

---

# Prediction-Centric Learning of Independent Cascade Dynamics from Partial Observations

---

Mateusz Wilinski<sup>1</sup> Andrey Y. Lokhov<sup>1</sup>

## Abstract

Spreading processes play an increasingly important role in modeling for diffusion networks, information propagation, marketing and opinion setting. We address the problem of learning of a spreading model such that the predictions generated from this model are accurate and could be subsequently used for the optimization, and control of diffusion dynamics. We focus on a challenging setting where full observations of the dynamics are not available, and standard approaches such as maximum likelihood quickly become intractable for large network instances. We introduce a computationally efficient algorithm, based on a scalable dynamic message-passing approach, which is able to learn parameters of the effective spreading model given only limited information on the activation times of nodes in the network. The popular Independent Cascade model is used to illustrate our approach. We show that tractable inference from the learned model generates a better prediction of marginal probabilities compared to the original model. We develop a systematic procedure for learning a mixture of models which further improves the prediction quality.

## 1. Introduction

Spreading processes on networks have seen an increasing interest in the domains related to modeling of infectious diseases, regulatory networks, marketing and opinion dynamics. In a realistic setting, the model and parameters, as well as propagation paths are unknown, and one is left at most with several observed diffusion traces in a form of nodes' activation times. Moreover, even this information is not fully available since the time-stamps of node activations may not be available from each node in the network.

---

<sup>1</sup>Theoretical Division, Los Alamos National Laboratory, Los Alamos, USA. Correspondence to: Mateusz Wilinski <mateusz@lanl.gov>, Andrey Y. Lokhov <lokhov@lanl.gov>.

This motivates the development of efficient algorithms that can infer the effective spreading models from incomplete information.

In the recent years, the problem of diffusion model recovery from full observations has been extensively addressed through a series of works focusing on heuristic and exact algorithms, largely based on maximum likelihood type approaches (Netrapalli & Sanghavi, 2012; Gomez-Rodriguez et al., 2012; Gripon & Rabbat, 2013; Abrahao et al., 2013; Pouget-Abadie & Horel, 2015; Braunstein et al., 2019). These methods allow one to recover the parameters of the model, as well as the structure of the diffusion graph. Comparably, much less has been done on the side of parameter recovery in the case of partial observations. This is primarily due to the hardness of the estimation in this case, where maximum likelihood marginalized over the hidden observations quickly becomes computationally prohibitive with the number of hidden nodes (Lokhov, 2016). The work (Amin et al., 2014) primarily considered the case of tree graphs, and studied the network reconstruction problem for the case where only the initial and the final states of the dynamics are observed. This setting was further developed in (Supeesun & Fakcharoenphol, 2017) and very recently in (Han et al., 2020) for more general graphs. In a different setting, (Sefer & Kingsford, 2015) addressed the structure learning problem in the case of missing information in time through relaxation optimization techniques under the assumption that the full probabilistic trace for each vertex is available. The case of noise in time data was analysed in (Hoffmann & Caramanis, 2019) and (Trouleau et al., 2019), where the latter assumed a Hawkes process driving the dynamics. An algorithm for learning model parameters based on estimating marginal probabilities through a dynamic message-passing approach has been suggested in (Lokhov, 2016). Although this method provides a low estimation error and is computationally efficient compared to the marginalized likelihood, it still suffers from a computational complexity that scales quadratically in the system size for each step of the gradient descent, which limits the applicability of the method to small systems with an order of dozens of nodes only.

Much of the work above focuses on an accurate estimation of model parameters. However, in many applications the

main reason for learning a model from data consists in its subsequent use for inference, optimization, or control tasks (Kempe et al., 2003; Nowzari et al., 2016). In this setting, sometimes referred to as “learning for inference”, or “inferning” (Heinemann & Globerson, 2014; Bresler & Karzand, 2020; Boix-Adsera et al., 2021)<sup>1</sup>, the main interest of the learning procedure is not focused on an accurate estimation of the true model parameters, but rather on the construction of an effective model such that the predictions generated from this model are accurate. The prediction task that is natural for the spreading dynamics is that of prediction of marginal probabilities of activation for each individual node that collectively quantify the expected size of the diffusion spread (Kempe et al., 2003). One way to address this problem was proposed in a line of works that address the problem of direct estimation of influence functions from samples (Du et al., 2014; Narasimhan et al., 2015; He et al., 2016; Makar et al., 2018). However, one could be interested in constructing an explicitly parametrized influence function with an associated inference technique that can be used for other tasks beyond prediction of spread. Similarly to how belief propagation is used for approximate inference in graphical models (Wainwright & Jordan, 2008), the task of an efficient approximate estimation of marginal probabilities is solved in the spreading context by dynamic message-passing (DMP) equations that avoid the computational overhead of repetitive Monte-Carlo sampling (Lokhov & Saad, 2019). Moreover, explicit equations that compute marginals has proven useful in inference tasks such as localization of the origin of the spreading process (Lokhov et al., 2014), or optimization tasks such as control of spreading through resource allocation (Lokhov & Saad, 2017). Therefore, it is natural to consider a learning procedure that would construct an effective model for which approximate inference techniques would yield an accurate prediction of the marginals.

In this paper, we make two contributions to the problems above. First, we revisit the objective function introduced in (Lokhov, 2016), and develop an exact method that computes its gradient in a time linear with respect to the dimension of the problem, thus overcoming the computational limitation of the DMPrec algorithm presented in that work. This enables application of model learning from partial observations to large realistic size networks. We also provide an interpretation of this objective function as a Kullback-Leibler divergence between empirical and model marginal distributions. This leads to the second major contribution of this work: considering dynamic message-passing as an approximate inference method, we show that inference from the effective model learned with our method generates a better prediction of marginal probabilities than inference

from the original model. This observation is in agreement with previous work in learning of graphical models (Wainwright, 2006; Heinemann & Globerson, 2014) that showed that an inconsistent parameter estimation that takes approximate inference (such as belief propagation) into account is provably beneficial for the prediction task. We capitalize on this observation and develop a novel procedure for learning a mixture of spreading models on several replicas of the graph which further improves prediction quality of marginal probabilities. Although our approach is general and can be adapted to a large class of stochastic spreading models, we illustrate our approach on the popular Independent Cascade (IC) model (Goldenberg et al., 2001; Kempe et al., 2003).

## 2. Model

We define the diffusion network as  $G = (V, E)$ , where  $V$  is the set of nodes and  $E$  is the set of edges. In the IC model, each node can be in either of two states: active or inactive. Node  $i$  activated at time  $t$  has a single chance to activate its inactive neighbor  $j$  at time  $t + 1$  with probability  $\alpha_{ij}$  associated with edge  $(i, j) \in E$ . A cascade starts with a set of initially active nodes and continues until the dynamics dies out or at most for a predefined number of steps  $T$ . Since each node can be activated no more than once, every cascade  $c$  can be fully described by a set of activation times  $\{\tau_i^c\}_{i \in V}$ . Additionally, if node  $i$  is not activated during the dynamics, we assign  $T$  as its activation time by definition. We assume that we observe  $M$  independent cascades with potentially different initial seeds. In general, there can be one or several seeds at initial time, or even stochastic initial conditions, where node  $i$  is independently activated with probability  $p_i^c(0)$  at initial time in each cascade  $c$ .

