCMC) algorithms, such as the Metropolis-Hastings algorithm, are used to sample from posterior distributions of model parameters (Jewell et al., 2009; Hu et al., 2017). Bayesian frameworks allow, in particular for uncertainty quantification and model selection. However, to the best of our knowledge, no statistical guarantees have yet been provided for the recovery of the parameters.

Sampling and contact tracing data are essential for estimating epidemic parameters. Lloyd-Smith et al. (2005) used network-based contact tracing to reconstruct transmission chains in emerging infectious diseases. This approach provides insights into superspreading events and helps identify high-risk nodes. Statistical techniques for inferring contact networks from observed epidemic data have been also developed. For instance, Rocha et al. (2011) introduced a likelihood-based method for reconstructing contact networks using observed infection times. Such methods are crucial for understanding the underlying network structure in real-world epidemics.

The study of infectious disease spread in the context of partially observed epidemics is essential in real-world situations where complete information about individuals' health status is often unavailable or uncertain. Research in this area has addressed various aspects of the problem, including modeling, inference, and control strategies (for a review of the latter, see for instance Britton & Giardina (2016)). Several models have been developed that incorporate partial information. For example, the Susceptible-Infectious-Unknown (SIU) model extends the classic Susceptible-Infectious-Recovered (SIR) model to account for individuals with unknown health status.

Within the broader context of partially-observed Markov processes, several techniques have been incorporated into the R package POMP (King et al., 2016). Notably, one of these methods is maximum iterated filtering (MIF) (Ionides et al., 2011). Other approaches include Gaussian approximation of the epidemic density-dependent Markovian jump process (Narci et al., 2021). Additionally, likelihood-free approaches such as approximate Bayesian computation based on sequential Monte Carlo (Toni et al., 2009; Sisson et al., 2007) and particle Markov chain Monte Carlo (Andrieu et al., 2010) have emerged. These approaches aim to infer the hidden states of individuals (infected or healthy) based on observed data, such as reported cases or symptom data.

ACKNOWLEDGMENTS

Claire Donnat gratefully acknowledges support from NSF Award Number RI:2238616, as well as the resources provided by the University of Chicago's Research Computing Center. The work of Olga Klopp was funded by CY Initiative (grant "Investissements d'Avenir" ANR-16-IDEX-0008) and Labex MME-DII (ANR11-LBX-0023-01). This work was done while O. Klopp and C. Donnat were visiting the Simons Institute for the Theory of Computing. The work of N. Verzelen has been partially supported by ANR-21-CE23-0035 (ASCAI).

SUPPLEMENTARY MATERIAL

The Supplementary Material includes the proofs of the theorems and proposition established in this paper, as well as further results stemming from our synthetic experiments.

A. PROOF OF THEOREM 1

The proof of Theorem 1 is close in spirit to the proof of the Theorem 2 in (Hütter & Rigollet, 2016) the main difference being the control of the noise term.

Applying the chain rule, we have that the subdifferential of the l_1 term is given by

$$\partial \|Dp\|_1 = D^T \text{sign}(Dp)$$

where

$$\text{sign}(x)_i = \begin{cases} 1 & \text{if } x_i > 0, \\ [-1, 1] & \text{if } x_i = 0, \\ -1 & \text{if } x_i < 0. \end{cases}$$

Using first order optimality condition for the convex problem (2), we derive that there exists $z \in \text{sign}(D\widehat{p})$ such that for any $\bar{p} \in \mathbb{R}^n$, we have

$$\frac{2}{n} \langle \bar{p}, Y - \widehat{p} \rangle = \lambda \langle \bar{p}, D^T z \rangle = \lambda \langle D\bar{p}, z \rangle. \quad (\text{A1})$$

This implies that

$$\frac{2}{n} \langle \widehat{p}, Y - \widehat{p} \rangle = \lambda \|D\widehat{p}\|_1 \quad \text{and}$$

$$\frac{2}{n} \langle p^*, Y - \widehat{p} \rangle \leq \lambda \|Dp^*\|_1.$$

Subtracting the two terms above, we get

$$\frac{2}{n} \|p^* - \widehat{p}\|_2^2 \leq \frac{2}{n} \langle \xi, \widehat{p} - p^* \rangle + \lambda \|Dp^*\|_1 - \lambda \|D\widehat{p}\|_1.$$

Next, we control the error term. Let Π denote the projection on $\ker D$, the kernel of D . Then, by the definition of $D^\dagger$, $D^\dagger D = I - \Pi$ is the projection on $(\ker D)^\perp$. We have that $\ker D = \ker(D^T D) = \ker L$. We assume that G is connected, which implies that zero is an eigenvalue of L of multiplicity one and $\dim \ker D = 1$. We can write

$$\begin{aligned} \langle \xi, \widehat{p} - p^* \rangle &= \langle \Pi \xi, \widehat{p} - p^* \rangle + \langle (I - \Pi) \xi, \widehat{p} - p^* \rangle \\ &= \langle \Pi \xi, \Pi(\widehat{p} - p^*) \rangle + \langle D^\dagger D \xi, \widehat{p} - p^* \rangle \\ &= \langle \Pi \xi, \Pi(\widehat{p} - p^*) \rangle + \langle \xi, D^\dagger D(\widehat{p} - p^*) \rangle \\ &= \langle \Pi \xi, \Pi(\widehat{p} - p^*) \rangle + \left\langle \left(D^\dagger \right)^T \xi, D(\widehat{p} - p^*) \right\rangle. \end{aligned} \quad (\text{A2})$$

To bound the right hand side in (A2), we first use the Hölder inequality to get

$$\langle \Pi \xi, \Pi(\widehat{p} - p^*) \rangle \leq \|\Pi \xi\|_2 \|\Pi(\widehat{p} - p^*)\|_2.$$

Note that $L \mathbf{1}_n = \mathbf{0}_n$ and, for any vector $v \in \mathbb{R}^n$, $\Pi v = \frac{1}{n} \langle v, \mathbf{1}_n \rangle \mathbf{1}_n$ which implies

$$\|\Pi \xi\|_2 = \frac{1}{\sqrt{n}} |\langle \xi, \mathbf{1}_n \rangle|.$$

Now, using Hoeffding's inequality we get that, with probability at least $1 - \delta/2$,

$$|\langle \xi, \mathbf{1}_n \rangle| = \left| \sum_{i=1}^n \xi_i \right| \leq \sqrt{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)} \quad \text{and} \quad \|\Pi \xi\|_2 \leq \sqrt{\frac{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)}{n}}. \quad (\text{A3})$$

On the other hand, we get

$$\|\Pi(\widehat{p} - p^*)\|_2 = \frac{1}{\sqrt{n}} |\langle \widehat{p} - p^*, \mathbf{1}_n \rangle| \leq \frac{1}{\sqrt{n}} \|\widehat{p} - p^*\|_1 \quad (\text{A4})$$

and (A3–A4) imply

$$\langle \Pi \xi, \Pi(\widehat{p} - p^*) \rangle \leq \frac{\sqrt{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)}}{n} \|\widehat{p} - p^*\|_1. \quad (\text{A5})$$

