

Deep Demixing: Reconstructing the Evolution of Epidemics using Graph Neural Networks

Gojko Čutura, Boning Li, Ananthram Swami, and Santiago Segarra

Abstract—We study the temporal reconstruction of epidemics evolving over networks. Given partial or aggregated temporal information of the epidemic, our goal is to estimate the complete evolution of the spread leveraging the topology of the network but being agnostic to the precise epidemic model. We overcome this lack of model awareness through a data-driven solution to the inverse problem at hand. In particular, we propose DDmix, a graph conditional variational autoencoder that can be trained from past epidemic spreads and whose latent space seeks to capture key aspects of the underlying (unknown) spreading dynamics. We illustrate the accuracy and generalizability of DDmix and compare it with non-graph-aware learning algorithms through numerical experiments on epidemic spreads simulated on synthetic and real-world networks.

Index Terms—Network, inverse problem, epidemics, graph neural network, variational autoencoder

I. INTRODUCTION

Networks or graphs have emerged as useful models to represent complex interconnected systems and data defined on them. These representations have wide applicability across multiple domains, including neuroscience [1], sociology [2], and urban planning [3]. In an attempt to better understand and learn from data defined on networks, classical signal processing and machine learning methods have recently been extended to encompass this data type [4]. These novel tools have shown remarkable performance in a variety of popular network science problems such as node classification [5], link prediction [6], and several inference tasks related to partially observed network processes [7], [8].

Within the range of network processes, the use of graphs for modeling and understanding *epidemics* is a fertile subfield [9]. Networks provide versatile modelling tools where nodes might represent anything from single individuals [10] to large cities [11] and edges can encode different mechanisms of disease propagation – airborne, contact, or vector transmission – across the nodes of the network. Classical graph features such as node centralities and connectivity measures can then be used to inform public health measures, e.g., immunization strategies and lockdown procedures [12], [13].

In this paper, we go beyond established classical methods and propose a novel neural network architecture to solve the

G. Čutura is with the University of Belgrade, Serbia. B. Li and S. Segarra are with Rice University, USA. A. Swami is with the US Army Research Laboratory, USA. Research was sponsored by the Army Research Office and was accomplished under Cooperative Agreement Number W911NF-19-2-0269. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. Emails: cg203066m@student.etf.bg.ac.rs, {boning.li, segarra}@rice.edu, ananthram.swami.civ@mail.mil.

challenging problem of temporally reconstructing an epidemic. More precisely, given partial or aggregated temporal information of the epidemic, we want to infer the evolution of the spread. Being an inherently ill-posed inverse problem, one might rely on precise knowledge of the network process or structural features of the initial condition to solve this underdeterminacy [7]. By contrast, we propose a *model-inspired data-driven* solution. Namely, we put forth a conditional variational auto-encoder (CVAE) [14] based on graph neural networks (GNNs) [5] that is trained to solve the inverse problem from available data and, importantly, whose architecture and training loss are inspired by the locality of the spreading.

Related work. We can model an epidemic spread as the temporal evolution of a signal defined on the nodes of a graph. From this viewpoint, epidemic reconstruction boils down to different well-studied problems depending on the observation model. More precisely, if the state of some nodes is observed and we want to infer the state of the remaining nodes, the problem can be modeled as the interpolation of graph signals [15]; if the final state of every node is observed, the problem resembles blind deconvolution on graphs [7]; and if a temporally aggregated signal for each node is observed, then the problem boils down to *demixing* of graph signals [16]. This last body of work is the one that best fits our observation model. However, existing tools [16] were derived for *linear* network processes that cannot accurately model epidemic spreads, thus prompting the need for *non-linear* methods as DDmix, the one derived here.

From the perspective of computational epidemiology, temporal reconstruction has also received attention [17]. For example, under the assumption that an epidemic spread follows the SI model, [18] recovers multiple source nodes from a single snapshot of the complete graph. Similarly, [19] focuses on identifying key spreaders under the independent-cascade model. In general, this body of work largely relies on the precise knowledge of the epidemic model or the time of contagion for a subset of nodes. We depart from this paradigm and, instead, rely on observed past data to learn how to solve the reconstruction problem.

Related to our proposed architecture, CVAEs [14] have been shown to be effective in solving prediction and data generation problems in non-graph settings [20] and have been used for the reconstruction of videos from temporally aggregated data [21]. Motivated by this success, some of these tools have been extended to graph settings [22] and mostly applied as node embedding procedures [23]. To the best of our knowledge, this is the first implementation of a graph CVAE for the study of epidemics and, generally, for the solution of inverse problems related to network processes.

Contribution. The contributions of our paper are twofold:
 i) We propose DDmix, a novel graph CVAE architecture for the temporal reconstruction of partially-observed network processes; and
 ii) We successfully implement DDmix to infer the evolution of epidemics surpassing non-graph-aware deep learning methods in terms of accuracy and generalizability.