In practice, full observation of cascades is rarely available, and hence we split the activation times into observed and unobserved parts. Our goal is to learn an effective dynamic model from the set of observed activation times. We will be interested in two tasks: (i) learning model parameters, i.e. estimating  $\{\alpha_{ij}\}_{(i,j) \in E}$  so that they are close to the parameters of the original model used to generate data, denoted as  $\{\alpha_{ij}^*\}_{(i,j) \in E}$ ; and (ii) prediction of marginal probabilities  $\{p_i(t)\}_{i \in V}$ , where  $p_i^c(t)$  is the probability that node  $i$  is active at time  $t$ , so that they match the marginal probabilities of the true model denoted as  $\{p_i^{c*}(t)\}_{i \in V}$ .

## 3. Methods

In this section, we first discuss the dynamic message-passing approximate inference algorithm for estimating marginal probabilities for the IC model, and then introduce our algorithm for learning model parameters from data.

<sup>1</sup>See also the ICML 2013 workshop on “Inferning: Interactions between Inference and Learning”: <http://inferning.cs.umass.edu/2012>.

### 3.1. Dynamic Message Passing Equations

Computation of marginal probabilities of activation and of the influence function – sum of marginals at final time  $T$ , which quantifies the expected spread – is #P-hard in the IC model (Chen et al., 2010; Shapiro & Delgado-Eckert, 2012). In practice, one could estimate the marginal probabilities using Monte-Carlo sampling, however the required number of effective runs typically scales at least linearly with the size of the system (Kempe et al., 2003; Chen et al., 2009; Du et al., 2013; Cohen et al., 2014; Lucier et al., 2015; Nguyen et al., 2017). Here, we present dynamic message-passing (DMP) equations, a low complexity approximate inference algorithm that does not rely on sampling, and allows for an efficient estimation of marginal probabilities that is exact on tree graphs and asymptotically exact on random graphs.

DMP equations approximate  $\{p_i^c(t)\}_{i \in V}$  – the set of marginal probabilities that each node  $i \in V$  is active at time  $t$  under a given initial condition in the cascade  $c$ . The computation is performed with additional quantities called messages  $p_{j \rightarrow i}^c(t)$ , which have a meaning of probabilities of node  $j$  being active at time  $t$  on an auxiliary graph where  $i$  has been removed. These probabilities can be calculated by solving a following system of equations:

$$p_i^c(t) = 1 - (1 - \bar{p}_i^c) \prod_{k \in \partial i} (1 - \alpha_{ki} \cdot p_{k \rightarrow i}^c(t-1)), \quad (1)$$

$$p_{j \rightarrow i}^c(t) = 1 - (1 - \bar{p}_j^c) \prod_{k \in \partial j \setminus i} (1 - \alpha_{kj} \cdot p_{k \rightarrow j}^c(t-1)), \quad (2)$$

where  $\bar{p}_i^c = p_i^c(0)$  is the initial probability of being active for node  $i$  under initial conditions of cascade  $c$ . Properties of these equations have been extensively studied in (Lokhov & Saad, 2019), and are related to the DMP equations for the popular SIR model in the limit of deterministic recovery (Karrer & Newman, 2010; Lokhov et al., 2015), while the large-time form of these equations have been previously derived in (Abbe et al., 2017; Burkholz & Quackenbush, 2019). Notice that unlike belief propagation, DMP equations do not need to be iterated until convergence and estimate marginals on any graph. Although exactness of DMP predictions are guaranteed on tree graphs only, they often provide accurate approximations on sparse but loopy graphs. Moreover, for general graphs, the solution of above equations provides an upper bound for marginal probabilities (Lokhov & Saad, 2019). Note that for given set of parameters, the solution depends only on the initial state  $\bar{p}_i^c$ .

### 3.2. Constrained Optimization Based Learning Framework

Next, we explain how DMP equations combined with a Lagrangian formulation of the learning problem can be used to estimate model parameters  $\{\alpha_{ij}^*\}_{(i,j) \in E}$  even in the case

of partial observations. Our approach is based on the maximization of the following objective function:

$$\mathcal{O} = \sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{O}} \log \mu_i^c(\tau_i^c), \quad (3)$$

where  $\mathcal{C}$  is the set of cascades, and  $\mu_i^c(\tau_i^c)$  are marginal probabilities that node  $i$  in a cascade  $c$  is activated precisely at time  $\tau_i^c$ . Under the conventions introduced above, these quantities can be approximately expressed through the marginals estimated with the DMP equations as  $\mu_i^c(\tau_i^c) = p_i^c(\tau_i^c) \cdot \mathbb{1}_{(\tau_i^c < T)} - p_i^c(\tau_i^c - 1) \cdot \mathbb{1}_{(\tau_i^c > 0)} + \mathbb{1}_{(\tau_i^c = T)}$ . All cascades can be divided into classes according to their initial conditions. Let us denote the set of such classes as  $\mathcal{S}$ . Notice that for given parameters, the number of different sets of messages and marginals is of the same size as  $\mathcal{S}$ , because they only differ by initial conditions. As a result, we can rewrite the objective as follows:

$$\mathcal{O} = \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{O}} \sum_{\tau_i^s} m^{\tau_i^s} \log \mu_i^s(\tau_i^s), \quad (4)$$

where  $m^{\tau_i^s}$  is the number of cascades in class  $s$  for which node  $i$  was activated at time  $\tau_i^s$ , computed from available samples. The objective (4) is related to the Kullback-Leibler (KL) divergence between empirical and model marginal distributions (computed with DMP) for the observed nodes. The key premise of this objective function that enables reconstruction of model parameters even in the case of partial observations consists in the fact that marginal probabilities on observed nodes depend on all model parameters, including on those that are adjacent to hidden nodes. This objective has been originally proposed in (Lokhov, 2016) with a different interpretation, with its form inspired by a mean-field like approximation to the full likelihood function. This work also included a proof of the asymptotic consistency of DMP-based recovery of model parameters using the objective (4) on tree graphs. The DMPrec algorithm proposed in (Lokhov, 2016) consists in estimating the gradient of the objective via a direct derivative of the DMP equations. This leads to a single optimisation step complexity  $O(|E|^2 TM)$ , where  $M$  is the number of cascades. Here, we propose a more efficient algorithm that computes the exact gradient of the objective function in time  $O(|E| TM)$  at each step of the optimization, which will allow to apply it to large-sized networks.