For the second term in (A2) we have that, with probability at least $1 - \delta/2$,

$$\begin{aligned} \langle (D^\dagger)^T \xi, D(\widehat{p} - p^*) \rangle &\leq \|(D^\dagger)^T \xi\|_\infty \|D(\widehat{p} - p^*)\|_1 \\ &\leq \sqrt{2}\rho \log\left(\frac{4n^2}{\delta}\right) \|D(\widehat{p} - p^*)\|_1, \end{aligned} \quad (\text{A6})$$

where the last inequality follows from Lemma A1. Putting together (A5) and (A6) we get that, with probability at least $1 - \delta$,

$$\langle \xi, \widehat{p} - p^* \rangle \leq \sqrt{2}\rho \log\left(\frac{4n^2}{\delta}\right) \|D(\widehat{p} - p^*)\|_1 + \frac{\sqrt{2\|p^*\|_0}}{n} \log\left(\frac{4}{\delta}\right) \|\widehat{p} - p^*\|_1,$$

since $\|p^*\|_1 \leq \|p^*\|_0$. This leads us to

$$\begin{aligned} \|p^* - \widehat{p}\|_2^2 &\leq \sqrt{2}\rho \log\left(\frac{4n^2}{\delta}\right) \|D(\widehat{p} - p^*)\|_1 + n\lambda \|Dp^*\|_1 - n\lambda \|D\widehat{p}\|_1 \\ &\quad + \frac{\sqrt{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)}}{n} \|\widehat{p} - p^*\|_1. \end{aligned}$$

Using $n\lambda = \sqrt{2}\rho \log\left(\frac{4n^2}{\delta}\right)$ and the triangle inequality we get

$$\|p^* - \widehat{p}\|_2^2 \leq 2n\lambda (\|D(\widehat{p} - p^*)\|_1 + \|(Dp^*)_{T^c}\|_1) + \frac{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)}{n} + \frac{1}{4} \|\widehat{p} - p^*\|_2^2.$$

Now, applying Definition 1 of compatibility factor we can write

$$\frac{1}{2} \|p^* - \widehat{p}\|_2^2 \leq \frac{8\rho^2 |T|}{\kappa_T^2} \log^2\left(\frac{4n^2}{\delta}\right) + 2\sqrt{2}\rho \log\left(\frac{4n^2}{\delta}\right) \|(Dp^*)_{T^c}\|_1 + \frac{2\|p^*\|_0 \log\left(\frac{4}{\delta}\right)}{n}$$

which proves the statement of Theorem 1.

A.1. Lemma A1

LEMMA A1. Fix $\delta \in [0, 1]$. For $i \in [n]$, let z_i be independent centered Bernoulli random variables with parameter q_i and $z = (z_i)_{i \in [n]}$. Then, with probability at least $1 - \delta$,

$$\|(D^\dagger)^T z\|_\infty \leq \sqrt{2}\rho \log\left(\frac{2n^2}{\delta}\right).$$

Proof. We start by applying Bernstein's inequality to $\sum_{i=1}^n D_{ij}^\dagger z_i$. Note that $|D_{ij}^\dagger z_i| \leq \rho$ and $\sum_{i=1}^n (D_{ij}^\dagger)^2 E z_i^2 \leq \sum_{i=1}^n (D_{ij}^\dagger)^2 q_i \leq \rho^2$. Then, the Bernstein's inequality implies that, with probability at least $1 - \frac{1}{n^2 \delta}$,

$$\left| \sum_{i=1}^n D_{ij}^\dagger z_i \right| \leq \sqrt{2} \rho \log \left(\frac{2n^2}{\delta} \right).$$

Now, applying the union bound and using $m \leq n^2$ we get the result of Lemma A1. $\square$

B. PROOF OF PROPOSITION 2

Given a $p \in \mathbb{R}^n$, let $\mathcal{P}(p) = (\mathcal{P}(p_i))_{i=1}^n \in [0, 1]^n$ be defined as follows

$$\mathcal{P}(p_i) = \begin{cases} p_i & \text{if } p_i \in [0, 1], \\ 0 & \text{if } p_i < 0, \\ 1 & \text{if } p_i > 0. \end{cases}$$

Let $\mathcal{F} = \frac{1}{n} \sum_{i=1}^n (y_i^{k_0} - p_i)^2 + \lambda \|Dp\|_1$. Note that, for any p , $\mathcal{F}(\mathcal{P}(p)) \leq \mathcal{F}(p)$. So, without loss of generality, we assume that $\widehat{p}_i \in [0, 1]^n$. Now take $\bar{p} = (\bar{p}_i)_{i=1}^n$ such that

$$\bar{p}_i = \begin{cases} -1 & \text{if } i \in S^c, \\ 0 & \text{if not.} \end{cases}$$

Then, the first order optimality condition (A1) implies

$$\langle \bar{p}, (p^* - \widehat{p})_S \rangle - \langle \bar{p}, \widehat{p}_{S^c} \rangle + \langle \bar{p}, \xi \rangle = \frac{n\lambda}{2} \langle D\bar{p}, z \rangle.$$

Note that $\xi_{S^c} = 0$ which implies

$$\|\widehat{p}_{S^c}\|_1 = \frac{n\lambda}{2} \langle D\bar{p}, z \rangle \leq \frac{n\lambda}{2} \|D\bar{p}\|_1.$$

We have that for any $e = (i, j) \in E$, $(D\bar{p})_e = 0$ if $(i, j) \in S \times S$ or if $(i, j) \in S^c \times S^c$ which implies that $\|D\bar{p}\|_1 \leq 2sd_{\max}(S)$. Then we compute

$$\|p^* - \widehat{p}\|_1 = \|(p^* - \widehat{p})_S\|_1 + \|\widehat{p}_{S^c}\|_1 \leq n\lambda sd_{\max}(S) + \sqrt{s} \|p^* - \widehat{p}\|_2.$$

Combining this bound with Theorem 1 we get the statement of Proposition 2.

C. PROOF OF THEOREM 2

For simplicity, we drop the *miss* index, that is we write $\widehat{p}_{miss} = \widehat{p}$. By first order optimality condition for the convex problem (11), there exists $z \in \text{sign}(D\widehat{p})$ such that for any $\bar{p} \in \mathbb{R}$, we have

$$\frac{2}{n} \left\langle \bar{p}, \mathfrak{M} \circ (\widetilde{Y} - \widehat{p}) \right\rangle = \lambda \langle \bar{p}, D^T z \rangle = \lambda \langle D\bar{p}, z \rangle, \quad (\text{C1})$$

which implies

$$\frac{2}{n} \left\langle \widehat{p}, \mathfrak{M} \circ (\widetilde{Y} - \widehat{p}) \right\rangle = \lambda \|D\widehat{p}\|_1 \quad \text{and}$$

$$\frac{2}{n} \left\langle p^*, \mathfrak{M} \circ (\widetilde{Y} - \widehat{p}) \right\rangle \leq \lambda \|Dp^*\|_1.$$

Subtracting the above two, we obtain

$$\frac{2}{n} \|\mathfrak{M} \circ (p^* - \hat{p})\|_2^2 \leq \frac{2}{n} \langle \mathfrak{M} \circ \xi, \hat{p} - p^* \rangle + \lambda \|Dp^*\|_1 - \lambda \|D\hat{p}\|_1.$$

The control of the error term is similar to the case with complete observations and we only sketch it. Let $\eta = \mathfrak{M} \circ \xi = (\eta_i)_{i \in [n]}$. Note that the η_i 's are independent Bernoulli random variables with parameter $\pi_i p_i^*$. Then, following similar argument as in the proof of Theorem 1 and using Bernstein's inequality, we get with probability at least $1 - \delta/2$,

$$\begin{aligned} \langle \Pi \eta, \Pi (\hat{p} - p^*) \rangle &\leq \|\Pi \eta\|_2 \|\Pi (\hat{p} - p^*)\|_2 \\ &\leq \frac{\sqrt{2\|\pi \circ p^*\|_1 \log\left(\frac{4}{\delta}\right)}}{n} \|\hat{p} - p^*\|_1 \\ &\leq \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1 \log\left(\frac{4}{\delta}\right)}}{n} \|\hat{p} - p^*\|_{l_2(\pi)} \end{aligned} \quad (\text{C2})$$

Similarly, we have that with probability $1 - \delta/2$

$$\begin{aligned} \left\langle \left(D^\dagger\right)^T \eta, D(\hat{p} - p^*) \right\rangle &\leq \left\| \left(D^\dagger\right)^T \eta \right\|_\infty \|D(\hat{p} - p^*)\|_1 \\ &\leq \sqrt{2\rho} \log\left(\frac{4n^2}{\delta}\right) \|D(\hat{p} - p^*)\|_1 \end{aligned} \quad (\text{C3})$$

where the last inequality follows from Lemma A1. Putting together (C2) and (C3) we get that, with probability at least $1 - \delta$,

$$\begin{aligned} \langle \mathfrak{M} \circ \xi, \hat{p} - p^* \rangle &\leq \sqrt{2\rho} \log\left(\frac{4n^2}{\delta}\right) \|D(\hat{p} - p^*)\|_1 \\ &\quad + \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1 \log\left(\frac{4}{\delta}\right)}}{n} \|\hat{p} - p^*\|_{l_2(\pi)} \end{aligned}$$

and

$$\begin{aligned} \|\mathfrak{M} \circ (p^* - \hat{p})\|_2^2 &\leq \sqrt{2\rho} \log\left(\frac{4n^2}{\delta}\right) \|D(\hat{p} - p^*)\|_1 + \frac{n\lambda}{2} \|Dp^*\|_1 - \frac{n\lambda}{2} \|D\hat{p}\|_1 \\ &\quad + \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1 \log\left(\frac{4}{\delta}\right)}}{n} \|\hat{p} - p^*\|_{l_2(\pi)}. \end{aligned}$$

For $T = \text{supp}(Dp^*)$ and $\delta = 4/n$, using $n\lambda = 9\sqrt{2\rho} \log(n)$, we can write

$$\begin{aligned} \|\mathfrak{M} \circ (p^* - \hat{p})\|_2^2 + \frac{n\lambda}{6} \|D\hat{p}_{T^c}\|_1 &\leq \frac{5n\lambda}{6} \|D(p^* - \hat{p})_T\|_1 \\ &\quad + \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1 \log(n)}}{n} \|\hat{p} - p^*\|_{l_2(\pi)} \end{aligned} \quad (\text{C4})$$

and

$$\|(D\hat{p})_{T^c}\|_1 \leq 5 \|(D(p^* - \hat{p}))_T\|_1 + \frac{2\sqrt{\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1 \log(n)}}{3n\rho \log(n)} \|\hat{p} - p^*\|_{l_2(\pi)}. \quad (\text{C5})$$

Next, we need to provide a lower bound for $\|\mathfrak{M} \circ (p^* - \hat{p})\|_2^2$. It can be done on a set of vectors that satisfy (C5) if $\|p^* - \hat{p}\|_{l_2(\pi)}^2$ is not too small. We start by introducing the corresponding set of vectors. Let

$$\Delta = \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1}}{3n\rho}. \quad (\text{C6})$$

For a set of indices T and $\epsilon \geq \frac{\log(n)}{0.01 \log(6/5)}$, we denote by $C(T, \epsilon)$ the following set of vectors:

$$C(T, \epsilon) = \left\{ v \in [0, 1]^n : \|v\|_{l_2(\pi)}^2 \geq \epsilon, \| (Dv)_{T^c} \|_1 \leq 5 \| (Dv)_T \|_1 + \Delta \|v\|_{l_2(\pi)} \right\}.$$

Next result provides partial isometry for vectors in $C(T, \epsilon)$:

LEMMA C1. *With probability higher than $1 - 8/n$, we have for all $v \in C(T, \epsilon)$*

$$\|\mathfrak{M} \circ v\|_2^2 \geq \frac{\|v\|_{l_2(\pi)}^2}{2} - 77 \left\{ \frac{36\rho\sqrt{\kappa_\pi|T|\log(n)}}{\kappa_T} + \frac{\sqrt{\|\pi^{-1}\|_1}}{n} \left(5\sqrt{\|\pi\|_1} \log(n) + 1 \right) \right\}^2.$$

We now consider two cases, depending on whether the vector $p^* - \hat{p}$ belongs to the set $C(T, \epsilon)$ or not.

Case 1: Suppose first that $\|p^* - \hat{p}\|_{l_2(\pi)}^2 < \epsilon$, then the statement of Theorem 2 is true.

Case 2: It remains to consider the case $\|p^* - \hat{p}\|_{l_2(\pi)}^2 \geq \epsilon$. Let $T = \text{supp}(Dp^*)$. Then, (C5) implies that $p^* - \hat{p} \in C(T, \epsilon)$ with probability at least $1 - 4/n$. Then, combining (C4) and Lemma C1 we get

$$\begin{aligned} \frac{1}{2} \|p^* - \hat{p}\|_{l_2(\pi)}^2 &\leq 77 \left\{ \frac{36\rho\sqrt{\kappa_\pi|T|\log(n)}}{\kappa_T} + \frac{\sqrt{\|\pi^{-1}\|_1}}{n} \left(5\sqrt{\|\pi\|_1} \log(n) + 1 \right) \right\}^2 \\ &+ \frac{12\rho\sqrt{\kappa_\pi|T|\log(n)}}{\kappa_T} \|p^* - \hat{p}\|_{l_2(\pi)} + \frac{\sqrt{2\|\pi \circ p^*\|_1 \|\pi^{-1}\|_1} \log(n)}{n} \|\hat{p} - p^*\|_{l_2(\pi)}. \end{aligned} \quad (C7)$$

Now the statement of Theorem 2 follows applying $2ab \leq 8a^2 + \frac{1}{8}b^2$ to the last two terms in (C7).