We express the model estimation problem as a standard Lagrangian formulation for the constrained optimization problem, with the objective function (4) and the constraints expressing the fact that the marginals are estimated using DMP equations:

$$\mathcal{L} = \underbrace{\mathcal{O}}_{\text{objective}} + \underbrace{\mathcal{C}}_{\text{constraints}}. \quad (5)$$

The constraints are given by DMP equations (1)-(2) re-weighted by Lagrange multipliers, where  $\lambda_i^s(t)$  and  $\lambda_{i \rightarrow j}^s(t)$  are Lagrange multipliers corresponding respectively to equations (1) and (2).

The resulting expressions for all the quantities can be found by differentiating the Lagrangian. Differentiation over marginals yields expressions for  $\lambda_i^s(t)$  multipliers for all sources  $s$  and times  $t$ . Similarly, the  $\lambda_{i \rightarrow j}^s(t)$  multipliers can be obtained from the derivatives over messages, which give rise to message-passing equations in the dual space of Lagrange multipliers. Finally, derivatives over parameters  $\alpha_{ij}$  take a simple form:

$$\frac{\partial \mathcal{L}}{\partial \alpha_{ij}} = \frac{-1}{\alpha_{ij}} \sum_{s \in S} \sum_{t=0}^{T-1} (\lambda_{i \rightarrow j}^s(t) p_{i \rightarrow j}^s(t) + \lambda_{j \rightarrow i}^s(t) p_{j \rightarrow i}^s(t)), \quad (6)$$

for non-zero parameters  $\alpha_{ij}$ . Now, we can use gradient components (6) to update parameters  $\alpha_{ij}$ :

$$\alpha_{ij} \leftarrow \alpha_{ij} + \varepsilon \cdot \frac{\partial \mathcal{L}}{\partial \alpha_{ij}}, \quad (7)$$

where  $\varepsilon$  is a learning rate. The resulting algorithm can be expressed in a form of forward-backward iterations: (i) Start with initial parameter guess  $\alpha_{ij}$ ; (ii) Calculate marginals and messages using forward DMP equations (1)-(2); (iii) Use estimated marginals and messages to compute Lagrange multipliers through the ‘‘backward’’ equations; (iv) Use equations (6) and (7) to update parameters  $\alpha_{ij}$ ; (v) Go back to step 2, and repeat the process with updated  $\alpha_{ij}$ , until the global convergence of the algorithm. The complete form of the algorithm along with the full derivation is presented in the Supplementary Materials, section S1 (Wilinski & Lokhov, 2020), where we also discuss the selection of  $\varepsilon$  and its relation to the convergence of our learning procedure. We refer to the above algorithm as to the Scalable Learning of Independent Cascade Effective Representation (SLICER) (Wilinski & Lokhov, 2021).

Calculating marginals using DMP equations requires complexity  $O(|E|T)$  and needs to be done for each different initial point of all available cascades, resulting in the worst-case computational complexity  $O(|E|TM)$ . In the case of single-node initial conditions, this complexity is upper-bounded by  $O(|E|TN)$ , and for stochastic initial conditions is only  $O(|E|T)$ . Lagrange multipliers require the same computational complexity, since backward equations are equivalent to dynamic message-passing equations in the dual space. Finally, derivatives with respect to  $\{\alpha_{ij}\}_{(i,j) \in E}$ , as seen from equation (6), are computed with the same computational complexity. As a result, overall SLICER has a complexity  $\min [O(|E|TM), O(|E|T|S|)]$ . In the next section, we test the SLICER algorithm on synthetic data.

## 4. Estimation of Model Parameters

In this section, we focus on the problem of parameter estimation. We start by presenting a comparison with established methods, and then empirically evaluate the performance of our algorithm on a variety of topologies, including random graphs, as well as regular lattices and real-world networks with a very large number of loops. Naturally, except for the calibration comparisons, our main focus will be the case of partial observations, where maximum likelihood based approaches are not applicable.

In all the tests below, unless stated otherwise, (i) parameters  $\alpha_{ij}$  are sampled uniformly from  $[0, 1]$ ; (ii) each cascade is generated independently from the IC model with limited  $T$ , varying from 4 to 20; (iii) the source of every cascade is a single, randomly chosen, node and (iv) hidden nodes are chosen uniformly at random. A non-uniform distribution of seeds and hidden nodes, and a use of multiple sources do not qualitatively change the reconstruction trends as long as sufficient statistics of activations is available, see a related discussion along with additional experiments in the Supplementary Materials, section S2 (Wilinski & Lokhov, 2020). All points in the plots are averaged over multiple realisations of parameters  $\alpha_{ij}^*$  and network structures – where applicable, and the respective statistical confidence is presented through the error bars. We initialise the learning process with  $\alpha_{ij} = 0.5 \forall (ij) \in E$ , unless stated otherwise.

### 4.1. Comparison with Established Methods

We start with a few comparisons with existing methods to illustrate the basic properties of SLICER. DMPrec has been extensively compared to marginalized likelihood algorithms in terms of accuracy and running time in (Lokhov, 2016), and reconstructs model parameters from partial observations in an exact setting that we are considering here. Hence, it is natural to compare SLICER to this algorithm as a benchmark: from the discussion in the previous section, we expect to see an improvement in the algorithmic complexity. DMPrec has been originally introduced for a stochastic susceptible-infected model; in the Supplementary Materials, section S3 (Wilinski & Lokhov, 2020), we adapt this algorithm to the case of the IC model considered here. For illustrating the accuracy of reconstruction obtained with SLICER, we benchmark our algorithm against the maximum likelihood approach; we only provide this ‘‘calibration’’ comparison in the case of full observability, where maximum likelihood is computationally tractable (Lokhov, 2016). Precise maximum likelihood based formulation is presented in the Supplementary Materials, Section S4 (Wilinski & Lokhov, 2020). For the optimization of maximum likelihood, we use the optimization software Ipopt (Wächter & Biegler, 2006) within the Julia/JuMP modeling framework for mathematical optimization (Dunning et al., 2017).