C.1. Proof of Lemma C1

Set $\nu = \frac{\log(n)}{0.01 \log(6/5)}$, $\alpha = \frac{6}{5}$ and

$$\mathcal{E} = 77 \left\{ \frac{36\rho\sqrt{\kappa_\pi|T|\log(n)}}{\kappa_T} + \frac{\sqrt{\|\pi^{-1}\|_1}}{n} \left(5\sqrt{\|\pi\|_1} \log(n) + 1 \right) \right\}^2.$$

The proof of Lemma C1 is based on the peeling argument. Denote by $\mathcal{B}$ the set that contains the complement of the event that we are interested in:

$$\mathcal{B} = \left\{ \exists v \in C(T, \epsilon) \text{ such that } \left| \|\mathfrak{M} \circ v\|_2^2 - \|v\|_{l_2(\pi)}^2 \right| > \frac{1}{2} \|v\|_{l_2(\pi)}^2 + \mathcal{E} \right\}.$$

For $r > \nu$ we will consider the following set of vectors:

$$C(T, \epsilon, r) = \left\{ v \in C(T, \epsilon) : \|v\|_{l_2(\pi)}^2 \leq r \right\}$$

and, for $l \in \mathbb{N}$, the following events

$$\mathcal{B}_l = \left\{ \exists v \in C(T, \epsilon, \alpha^l \nu) : \left| \|\mathfrak{M} \circ v\|_2^2 - \|v\|_{l_2(\pi)}^2 \right| > \frac{5}{12} \alpha^l \nu + \mathcal{E} \right\},$$

$$S_l = \left\{ v \in C(T, \epsilon) : \alpha^{l-1} \nu \leq \|v\|_{l_2(\pi)}^2 \leq \alpha^l \nu \right\}.$$

If the event $\mathcal{B}$ holds for some vector $v \in C(T, \epsilon)$, then v belongs to some S_l and

$$\begin{aligned} \left| \|\mathfrak{M} \circ v\|_2^2 - \|v\|_{l_2(\pi)}^2 \right| &> \frac{1}{2} \|v\|_{l_2(\pi)}^2 + \mathcal{E} \\ &> \frac{1}{2} \alpha^{l-1} \nu + \mathcal{E} \\ &= \frac{5}{12} \alpha^l \nu + \mathcal{E} \end{aligned} \quad (C8)$$

which implies $\mathcal{B} \subset \cup \mathcal{B}_l$. Next, we apply the union bound using the following lemma:

LEMMA C2. *Let*

$$Z_r = \sup_{v \in C(T, \epsilon, r)} \left| \|\mathfrak{M} \circ v\|_2^2 - \|v\|_{l_2(\pi)}^2 \right|.$$

We have that

$$pr \left(Z_r \geq \frac{5}{12}r + \mathcal{E} \right) \leq 4e^{-c_1 r}$$

with $c_1 \geq 0.01$.

Lemma C2 implies that $\mathbb{P}(\mathcal{B}_l) \leq 4 \exp(-c_1 \alpha^l v)$. Using $e^x \geq x$, we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{B}) &\leq \sum_{l=1}^{\infty} \mathbb{P}(\mathcal{B}_l) \leq 4 \sum_{l=1}^{\infty} \exp(-c_1 \alpha^l v) \\ &\leq 4 \sum_{l=1}^{\infty} \exp(-c_1 v \log(\alpha) l) \leq \frac{4 \exp(-c_1 v \log(\alpha))}{1 - \exp(-c_1 v \log(\alpha))} \\ &= \frac{4 \exp(-\log(n))}{1 - \exp(-\log(n))}. \end{aligned}$$

This completes the proof of Lemma C1.

C.2. Proof of Lemma C2

We first provide an upper bound on $\mathbb{E}(Z_r)$ and then show that Z_r concentrates around its expectation. By definition of Z_r we have

$$Z_r = \sup_{v \in C(T, \epsilon, r)} \left| \sum_i m_i v_i^2 - E \left(\sum_i m_i v_i^2 \right) \right|.$$

Using a standard symmetrization argument (see e.g. (Ledoux, 2001)) we obtain

$$\begin{aligned} E(Z_r) &= E \left(\sup_{v \in C(T, \epsilon, r)} \left| \sum_i m_i v_i^2 - E \left(\sum_i m_i v_i^2 \right) \right| \right) \\ &\leq 2E \left(\sup_{v \in C(T, \epsilon, r)} \left| \sum_i \zeta_i m_i v_i^2 \right| \right) \end{aligned}$$

where $\{\zeta_i\}$ is an i.i.d. Rademacher sequence. Then, using that $v_i \leq 1$, the contraction inequality (see e.g. (Koltchinskii, 2011, Theorem 2.2)) yields

$$E(Z_r) \leq 8E \left(\sup_{v \in C(T, \epsilon, r)} \left| \sum_i \zeta_i m_i v_i \right| \right) = 8E \left(\sup_{v \in C(T, \epsilon, r)} |\langle \psi, v \rangle| \right)$$

where $\psi = (\zeta_i m_i)_{i \in [n]}$. We can write

$$\langle \psi, v \rangle = \langle \Pi \psi, v \rangle + \langle (I - \Pi) \psi, v \rangle = \langle \Pi \psi, \Pi v \rangle + \left\langle \left(D^\dagger \right)^T \psi, D v \right\rangle \quad (\text{C9})$$

where Π denote the projection on the kernel of D . To bound the first term on the right hand side in (C9), we first use the Hölder inequality:

$$|\langle \Pi \psi, \Pi v \rangle| \leq \|\Pi \psi\|_2 \|\Pi v\|_2.$$

Using the same argument as in the proof of Theorem 1, we have that

$$\|\Pi \psi\|_2 = \frac{1}{\sqrt{n}} |\langle \psi, \mathbf{1}_n \rangle| \quad \text{and} \quad \|\Pi v\|_2 = \frac{1}{\sqrt{n}} |\langle v, \mathbf{1}_n \rangle|.$$

For $v \in C(T, \epsilon, r)$ we get

$$\|\Pi v\|_2 \leq \frac{1}{\sqrt{n}} \|v\|_1 \leq \sqrt{\frac{\|\pi^{-1}\|_1}{n}} \|v\|_{l_2(\pi)} \leq \sqrt{\frac{\|\pi^{-1}\|_1 r}{n}}. \quad (\text{C10})$$

On the other hand, Bernstein's inequality implies that that, with probability at least $1 - \frac{1}{n}$, we have

$$\left| \sum_{i=1}^n \psi_i \right| \leq \sqrt{2\|\pi\|_1 \log(2n)}$$

and

$$E \left(\frac{1}{\sqrt{n}} |\langle \psi, \mathbf{1}_n \rangle| \right) \leq \sqrt{\frac{2\|\pi\|_1}{n}} \log(2n) + \frac{1}{\sqrt{n}} \quad (\text{C11})$$

where we used $|\langle \psi, \mathbf{1}_n \rangle| \leq n$. Putting together (C10) and (C11) we get that for all

$$\mathbb{E} \left[\sup_{v \in C(T, \epsilon, r)} |\langle \Pi \psi, \Pi v \rangle| \right] \leq \frac{\sqrt{\|\pi^{-1}\|_1 r}}{n} \left(\sqrt{2\|\pi\|_1 \log(2n)} + 1 \right). \quad (\text{C12})$$

For the second term in (C9), using again the Hölder inequality, we get

$$\left| \left\langle \left(D^\dagger \right)^T \psi, Dv \right\rangle \right| \leq \left\| \left(D^\dagger \right)^T \psi \right\|_\infty \|Dv\|_1. \quad (\text{C13})$$

Lemma A1 implies that, with probability at least $1 - 4/n$, we have

$$\left\| \left(D^\dagger \right)^T \psi \right\|_\infty \leq 3\sqrt{2}\rho \log(8n). \quad (\text{C14})$$