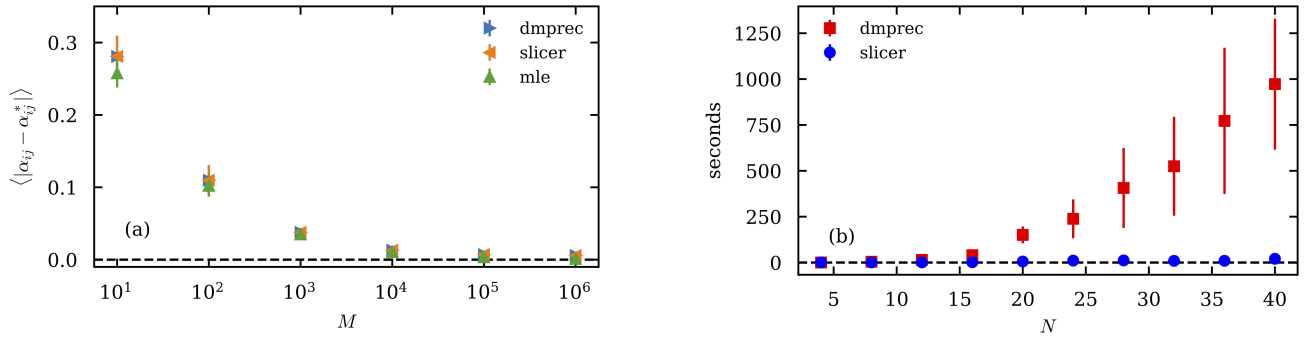


Figure 1. Comparison of accuracy and time complexity between SLICER, DMPrec and Maximum Likelihood estimators. (a) The average difference between inferred and real parameters  $\alpha_{ij}$ , as a function of number of available cascades for a 3-regular tree with  $N = 20$  nodes and cascades length  $T = 5$ . Results averaged over five different sets of parameters. (b) Run-times comparisons of a single gradient step evaluation between SLICER and DMPrec as a function of network size. Results for 3-regular tree and  $10^3$  cascades of length  $T = 5$  averaged over 200 gradient steps evaluations and five different sets of parameters.

Our experiments reveal that both SLICER and DMPrec produce identical results for parameter recovery, which is not surprising given that they optimize the same objective function. This phenomenon is illustrated in Fig. 1(a), where both algorithms exhibit exactly the same accuracy. The main difference is the time needed to evaluate a single gradient step. As shown in Fig. 1(b), a single-step gradient evaluation of DMPrec quickly reaches running time that are prohibitive for the application of the method to networks with more than several dozens of nodes. For this reason, we perform comparisons on relatively small networks with up to  $N = 40$  nodes, consistent with the maximum size of test cases considered in (Lokhov, 2016). As we show below, SLICER maintains a low computational complexity, and can be easily scaled up to large networks; additional tests on computational complexity of SLICER as a function of problem parameters are presented in the Supplementary Materials, section S1 (Wilinski & Lokhov, 2020).

Unlike maximum likelihood, SLICER optimizes the likelihood of marginal distributions (estimated with the DMP) only, which is the key to tractability in the case of partial observations. Still, as shown in Fig. 1(a), it demonstrates a comparable accuracy to maximum likelihood on tree graphs, where DMP is exact, in the case of full observations.

## 4.2. Learning Model Parameters

In order to get a flavor of typical reconstruction errors as a function of network topology, we calibrate our approach on the case of full observations. The results for this simplest setting, for different network types, are shown in Fig. 2. In particular, we test error decay on a regular tree of degree three, a scale-free tree, a random regular graph of degree three, an Erdős-Rényi network with average degree three and a square lattice. The apparent difference between left and right subplots of Fig. 2 is a result of loops affecting

the accuracy of DMP equations: When the observation time horizon is short, even though the graph is loopy, the spreading effectively occurs in a tree-like region where DMP equations are exact. As  $T$  increases, information spreads through loops that start to have an effect on the accuracy of predictions of marginals, and hence on the recovered parameters. For the adversarial case of a regular grid, the error saturates to a finite value for  $T = 10$  even for a large number of samples. For parameter estimation task, this is expected due to the presence of systematic short loops.

Let us now consider a more challenging case of reconstructing parameters when only part of the information is available. To test the accuracy of our method, we randomly pick a fraction of nodes  $\xi$  and hide their activation times throughout all cascades. Results for this setting are shown in Fig. 3. We test two types of networks, regular trees and regular random graphs, both for shorter,  $T = 5$ , and longer,  $T = 10$ , cascades. In a particular case of trees, we neglected the edges which connect the networks with unobserved leaves, since estimating their spreading parameters is impossible with any algorithm. These results show that  $\alpha_{ij}^*$  can be efficiently estimated to a good accuracy even with  $\xi = 25\%$  of unobserved nodes, as long as enough data is available. Moreover, the reconstruction quality improves with the growing observation time for these tree-like cases.

Finally, we test SLICER on loopy real-world networks. We use two web networks for numerical tests: one representing the snapshot of the structure of the Internet at the level of autonomous systems (Rossi & Ahmed, 2015), and the other obtained by a web-crawler (Boldi et al., 2004). These networks are not only loopy (46, 873 and 993, 421 triangles respectively), but also large ( $N = 22, 963$  and  $N = 121, 422$  nodes,  $M = 48, 436$  and  $M = 334, 419$  edges respectively), which allows us to test the scalability of SLICER. We use a stochastic initial condition, where instead of a

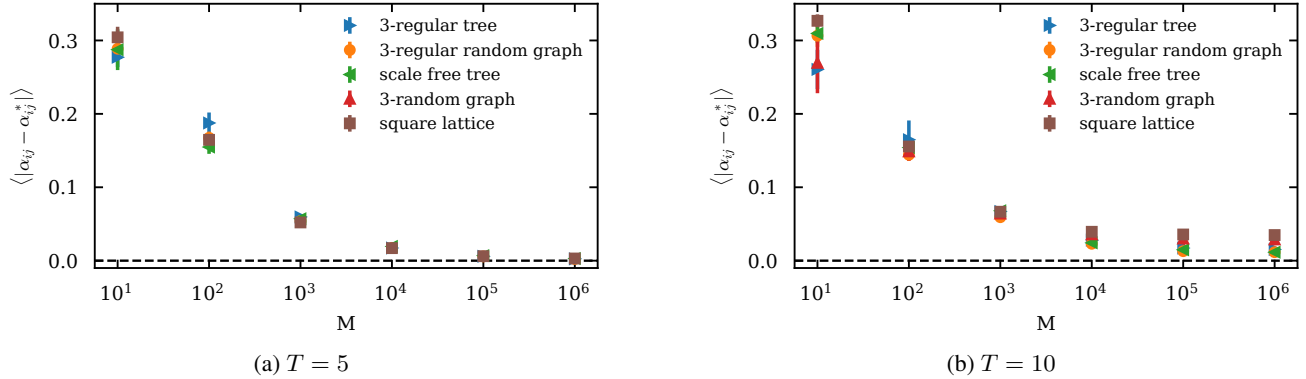


Figure 2. Average difference between inferred  $\alpha_{ij}$  and real  $\alpha_{ij}^*$  parameters in  $\ell_1$  norm, as a function of number of available cascades in the case of full observations of cascades, for two observation horizons: (a)  $T = 5$  and (b)  $T = 10$ . Each point is averaged over 5 different networks and 5 different sets of parameters  $\alpha_{ij}^*$  (sampled uniformly from  $[0, 1]$ ). All networks contain  $N = 100$  nodes.