On the other hand, we have that

$$\left\| \left(D^\dagger \right)^T \psi \right\|_\infty = \sup_j \left| \sum_{i=1}^n D_{ij}^\dagger \psi_i \right| \leq \sup_j \sqrt{\sum_i \left(D_{ij}^\dagger \right)^2} \sqrt{\sum_i \psi_i^2} \leq \rho \sqrt{n}$$

which together with (C14) imply

$$E \left(\left\| \left(D^\dagger \right)^T \psi \right\|_\infty \right) \leq 3\sqrt{2}\rho \log(n) + \frac{4\rho}{\sqrt{n}} \leq 6\rho \log(n) \quad (\text{C15})$$

for $n \geq 4$. Note that, as $v \in C(T, \epsilon, r)$, we also obtain

$$\begin{aligned} \|Dv\|_1 &\leq \|(Dv)_T\|_1 + \|(Dv)_{T^c}\|_1 \leq 6\|(Dv)_T\|_1 + \Delta\|v\|_{l_2(\pi)} \\ &\leq \frac{6\sqrt{|T|}}{\kappa_T} \|v\|_2 + \Delta\sqrt{r} \\ &\leq \frac{6\sqrt{\kappa_\pi|T|}}{\kappa_T} \|v\|_{l_2(\pi)} + \Delta\sqrt{r} \quad \text{using (10)} \\ &\leq \left(\frac{6\sqrt{\kappa_\pi|T|}}{\kappa_T} + \Delta \right) \sqrt{r}. \end{aligned} \quad (\text{C16})$$

Using (C15) and (C16) we get

$$E \left(\sup_{v \in C(T, \epsilon, r)} \left| \left\langle \left(D^\dagger \right)^T \psi, Dv \right\rangle \right| \right) \leq 6\rho\sqrt{r} \log(n) \left(\frac{6\sqrt{\kappa_\pi|T|}}{\kappa_T} + \Delta \right). \quad (\text{C17})$$

Putting together (C12), (C17) and using the definition (C6) of Δ , we get

$$E(Z_r) \leq 8\sqrt{r} \left\{ \frac{36\rho\sqrt{\kappa_\pi|T|}\log(n)}{\kappa_T} + \frac{\sqrt{\|\pi^{-1}\|_1}}{n} \left(5\sqrt{\|\pi\|_1}\log(n) + 1 \right) \right\}$$

where we used $\|\pi \circ p^*\|_1 \leq \|\pi\|_1$. Now, $2ab \leq a^2 + b^2$ implies

$$\begin{aligned} E(Z_r) &\leq \frac{5}{24}r + 77 \left\{ \frac{36\rho\sqrt{\kappa_\pi|T|}\log(n)}{\kappa_T} + \frac{\sqrt{\|\pi^{-1}\|_1}}{n} \left(5\sqrt{\|\pi\|_1}\log(n) + 1 \right) \right\}^2 \\ &= \frac{5}{24}r + \mathcal{E} \end{aligned} \tag{C18}$$

Now we can show that Z_r concentrates around its expectation using Talagrand's concentration inequality (Talagrand, 1996), which in the current can be obtained by inverting the tail bound in Theorem 3.3.16 in (Giné & Nickl, 2016).

THEOREM C1. *Let $(S, \mathcal{S})$ be a measurable space and let $n \in \mathbb{N}$. Let X_k , $k = 1, \dots, n$ be independent S -valued random variables and let $\mathcal{F}$ be a countable set of functions $f = (f_1, \dots, f_n) : S^n \rightarrow [-K, K]^n$ such that $\mathbb{E}f_k(X_k) = 0$ for all $f \in \mathcal{F}$ and $k = 1, \dots, n$. Set*

$$Z := \sup_{f \in \mathcal{F}} \sum_{k=1}^n f_k(X_k) .$$

Define the variance proxy

$$V_n := 2K\mathbb{E}Z + \sup_{f \in \mathcal{F}} \sum_{k=1}^n \mathbb{E}[(f_k(X_k))^2] .$$

Then, for all $t \geq 0$,

$$\text{pr}(Z - \mathbb{E}Z \geq t) \leq \exp\left(\frac{-t^2}{4V_n + (9/2)Kt}\right) .$$

Note that the functional $v \rightarrow \|\mathfrak{M} \circ v\|_2^2 - \|v\|_{l_2(\pi)}^2$ is continuous on the set of vectors $C(T, \epsilon, r)$. Hence it's enough to consider a dense countable subset of $C(T, \epsilon, r)$ and we may apply Talagrand's inequality to Z_r . We have for our particular case, since $\sup_{f \in \mathcal{F}} |f(X)| = \sup_{f \in \{\mathcal{F} \cup \{-\mathcal{F}\}\}} f(X)$, that

$$X_i = m_i - \mathbb{E}(m_i), \quad S = [-1, 1].$$

We can compute:

$$\sum_i \mathbb{E}[(f_k(X_k))^2] = \sum_i \mathbb{E}[m_i - \mathbb{E}(m_i)]^2 v_i^4 \leq \sum_i \pi_i v_i^2 = \|v\|_{l_2(\pi)}^2 \leq r$$

which implies that $V_n \leq \frac{17}{12}r + 2\mathcal{E}$. Hence, using (C18) and Talagrand's inequality with $t = \frac{5}{24}r + 0.5\mathcal{E}$, we obtain

$$\text{pr}\left(Z_r > \frac{5}{12}r + 1.5\mathcal{E}\right) \leq \exp(-c_1 r)$$

with $c_1 \geq 0.01$ which completes the proof of Lemma C2.

D. FURTHER RESULTS: SIMULATION STUDY

In this section, we report additional results of our simulation study.

D.1. *Algorithms and Compute time*

In our simulation study, we use the semismooth Newton based augmented Lagrangian method (SSNAL) proposed by Sun et al. (2021). SSNAL is a scalable algorithm that was originally designed to perform convex clustering to solve the following optimization problem:

$$\min_{X \in \mathbb{R}^{d \times n}} \frac{1}{2} \sum_{i=1}^n \|x_i - a_i\|^2 + \lambda \sum_{i < j} W_{ij} \|x_i - x_j\|_2$$

Note that in our case, since our vector p is one dimensional $\|p_i - p_j\|_2 = \sqrt{(p_i - p_j)^2} = |p_i - p_j|$, and our problem shares the same objective function as convex clustering. In their work, Sun et al. (2021) prove their algorithm to be theoretically efficient and practically efficient and robust. More specifically, the authors show that, as long as the dimension of the feature vector is small (i.e, in our case, equal to 1), this method has the potential of considerably improving upon existing solvers.

Evaluation. To highlight the scalability of SSNAL, we propose comparing here four different solvers on a set of synthetic experiments. More specifically, we compare the performances of:

1. *An off-the-shelf convex optimization solver (CVXPY),*
2. *The dual-coordinate descent algorithm* proposed by Tibshirani and Taylor Tibshirani & Taylor (2011). This method relies on solving the dual problem:

$$\text{minimize } \frac{1}{2} \|y - \Gamma^T u\|^2 \text{ subject to } \|u\|_\infty \leq \lambda,$$

yielding a simple algorithm.

3. *An ADMM-based approach:* ADMM is often the method of choice to deal with large scale constrained and has been suggested in multiple instance as the method of choice for graph-based penalties on large graphs Hallac et al. (2017),
4. *The SSNAL algorithm of Sun et al. (2021),* described in the previous paragraph.

In this first set of experiments, we propose comparing the performance of the different algorithms when the underlying is an Erdos Renyi random graph. The reason for considering this topology here is that this is a difficult setting, since the graph is quite dense. Since the complexity of the penalty depends on the number of edges, we plot the time requirements of the algorithm as the number of nodes increases and the penalty varies. The results are displayed in Figure 9 below.




Fig. 9: Computation time for the different algorithms




Fig. 10: Number of infections for the kNN graph as a function of β , n_{steps} and the parameter k in the k-NN graph. The healing rate is here fixed to $\gamma = 0.1$. Results are averaged over 100 independent experiments.

D.2. Synthetic Experiments: Additional Results

REFERENCES

- AHN, H. J. & HASSIBI, B. (2013). Global dynamics of epidemic spread over complex networks. In *52nd IEEE Conference on Decision and Control*.
- AIELLO, A. E., SIMANEK, A. M., EISENBERG, M. C., WALSH, A. R., DAVIS, B., VOLZ, E., CHENG, C., RAINEY, J. J., UZICANIN, A., GAO, H. et al. (2016). Design and methods of a social network isolation study for reducing respiratory infection transmission: The ex-flu cluster randomized trial. *Epidemics* **15**, 38–55.
- ANDERSSON, H. & BRITTON, T. (2012). *Stochastic epidemic models and their statistical analysis*, vol. 151. Springer Science & Business Media.
- ANDRIEU, C., DOUCET, A. & HOLENSTEIN, R. (2010). Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**, 269–342.
- ARNOLD, T. B. & TIBSHIRANI, R. J. (2016). Efficient Implementations of the Generalized Lasso Dual Path Algorithm. *Journal of Computational and Graphical Statistics* **25**, 1–27. Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/10618600.2015.1008638>.
- BALCAN, D., COLIZZA, V., GONÇALVES, B., HU, H., RAMASCO, J. J. & VESPIGNANI, A. (2009). Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the National Academy of Sciences* **106**, 21484–21489. Publisher: Proceedings of the National Academy of Sciences.
- BALL, F., BRITTON, T., PARDOUX, E., LARÉDO, C., SIRL, D. & TRAN, V. (2019). *Stochastic Epidemic Models with Inference*. Lecture Notes in Mathematics. Springer International Publishing.
- BICKEL, P., CHOI, D., CHANG, X. & ZHANG, H. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *The Annals of Statistics* **41**, 1922–1943. Publisher: Institute of Mathematical Statistics.
- BOWMAN, C., GUMEL, A., VAN DEN DRIESCHE, P., WU, J. & ZHU, H. (2005). A mathematical model for assessing control strategies against west nile virus. *Bulletin of mathematical biology* **67**, 1107–1133.
- BOYD, S., PARIKH, N. & CHU, E. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc.
- BRITTON, T. (2010). Stochastic epidemic models: a survey. *Mathematical biosciences* **225**, 24–35.
- BRITTON, T. & GIARDINA, F. (2016). Introduction to statistical inference for infectious diseases. *Journal de la société française de statistique* **157**, 53–70.
- BU, F., AIELLO, A. E., XU, J. & VOLFOVSKY, A. (2022). Likelihood-based inference for partially observed epidemics on dynamic networks. *Journal of the American Statistical Association* **117**, 510–526.
- CARCIONE, J. M., SANTOS, J. E., BAGAINI, C. & BA, J. (2020). A simulation of a covid-19 epidemic based on a deterministic seir model. *Frontiers in public health* **8**, 230.




Fig. 11: ℓ_1 error for various graph topologies as a function of β, n_{steps} . The healing rate is here fixed to $\gamma = 0.1$. Results are averaged over 100 independent experiments, and error bars denote interquartile ranges.




Fig. 12: ℓ_2 error for various graph topologies as a function of β , n_{steps} . The healing rate is here fixed to $\gamma = 0.1$. Results are averaged over 100 independent experiments, and error bars denote interquartile ranges.




Fig. 13: Number of infections for graphs with various topologies as a function of β , n_{steps} . The healing rate is here fixed to $\gamma = 0.1$. Results are averaged over 100 independent experiments, and error bars denote interquartile ranges.

- CELISSE, A., DAUDIN, J.-J. & PIERRE, L. (2012). Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electronic Journal of Statistics* **6**, 1847–1899. Publisher: Institute of Mathematical Statistics and Bernoulli Society.
- CHAKRABARTI, D., WANG, Y., WANG, C., LESKOVEC, J. & FALOUTSOS, C. (2008). Epidemic thresholds in real networks. *ACM Transactions on Information and System Security (TISSEC)* **10**, 1–26.
- CHAKRABORTY, T., GHOSH, I., MAHAJAN, T. & ARORA, T. (2022). Nowcasting of covid-19 confirmed cases: Foundations, trends, and challenges. *Modeling, Control and Drug Development for COVID-19 Outbreak Prevention*, 1023–1064.
- DALALYAN, A. S., HEBIRI, M. & LEDERER, J. (2017). On the prediction performance of the Lasso. *Bernoulli* **23**, 552–581. Publisher: [Bernoulli Society for Mathematical Statistics and Probability, International Statistical Institute (ISI)].
- DANON, L., FORD, A. P., HOUSE, T., JEWELL, C. P., KEELING, M. J., ROBERTS, G. O., ROSS, J. V., VERNON, M. C. et al. (2011). Networks and the epidemiology of infectious disease. *Interdisciplinary perspectives on infectious diseases* **2011**.
- DESAI, A. N., KRAEMER, M. U., BHATIA, S., CORI, A., NOUVELLET, P., HERRINGER, M., COHN, E. L., CARRION, M., BROWNSTEIN, J. S., MADOFF, L. C. et al. (2019). Real-time epidemic forecasting: challenges and opportunities. *Health security* **17**, 268–275.