single source, the nodes activate at time  $t = 0$  randomly with  $p_i^c(0) = \frac{10}{N}$  for each of the cascade. This means that each cascade starts with ten activated nodes on average. In the learning procedure, the stochastic initial condition has an advantage that the DMP equations need to be run only once. We chose the stochastic initial conditions for making sure that a large part of a network can be activated, thus providing information on model parameters. Due to the presence of high-degree nodes, uniform choice of initial conditions leads to a rapid full activation of the network. We hence chose the degree-dependent transmission probabilities  $\alpha_{ij}^* = [\max(k_i, k_j)]^{-\frac{1}{2}}$  that lead to a non-trivial dynamics, representing a trade-off between immediate activation of the full network, and a limited coverage of the graph. For this choice of parameters,  $T = 4$  turns out to be the time horizon beyond which DMP predicts the activation of the full network; we hence use the cascade data up to this time in our numerical experiments. In this test case, SLICER is initialized with small  $\alpha_{ij}$ . Results are shown in Fig. 4, and show a similar trend compared to the synthetic instances analyzed previously. It is however important to note that the maximum considered data size of  $10^6$  cascades represents a small statistics compared to the synthetic cases with  $N = 100$ . Indeed, the ratio between the available data ( $10^6$  cascades in both cases) and the number of parameters in real-world instances (48, 436 and 334, 419 parameters respectively) is much lower than in synthetic cases (150 to 200 parameters). Hence, we anticipate that a better accuracy can be achieved by collecting more data. Finally, structural particularities of realistic networks with interconnected hubs may render some edges non-identifiable in some regime of parameters. We did not filter our results based on the identifiability, and our error estimates may include random guesses for those links (identifiability is an interesting question but beyond the current scope). Interestingly enough, despite a large number of short loops, it appears that partial observ-

ability has a small effect on the quality of the reconstructed parameters: reconstruction quality does not seem to drop much even for  $\xi = 25\%$ . Overall, tests on real-world networks demonstrate that SLICER can scale to large networks with hundreds of thousands of nodes, a significant improvement compared to the DMPrec algorithm of (Lokhov, 2016). Learned model parameters even on loopy instances can then be used in the setting of “inferning” discussed next.

## 5. Prediction of Marginal Probabilities

As explained in the introduction, in many applications falling under the “learning for inference” setting (Heinemann & Globerson, 2014; Bresler & Karzand, 2020), we are less interested in an accurate recovery of the underlying model parameters, but instead want to learn an effective model that succeeds in correctly predicting the marginal distributions that can subsequently be used for control, optimization, or learning tasks. Here, we show that our approach naturally recovers the best effective parameters that lead to a better approximation of the marginals in the case where DMP is used as an inference method due to its attractive computational complexity comparable to a single run of the Monte-Carlo simulation (Lokhov & Saad, 2019).

### 5.1. Effective Model for Prediction of Marginals

As already mentioned, the primary source of error in the DMP equations is related to the presence of loops. Their presence may result in overestimating parameters to account for DMP neglecting information on some spreading paths. To study this effect in more details, we focused on the most challenging case of a square lattice with many short loops, in an adversarial case of very long cascades with  $T = 20$ . As shown in Fig. 5(a), this represents a challenging setting in terms of prediction of marginals using the DMP-based

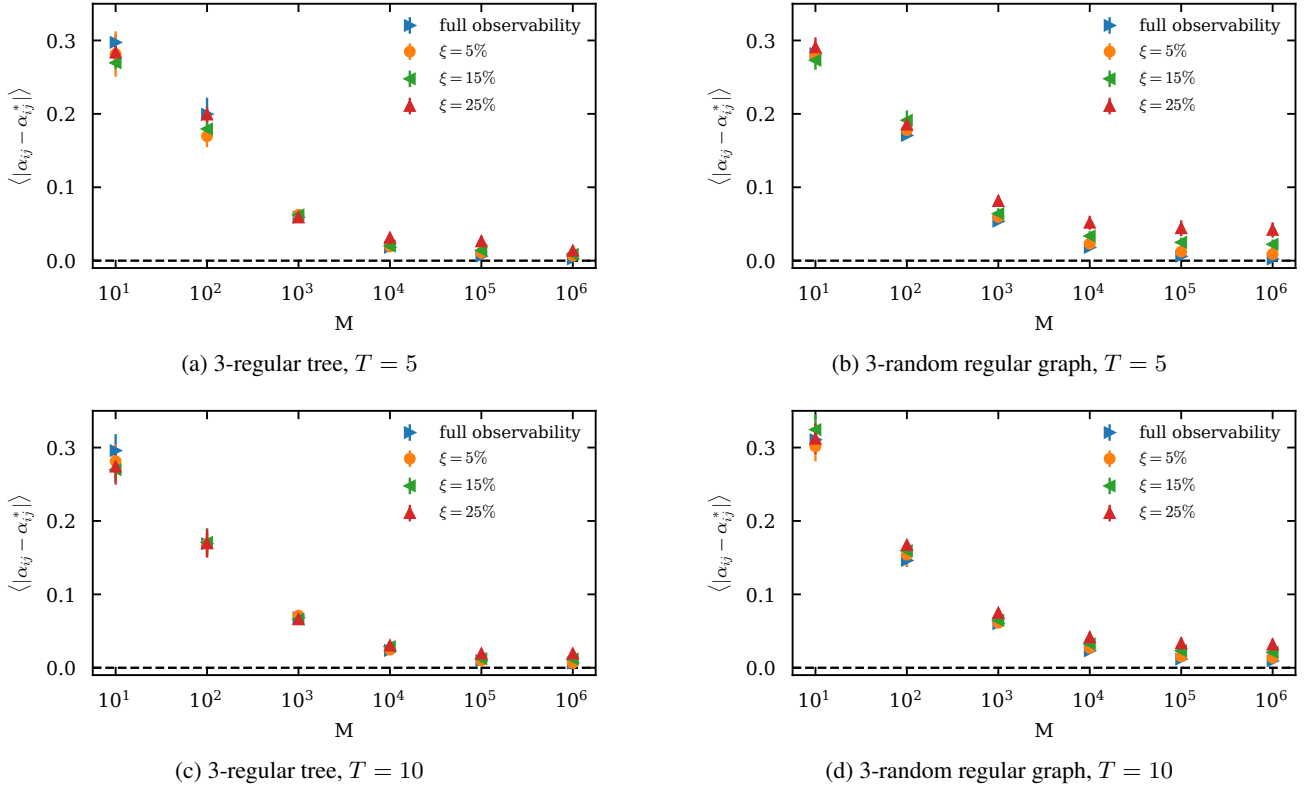


Figure 3. Average difference between inferred  $\alpha_{ij}$  and real  $\alpha_{ij}^*$  parameters in  $\ell_1$  norm, as a function of number of available cascades for different network types in the case where a fraction of nodes  $\xi$  is unobserved. Each point is averaged over five different networks and five different sets of parameters  $\alpha_{ij}^*$  (sampled from a uniform distribution between zero and one). All networks contain  $N = 100$  nodes. Unobserved nodes were picked at random. For tree graphs, edges adjacent to unobserved leaves were excluded.

algorithm, which results in a finite error for parameter estimation even for a large number of cascades. On the other hand, it is clear that when optimizing the objective function (4), the algorithm attempts to maximize the likelihood of observed marginal probabilities. Hence, the recovered effective parameters should be the ones that attempt to minimize the distance between marginal probabilities estimated from the model, and those observed in the data. Fig. 5(b) shows relative  $\ell_1$  distance between learned and “ground-truth” marginals

$$\delta_p = \frac{\langle |p_i^s(t) - p_i^{s*}(t)| \rangle}{\langle p_i^{s*}(t) \rangle}, \quad (8)$$

where  $p_i^s(t)$  are marginals estimated with DMP from the learned model;  $p_i^{s*}(t)$  are the “ground-truth” values obtained from the total of  $10^8$  Monte Carlo sampled cascades,  $10^6$  per each source  $s$ ; and averaging  $\langle \cdot \rangle$  is done for all nodes, times, and sources. Fig. 5(b) demonstrates that DMP equations run with reconstructed parameters produce a better approximation of real marginal probabilities, compared to DMP predictions using the ground-truth parameters  $\alpha_{ij}^*$ . This observation yields a promising avenue for building an

effective model, which may not reflect the real spreading parameters, but allows for a more accurate prediction of the dynamics, which can subsequently be used for tasks such as influence maximization (Kempe et al., 2003) or control (Nowzari et al., 2016).

These results further confirm the need for coherence in the “learning for inference” approach. Inference tasks are typically computationally harder compared to their learning counterpart. An ideal estimation procedure would perfectly reconstruct model parameters, that may yield inaccurate predictions of the dynamics outcome due to intrinsic errors in approximate inference algorithms. Our analysis shows that a coherent use of an approximate inference method in both learning and inference task leads to an effective model that has a much stronger prediction power. This observation is in agreement with previous work in learning of graphical models (Wainwright, 2006; Heinemann & Globerson, 2014) that showed that an inconsistent parameter estimation that takes approximate inference (such as belief propagation) into account is provably beneficial for the predictive task. This concept further motivates the extension of the search for effective models outside of the original class (for instance,

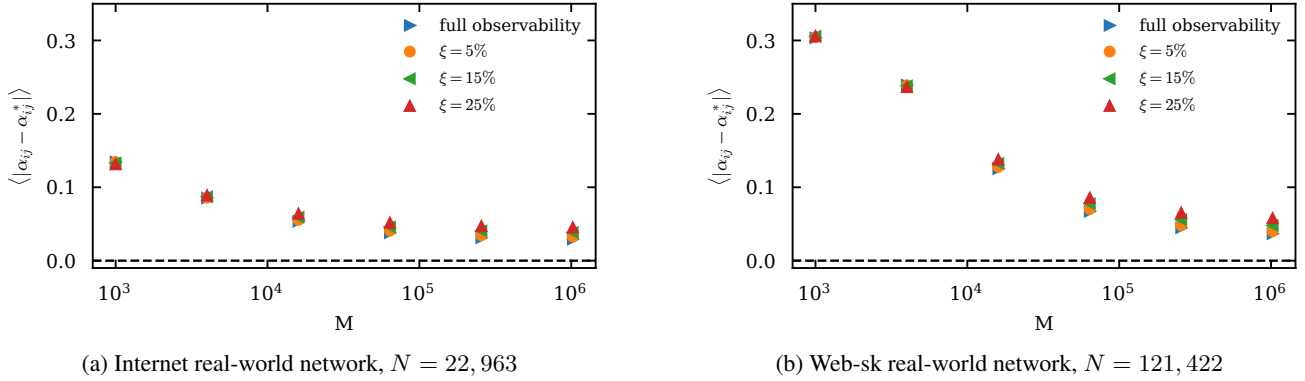


Figure 4. Average difference between inferred  $\alpha_{ij}$  and real  $\alpha_{ij}^*$  parameters in  $\ell_1$  norm, as a function of number of available cascades for two different empirical networks, where a fraction of nodes  $\xi$  is unobserved. Unobserved nodes were picked at random.

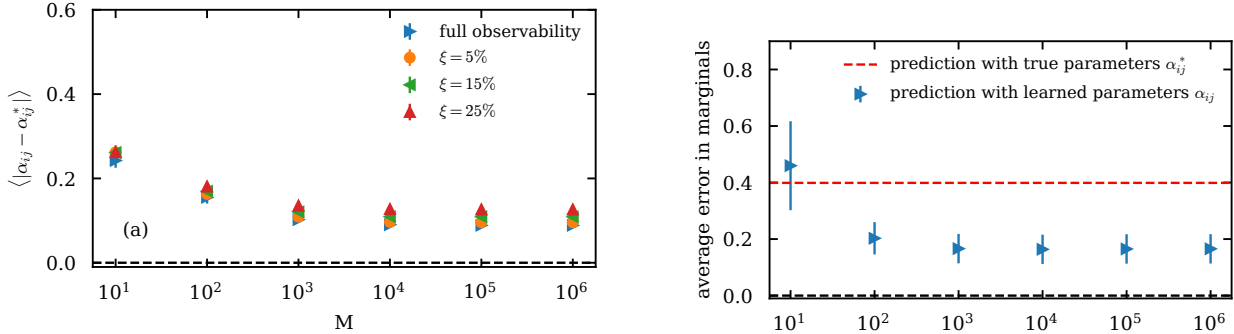


Figure 5. Illustration of the effective model concept, where our algorithm recovers the set of parameters that try to maximize the likelihood of marginal probabilities. Both panels correspond to an adversarial for DMP square lattice topology with many short loops,  $N = 100$  nodes, and cascades of length  $T = 20$ . (a) Distance between true  $\alpha_{ij}^*$  and estimated  $\alpha_{ij}$  parameters. (b) Relative distance between “ground-truth” (estimated through Monte-Carlo) and estimated (via DMP) marginals. The red line represents the relative distance between real marginals and the ones obtained by using the ground-truth parameters  $\alpha_{ij}^*$  within the DMP equations.

allowing for an extended set of parameters), as long as the goal consists in an accurate prediction of the dynamics.

## 5.2. Extension to Multi-Replica Effective Model

Here, we capitalize on the concept of effective models for inference and develop a novel procedure for learning a mixture of spreading models on several replicas of the graph which further improves prediction quality of marginal probabilities. Our proposed approach for introducing a mixture of replicas is motivated by the fact that DMP on a given graph instance is too constrained. Indeed, one may hope to adjust model parameters to compensate for the approximation error. However, in the context of other approximate inference algorithms such as belief propagation, it has been shown that some sets of valid marginals cannot be reached by any parameters or any learning procedure (Pitkow et al., 2011); a similar picture is likely to apply to DMP on a given network. One could naturally try to look at extended graph topologies that could potentially bring some flexibility, however in the context of dynamic models considered here, addi-

tional edges do not benefit the effective model constructions since they are likely inconsistent with the observed cascades. However, several copies of the original topology are always consistent with the data, while they may compensate for errors in marginal prediction by adjusting model parameters in each replica. Notice that introduction of a mixture of replicas in our approach is a mean to increase the expressiveness of the inference model and does not mean that the cascade data is actually produced by the mixture of graphs (Hoffmann et al., 2020).