- DIAMOND, S. & BOYD, S. (2016). Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* **17**, 1–5.
- DRAIEF, M. & MASSOULIÉ, L. (2009). *Epidemics and Rumours in Complex Networks*. London Mathematical Society Lecture Note Series. Cambridge University Press.
- DU, Z., XU, X., WU, Y., WANG, L., COWLING, B. J. & MEYERS, L. A. (2020). Serial interval of covid-19 among publicly reported confirmed cases. *Emerging infectious diseases* **26**, 1341.
- EIKENBERRY, S. E., MANCUSO, M., IBOI, E., PHAN, T., EIKENBERRY, K., KUANG, Y., KOSTELICH, E. & GUMEL, A. B. (2020). To mask or not to mask: Modeling the potential for face mask use by the general public to curtail the covid-19 pandemic. *Infectious disease modelling* **5**, 293–308.
- GANESH, A., MASSOULIÉ, L. & TOWSLEY, D. (2005). The effect of network topology on the spread of epidemics. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 2. IEEE.
- GAO, C., LU, Y. & ZHOU, H. H. (2015). Rate-Optimal Graphon Estimation. *The Annals of Statistics* **43**, 2624–2652. Publisher: Institute of Mathematical Statistics.
- GAUCHER, S. & KLOPP, O. (2021). Maximum likelihood estimation of sparse networks with missing observations. *Journal of Statistical Planning and Inference* **215**, 299–329.
- GINÉ, E. & NICKL, R. (2016). Mathematical foundations of infinite-dimensional statistical methods. *Cambridge University Press*.
- GRAHAM, F. (2020). Daily briefing: Why the coronavirus death toll in wuhan, china, has just jumped by 50. *Nature (Lond.)*.
- GÓMEZ, S., ARENAS, A., BORGE-HOLTHOEFER, J., MELONI, S. & MORENO, Y. (2010). Discrete-time Markov chain approach to contact-based disease spreading in complex networks. *Europhysics Letters* **89**, 38009.
- HACKL, J. & DUBERNET, T. (2019). Epidemic spreading in urban areas using agent-based transportation models. *Future internet* **11**, 92.
- HALLAC, D., PARK, Y., BOYD, S. & LESKOVEC, J. (2017). Network inference via the time-varying graphical lasso. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- HOLME, P. & SARAMÄKI, J. (2012). Temporal networks. *Physics Reports* **519**, 97–125.
- HU, B., GONZALES, J. L. & GUBBINS, S. (2017). Bayesian inference of epidemiological parameters from transmission experiments. *Scientific Reports* **7**, 16774.
- HUANG, J., MORSOMME, R., DUNSON, D. & XU, J. (2022). Detecting changes in the transmission rate of a stochastic epidemic model. *arXiv preprint arXiv:2211.14691*.
- HÜTTER, J.-C. & RIGOLLET, P. (2016). Optimal rates for total variation denoising. In *Conference on Learning Theory*. PMLR. ISSN: 1938-7228.
- IONIDES, E. L., BHADRA, A., ATCHADÉ, Y. & KING, A. (2011). Iterated filtering. *The Annals of Statistics* **39**, 1776–1802.
- JEWELL, C. P., KYPRAIOS, T., NEAL, P. & ROBERTS, G. O. (2009). Bayesian analysis for emerging infectious diseases. *Bayesian Analysis* **4**, 465–496. Publisher: International Society for Bayesian Analysis.
- JOHNS HOPKINS CRC, . (2024). New covid-19 cases worldwide. Available online: <https://coronavirus.jhu.edu/data/new-cases>. Accessed on 29 April 2024.
- KEELING, M. (2005). The implications of network structure for epidemic dynamics. *Theoretical population biology* **67**, 1–8.
- KEELING, M. J. & EAMES, K. T. (2005a). Networks and epidemic models. *Journal of the royal society interface* **2**, 295–307.
- KEELING, M. J. & EAMES, K. T. (2005b). Networks and epidemic models. *Journal of The Royal Society Interface* **2**, 295–307. Publisher: Royal Society.
- KENAH, E. & ROBINS, J. M. (2007). Network-based analysis of stochastic SIR epidemic models with random and proportionate mixing. *Journal of theoretical biology* **249**, 706–722.
- KERMACK, W. O. & MCKENDRICK, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character* **115**, 700–721.
- KIM, K., YOO, S., LEE, S., LEE, D. & LEE, K.-H. (2021). Network analysis to identify the risk of epidemic spreading. *Applied Sciences* **11**, 2997.
- KING, A. A., NGUYEN, D. & IONIDES, E. L. (2016). Statistical inference for partially observed markov processes via the r package pomp. *Journal of Statistical Software* **69**, 1–43.
- KLOPP, O., TSYBAKOV, A. B. & VERZELEN, N. (2017). Oracle Inequalities for Network Models and Sparse Graphon Estimation. *The Annals of Statistics* **45**, 316–354. Publisher: Institute of Mathematical Statistics.
- KLOPP, O. & VERZELEN, N. (2019). Optimal graphon estimation in cut distance. *Probability Theory and Related Fields* **174**, 1033–1090.
- KOLTCHINSKII, V. (2011). *Oracle inequalities in empirical risk minimization and sparse recovery problems*, vol. 2033 of *Lecture Notes in Mathematics*. Springer, Heidelberg. Lectures from the 38th Probability Summer School held in Saint-Flour, 2008, École d’Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School].
- KUMAR, A., CHOI, T.-M., WAMBA, S. F., GUPTA, S. & TAN, K. H. (2021). Infection vulnerability stratification risk modelling of covid-19 data: a deterministic seir epidemic model analysis. *Annals of Operations Research*, 1–27.