To test this idea, we introduce several replicas of the model with independent parameters, and where the predicted marginal is given by a uniform mixture over these replicated models. Formally, we consider the objective:

$$\mathcal{O}^{\text{mixture}} = \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{O}} \sum_{\tau_i^s} m^{\tau_i^s} \log \left( \frac{1}{|R|} \sum_{r \in R} \mu_{i,r}^s(\tau_i^s) \right), \quad (9)$$

where  $R$  is the set of replicas, and  $\mu_{i,r}^s(t)$  is the marginal probability of node  $i$  being activated at time  $t$  for a cascade



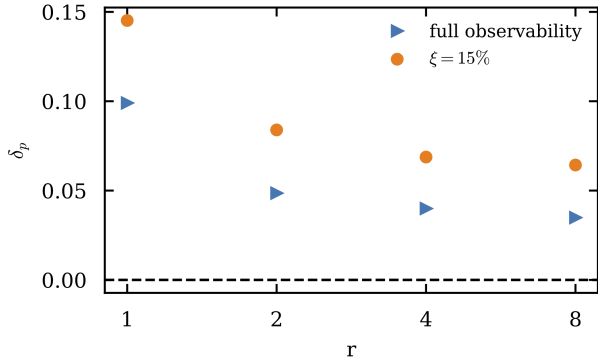


Figure 6. Improvement of the marginal prediction error using a multi-replica effective model as a function of the number of replicas for full  $\xi = 0\%$  and partial  $\xi = 15\%$  observability. The results are obtained for a square lattice with  $N = 100$  nodes and  $M = 10^6$  cascades of length  $T = 20$ .

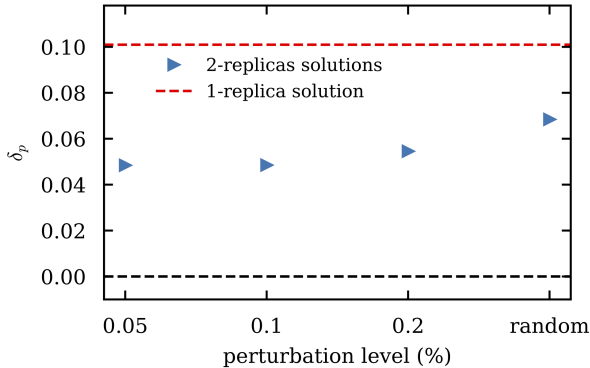


Figure 7. Dependence of the improvement in prediction of marginal probabilities by a 2-replica effective model compared to a single-replica model learned by SLICER as a function of perturbation of the previously obtained solution. The last point corresponds to a search of optimal parameters from scratch using random initialization. The results were obtained for a square lattice with  $N = 100$  nodes and  $M = 10^6$  cascades of length  $T = 20$ .

with source  $s$  in replica model  $r$ . SLICER can be straightforwardly generalized to the objective function (9) with a mixture of replica models. Thus obtained effective model is more expressive since the number of parameters is increased from  $|E|$  to  $|E| \cdot |R|$ ; the computational complexity of the algorithm for this case changes to  $O(|R||E|T|S|)$ . We propose the following scheme for a systematic improvement of the prediction of marginal probabilities by introducing more replicas compared to the solution obtained at the previous step, where the original effective model obtained by SLICER represents a single-replica case. It consists in finding a solution for an effective model with an increased number of replicas by running SLICER initialized with a slightly perturbed solution (to break symmetry between replicas), previously obtained for a model with a smaller number of replicas. The results of this procedure are given in Fig. 6 for the case of a square lattice with long,  $T = 20$ , observation time. We see that subsequent introduction of replicas pro-

gressively and substantially improves the prediction quality of marginal probabilities, both in the case of full and partial observations. Fig. 7 shows that small perturbations from previously obtained solutions produce better results and guarantee improved prediction compared to a search for optimal parameters from scratch. A more detailed discussion of this phenomenon along with supporting experiments are given in the Supplementary Materials, section S5 (Wilinski & Lokhov, 2020). This novel effective mixture model construction thus allows one to systematically overcome structural constraints of the DMP equations, and produce a model with higher-quality predictions.

## 6. Conclusions

Incomplete information about the dynamics significantly complicates the tasks of parameter estimation and learning of a model for generating accurate predictions. In this paper, we propose a novel approximate learning algorithm, SLICER, based on a dynamic message-passing approach. SLICER achieves a significantly better scalability compared to existing approximate methods. Moreover, even in a loopy regime which is adversarial for message-passing, our method explores the space of parameters in search of an effective model that achieves a better prediction of the dynamics. Finally, we propose a novel procedure for estimating a multi-replica effective model that systematically further improves prediction quality of marginal probabilities. We anticipate that scalability and prediction qualities of our algorithm will be useful for real-world applications where large amounts of incomplete data is available and a need of reliable scenario testing is essential. Although our study used the standard discrete-time IC model for illustration, DMP can be generalised to more complex dynamical models (Lokhov et al., 2015), as well as to models with time varying parameters  $\alpha_{ij}(t)$ , adaptive spreading settings, continuous time dynamics or temporal networks. It would be also insightful to extend our approach to the point processes on networks (Upadhyay et al., 2018; Farajtabar et al., 2017). Finally, we believe that our approach will be helpful in the task of optimal resource allocation for both monitoring and containment of spreading processes.

## Code and Supplementary Material

A full implementation of our algorithms is available at (Wilinski & Lokhov, 2021), whereas the supplementary materials can be found in (Wilinski & Lokhov, 2020).

## Acknowledgments

Authors acknowledge support from the Laboratory Directed Research and Development program of Los Alamos National Laboratory under projects numbers 20200121ER and 20210529CR.