- LEDoux, M. (2001). *The concentration of measure phenomenon*, vol. 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI.
- LIPSITCH, M., DONNELLY, C. A., FRASER, C., BLAKE, I. M., CORI, A., DORIGATTI, I., FERGUSON, N. M., GARSKE, T., MILLS, H. L., RILEY, S. et al. (2015). Potential biases in estimating absolute and relative case-fatality risks during outbreaks. *PLoS neglected tropical diseases* **9**, e0003846.
- LLOYD-SMITH, J. O., SCHREIBER, S. J., KOPP, P. E. & GETZ, W. M. (2005). Superspreading and the effect of individual variation on disease emergence. *Nature* **438**, 355–359. Number: 7066 Publisher: Nature Publishing Group.
- MAMMEN, E. & GEER, S. v. d. (1997). Locally adaptive regression splines. *The Annals of Statistics* **25**, 387–413. Publisher: Institute of Mathematical Statistics.
- MANZO, G. (2020). Complex social networks are missing in the dominant covid-19 epidemic models. *Sociologica* **14**, 31–49.
- MCGOUGH, S. F., JOHANSSON, M. A., LIPSITCH, M. & MENZIES, N. A. (2020). Nowcasting by bayesian smoothing: A flexible, generalizable model for real-time epidemic tracking. *PLoS computational biology* **16**, e1007735.
- MILLER, J. C. & VOLZ, E. M. (2013). Incorporating Disease and Population Structure into Models of SIR Disease in Contact Networks. *PLOS ONE* **8**, e69162. Publisher: Public Library of Science.
- MORENO, Y., PASTOR-SATORRAS, R. & VESPIGNANI, A. (2002). Epidemic outbreaks in complex heterogeneous networks. *The European Physical Journal B-Condensed Matter and Complex Systems* **26**, 521–529.
- MORSOMME, R. & XU, J. (2022). Exact inference for stochastic epidemic models via uniformly ergodic block sampling. *arXiv preprint arXiv:2201.09722*.
- MUBARAK, M., BERNEBURG, J. & NOWZARI, C. (2021). Stochastic vs. deterministic modeling for the spread of covid-19 in small networks. In *2021 American Control Conference (ACC)*. IEEE.
- NARCI, R., DELATTRE, M., LARÉDO, C. & VERGU, E. (2021). Inference for partially observed epidemic dynamics guided by kalman filtering techniques. *Computational Statistics and Data Analysis* **164**, 107319.
- NEEDELL, D. & WARD, R. (2013). Near-Optimal Compressed Sensing Guarantees for Total Variation Minimization. *IEEE Transactions on Image Processing* **22**, 3941–3949. Conference Name: IEEE Transactions on Image Processing.
- NEWMAN, M. E. (2002a). Spread of epidemic disease on networks. *Physical review E* **66**, 016128.
- NEWMAN, M. E. J. (2002b). Spread of epidemic disease on networks. *Physical Review E* **66**, 016128. Publisher: American Physical Society.
- OTTAVIANO, S., DE PELLEGRINI, F., BONACCORSI, S. & VAN MIEGHEM, P. (2018). Optimal curing policy for epidemic spreading over a community network with heterogeneous population. *Journal of Complex Networks* **6**, 800–829.
- PADILLA, O. H. M., SHARPBACK, J., SCOTT, J. G. & TIBSHIRANI, R. J. (2018). The DFS Fused Lasso: Linear-Time Denoising over General Graphs. *Journal of Machine Learning Research* **18**, 1–36.
- PARÉ, P. E., BECK, C. L. & BAŞAR, T. (2020a). Modeling, estimation, and analysis of epidemics over networks: An overview. *Annual Reviews in Control* **50**, 345–360.
- PARÉ, P. E., LIU, J., BECK, C. L., KIRWAN, B. E. & BAŞAR, T. (2020b). Analysis, Estimation, and Validation of Discrete-Time Epidemic Processes. *IEEE Transactions on Control Systems Technology* **28**, 79–93. Conference Name: IEEE Transactions on Control Systems Technology.
- PASTOR-SATORRAS, R., CASTELLANO, C., VAN MIEGHEM, P. & VESPIGNANI, A. (2015a). Epidemic processes in complex networks. *Reviews of modern physics* **87**, 925.
- PASTOR-SATORRAS, R., CASTELLANO, C., VAN MIEGHEM, P. & VESPIGNANI, A. (2015b). Epidemic processes in complex networks. *Rev. Mod. Phys.* **87**, 925–979.
- PASTOR-SATORRAS, R. & VESPIGNANI, A. (2001). Epidemic Spreading in Scale-Free Networks. *Physical Review Letters* **86**, 3200. Publisher: American Physical Society.
- PELLIS, L., BALL, F., BANSAL, S., EAMES, K., HOUSE, T., ISHAM, V. & TRAPMAN, P. (2015). Eight challenges for network epidemic models. *Epidemics* **10**, 58–62.
- PEREZ, L. & DRAGICEVIC, S. (2009). An agent-based approach for modeling dynamics of contagious disease spread. *International journal of health geographics* **8**, 1–17.
- ROCHA, L. E. C., LILJEROS, F. & HOLME, P. (2011). Simulated Epidemics in an Empirical Spatiotemporal Network of 50,185 Sexual Contacts. *PLoS Computational Biology* **7**, e1001109.
- ROSENFELD, R. & TIBSHIRANI, R. J. (2021). Epidemic tracking and forecasting: Lessons learned from a tumultuous year. *Proceedings of the National Academy of Sciences* **118**, e2111456118.
- ROSSI, R. A. & AHMED, N. K. (2015). The network data repository with interactive graph analytics and visualization. In *AAAI*.
- RUDIN, L. I., OSHER, S. & FATEMI, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60**, 259–268.
- SANCHEZ, F., CALVO, J. G., MERY, G., GARCÍA, Y. E., VÁSQUEZ, P., BARBOZA, L. A., PÉREZ, M. D. & RIVAS, T. (2022). A multilayer network model of covid-19: Implications in public health policy in costa rica. *Epidemics* **39**, 100577.
- SHIRLEY, M. D. & RUSHTON, S. P. (2005). The impacts of network topology on disease spread. *Ecological Complexity* **2**, 287–299.
- SIEGENFELD, A. F., TALEB, N. N. & BAR-YAM, Y. (2020). What models can and cannot tell us about covid-19. *Proceedings of the National Academy of Sciences* **117**, 16092–16095.

- SIETTOS, C., ANASTASSOPOULOU, C., RUSSO, L., GRIGORAS, C. & MYLONAKIS, E. (2015). Modeling the 2014 ebola virus epidemic—agent-based simulations, temporal analysis and future predictions for liberia and sierra leone. *PLoS currents* **7**.
- SISSON, S. A., FAN, Y. & TANAKA, M. M. (2007). Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences* **104**, 1760–1765.
- SPRICER, K. & BRITTON, T. (2019). An sir epidemic on a weighted network. *Network Science* **7**, 556–580.
- SUN, D., TOH, K.-C. & YUAN, Y. (2021). Convex clustering: Model, theoretical guarantee and efficient algorithm. *The Journal of Machine Learning Research* **22**, 427–458.
- TALAGRAND, M. (1996). New concentration inequalities in product spaces. *Inventiones Mathematicae* **126**, 505–563.
- TEWELDEMEDHIN, E., MARWALA, T. & MUELLER, C. (2004). Agent-based modelling: a case study in hiv epidemic. In *Fourth International Conference on Hybrid Intelligent Systems (HIS'04)*. IEEE.
- TIBSHIRANI, R. J. & TAYLOR, J. (2011). The solution path of the generalized lasso. *The annals of statistics* **39**, 1335–1371.
- TOLLES, J. & LUONG, T. (2020). Modeling epidemics with compartmental models. *Jama* **323**, 2515–2516.
- TONI, T., WELCH, D., STRELKOWA, N., IPSEN, A. & STUMPF, M. P. H. (2009). Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society, Interface* **6**, 187–202.
- VOLZ, E. & MEYERS, L. A. (2007). Susceptible–infected–recovered epidemics in dynamic contact networks. *Proceedings of the Royal Society B: Biological Sciences* **274**, 2925–2933.
- VRABAC, D., PARÉ, P. E., SANDBERG, H. & JOHANSSON, K. H. (2020). Overcoming Challenges for Estimating Virus Spread Dynamics from Data. In *2020 54th Annual Conference on Information Sciences and Systems (CISS)*.
- WAN, X., LIU, J., CHEUNG, W. K. & TONG, T. (2014). Inferring epidemic network topology from surveillance data. *PLoS One* **9**, e100661.
- WANG, H., LI, Q., D'AGOSTINO, G., HAVLIN, S., STANLEY, H. E. & VAN MIEGHEM, P. (2013). Effect of the interconnected network structure on the epidemic threshold. *Physical Review E* **88**, 022801.
- WANG, Y.-X., SHARPNACK, J., SMOLA, A. J. & TIBSHIRANI, R. J. (2016). Trend Filtering on Graphs. *Journal of Machine Learning Research* **17**, 1–41.
- WORLD HEALTH ORGANIZATION, . (2024). Daily covid-19 cases and deaths. Data processed by Our World in Data. Available online: <https://ourworldindata.org/grapher/daily-covid-cases-deaths-7-day-ra>. Accessed: 29 April 2024.
- WU, J. T., LEUNG, K., LAM, T. T., NI, M. Y., WONG, C. K., PEIRIS, J. M. & LEUNG, G. M. (2021). Nowcasting epidemics of novel pathogens: lessons from covid-19. *Nature medicine* **27**, 388–395.
- ZHANG, J. & MOURA, J. M. F. (2014). Diffusion in Social Networks as SIS Epidemics: Beyond Full Mixing and Complete Graphs. *IEEE Journal of Selected Topics in Signal Processing* **8**, 537–551. Conference Name: IEEE Journal of Selected Topics in Signal Processing.