## References

- Abbe, E., Kulkarni, S., and Lee, E. J. Nonbacktracking bounds on the influence in independent cascade models. In *Advances in Neural Information Processing Systems*, pp. 1407–1416, 2017.
- Abrahamo, B., Chierichetti, F., Kleinberg, R., and Panconesi, A. Trace complexity of network inference. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 491–499, 2013.
- Amin, K., Heidari, H., and Kearns, M. Learning from contagion (without timestamps). In *International Conference on Machine Learning*, pp. 1845–1853, 2014.
- Boix-Adsera, E., Bresler, G., and Koehler, F. Chow-liu++: Optimal prediction-centric learning of tree ising models. *arXiv preprint arXiv:2106.03969*, 2021.
- Boldi, P., Codenotti, B., Santini, M., and Vigna, S. Ubi-crawler: A scalable fully distributed web crawler. *Software: Practice and Experience*, 34(8):711–726, 2004.
- Braunstein, A., Ingrosso, A., and Muntoni, A. P. Network reconstruction from infection cascades. *Journal of the Royal Society Interface*, 16(151):20180844, 2019.
- Bresler, G. and Karzand, M. Learning a tree-structured ising model in order to make predictions. *Annals of Statistics*, 48(2):713–737, 2020.
- Burkholz, R. and Quackenbush, J. Cascade size distributions and why they matter. *arXiv preprint arXiv:1909.05416*, 2019.
- Chen, W., Wang, Y., and Yang, S. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 199–208, 2009.
- Chen, W., Wang, C., and Wang, Y. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1029–1038, 2010.
- Cohen, E., Delling, D., Pajor, T., and Werneck, R. F. Sketch-based influence maximization and computation: Scaling up with guarantees. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pp. 629–638. ACM, 2014.
- Du, N., Song, L., Rodriguez, M. G., and Zha, H. Scalable influence estimation in continuous-time diffusion networks. In *Advances in neural information processing systems*, pp. 3147–3155, 2013.
- Du, N., Liang, Y., Balcan, M., and Song, L. Influence function learning in information diffusion networks. In *International Conference on Machine Learning*, pp. 2016–2024, 2014.
- Dunning, I., Huchette, J., and Lubin, M. Jump: A modeling language for mathematical optimization. *SIAM review*, 59(2):295–320, 2017.
- Farajtabar, M., Wang, Y., Gomez-Rodriguez, M., Li, S., Zha, H., and Song, L. Coevolve: A joint point process model for information diffusion and network evolution. *The Journal of Machine Learning Research*, 18(1):1305–1353, 2017.
- Goldenberg, J., Libai, B., and Muller, E. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 12(3):211–223, 2001.
- Gomez-Rodriguez, M., Leskovec, J., and Krause, A. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4): 1–37, 2012.
- Gripon, V. and Rabbat, M. Reconstructing a graph from path traces. In *2013 IEEE International Symposium on Information Theory*, pp. 2488–2492. IEEE, 2013.
- Han, K., Tian, Y., Zhang, Y., Han, L., Huang, H., and Gao, Y. Statistical estimation of diffusion network topologies. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pp. 625–636. IEEE, 2020.
- He, X., Xu, K., Kempe, D., and Liu, Y. Learning influence functions from incomplete observations. In *Advances in Neural Information Processing Systems*, pp. 2073–2081, 2016.
- Heinemann, U. and Globerson, A. Inferring with high girth graphical models. In *International Conference on Machine Learning*, pp. 1260–1268. PMLR, 2014.
- Hoffmann, J. and Caramanis, C. Learning graphs from noisy epidemic cascades. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(2): 1–34, 2019.
- Hoffmann, J., Basu, S., Goel, S., and Caramanis, C. Learning mixtures of graphs from epidemic cascades. In *International Conference on Machine Learning*, pp. 4342–4352. PMLR, 2020.
- Karrer, B. and Newman, M. E. Message passing approach for general epidemic models. *Physical Review E*, 82(1): 016101, 2010.
- Kempe, D., Kleinberg, J., and Tardos, É. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference*

- on *Knowledge discovery and data mining*, pp. 137–146, 2003.
- Lokhov, A. Reconstructing parameters of spreading models from partial observations. In *Advances in Neural Information Processing Systems*, pp. 3467–3475, 2016.
- Lokhov, A. Y. and Saad, D. Optimal deployment of resources for maximizing impact in spreading processes. *Proceedings of the National Academy of Sciences*, 114(39):E8138–E8146, 2017.
- Lokhov, A. Y. and Saad, D. Scalable influence estimation without sampling. *arXiv preprint arXiv:1912.12749*, 2019.
- Lokhov, A. Y., Mézard, M., Ohta, H., and Zdeborová, L. Inferring the origin of an epidemic with a dynamic message-passing algorithm. *Physical Review E*, 90(1):012801, 2014.
- Lokhov, A. Y., Mézard, M., and Zdeborová, L. Dynamic message-passing equations for models with unidirectional dynamics. *Physical Review E*, 91(1):012811, 2015.
- Lucier, B., Oren, J., and Singer, Y. Influence at scale: Distributed computation of complex contagion in networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 735–744. ACM, 2015.
- Makar, M., Guttag, J., and Wiens, J. Learning the probability of activation in the presence of latent spreaders. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Narasimhan, H., Parkes, D. C., and Singer, Y. Learnability of influence in networks. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*, pp. 3186–3194. Curran Associates, Inc., 2015.
- Netrapalli, P. and Sanghavi, S. Learning the graph of epidemic cascades. *ACM SIGMETRICS Performance Evaluation Review*, 40(1):211–222, 2012.
- Nguyen, H. T., Nguyen, T. P., Vu, T. N., and Dinh, T. N. Outward influence and cascade size estimation in billion-scale networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1):20, 2017.
- Nowzari, C., Preciado, V. M., and Pappas, G. J. Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems Magazine*, 36(1):26–46, 2016.
- Pitkow, X., Ahmadian, Y., and Miller, K. D. Learning unbelievable probabilities. *Advances in neural information processing systems*, 24:738, 2011.
- Pouget-Abadie, J. and Horel, T. Inferring graphs from cascades: A sparse recovery framework. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML15*, pp. 977986. JMLR.org, 2015.
- Rossi, R. and Ahmed, N. The network data repository with interactive graph analytics and visualization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Sefer, E. and Kingsford, C. Convex risk minimization to infer networks from probabilistic diffusion data at multiple scales. In *2015 IEEE 31st International Conference on Data Engineering*, pp. 663–674. IEEE, 2015.
- Shapiro, M. and Delgado-Eckert, E. Finding the probability of infection in an sir network is np-hard. *Mathematical biosciences*, 240(2):77–84, 2012.
- Supeesun, A. and Fakcharoenphol, J. Learning network structures from contagion. *Information Processing Letters*, 121:11–16, 2017.
- Trouleau, W., Etesami, J., Grossglauser, M., Kiyavash, N., and Thiran, P. Learning hawkes processes under synchronization noise. In *International Conference on Machine Learning*, pp. 6325–6334. PMLR, 2019.
- Upadhyay, U., De, A., and Gomez-Rodriguez, M. Deep reinforcement learning of marked temporal point processes. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 3172–3182, 2018.
- Wächter, A. and Biegler, L. T. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.
- Wainwright, M. J. Estimating the “wrong” graphical model: Benefits in the computation-limited setting. *Journal of Machine Learning Research*, 7(9), 2006.
- Wainwright, M. J. and Jordan, M. I. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- Wilinski, M. and Lokhov, A. SLICER repository, 2021. URL <https://github.com/mateuszwilinski/dynamic-message-passing/>.
- Wilinski, M. and Lokhov, A. Y. Scalable learning of independent cascade dynamics from partial observations. *arXiv preprint arXiv:2007.06557*, 2020.