II. SYSTEM MODEL AND PROBLEM STATEMENT

We model our inter-connected system as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes in \mathcal{V} representing individuals and edges $(i, j) \in \mathcal{E}$ encoding the possibility of contagion between nodes i and j . The graph structure can be represented using the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, where $A_{ij} = A_{ji} = 1$ for all $(i, j) \in \mathcal{E}$, and $A_{ij} = A_{ji} = 0$ otherwise. We model the state of the nodes at any given time using graph signals, i.e., maps $x : \mathcal{V} \rightarrow \mathbb{R}$ from the node set into the reals. Graph signals can be conveniently represented as vectors $\mathbf{x} \in \mathbb{R}^N$, where the entry x_i collects the value of the graph signal at node i . Specifically, we consider a time series of graph signals $\{\mathbf{y}^{(t)}\}_{t=1}^T$ where $\mathbf{y}^{(t)} \in \{0, 1\}^N$ and $y_i^{(t)} = 1$ indicates that node i was in the infected state at time instant t while $y_i^{(t)} = 0$ indicates that it was not infected at that time.

Although not assumed to be known by our proposed solution, for simplicity in this paper we focus on a particular epidemic model, namely the well-established SIRS model [10]. This is a parametric stochastic model that, given $\mathbf{y}^{(1)}$, determines the probability distribution of the signals $\mathbf{y}^{(t)}$ for $t > 1$. More precisely, the SIRS model is defined by three parameters – the infection probability β , the healing probability δ , and the probability of losing immunity γ – and every node is at one of three states – susceptible (S), infected (I), or recovered (R). At every discrete time step t : i) Susceptible nodes can get infected independently with probability β by any of their infected neighbors in \mathcal{G} , ii) Infected nodes recover with probability δ , and iii) Recovered nodes become susceptible with probability γ .

Under an (unknown) SIRS spread $\{\mathbf{y}^{(t)}\}_{t=1}^T$ of interest, our partial information is given by a single graph signal $\mathbf{x} = f(\{\mathbf{y}^{(t)}\}_{t=1}^T)$. In this paper we focus on the case where \mathbf{x} represents a temporal aggregation of the reported infections, but adopting the assumption that individuals do not report back to the hospital when they get well. More precisely, defining $\mathbf{q}^{(t)} = \max_{i \in \{1, \dots, t\}} \mathbf{y}^{(i)}$, where the maximum is taken elementwise, we have

$$\mathbf{x} = \frac{1}{T} \sum_{t=1}^T \mathbf{q}^{(t)}. \quad (1)$$

With this notation in place, we can formally state our problem.

Problem 1: Consider an unknown epidemic $\{\mathbf{y}^{(t)}\}_{t=1}^T$ that spreads over a known graph \mathcal{G} . Given the temporally aggregated information \mathbf{x} as in (1), estimate the complete evolution of the spread $\{\mathbf{y}^{(t)}\}_{t=1}^T$.

Before presenting our proposed solution to Problem 1, a few comments are in order. First, the observation model in (1) is realistic and reasonable in practice. We get to observe if

a person was infected or not over a given period of time as well as the length of the infection (either self-reported or estimated by proxies such as serology tests), but we do not have precise information of when the infection started and ended. Second, Problem 1 can be reinterpreted as one of signal demixing in graphs since we are given the (temporal) aggregation of many signals and we want to tell them apart leveraging the graph structure. This is a challenging and highly underdetermined problem even in the simpler setting of linear network processes. Finally, to solve this ill-posed and non-linear problem, most existing approaches assume precise knowledge of the epidemic model [18], [19]. By contrast, we take a data-driven perspective where we assume nothing about the underlying epidemic model other than it being driven by the topology of \mathcal{G} , and we use (few) past known observation pairs $(\mathbf{x}, \{\mathbf{y}^{(t)}\}_{t=1}^T)$ to train our model.

III. GRAPH CVAE FOR TEMPORAL RECONSTRUCTION

Assume that we have access to M pairs (\mathbf{x}, \mathbf{Y}) where, for notation simplicity, we have collected the T vectors in $\{\mathbf{y}^{(t)}\}_{t=1}^T$ as columns of $\mathbf{Y} \in \mathbb{R}^{N \times T}$. Our goal is to leverage these observations to estimate the conditional distribution $p(\mathbf{Y}|\mathbf{x})$ for the particular scenario of interest. In this way, given a new observation \mathbf{x} , we can sample our candidate temporal evolution \mathbf{Y} from $p(\mathbf{Y}|\mathbf{x})$. In determining $p(\mathbf{Y}|\mathbf{x})$, we adopt a CVAE probabilistic model [14] with a latent variable \mathbf{z} that seeks to model features of the temporal variation of the epidemic that have been collapsed in \mathbf{x} . Intuitively, we want \mathbf{z} to capture key aspects of the underlying (unknown) spreading dynamics. We model \mathbf{z} as conditionally Gaussian given \mathbf{x} , i.e., $p_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mu_\phi(\mathbf{x}), \sigma_\phi^2(\mathbf{x}))$. Notice that we have made explicit our focus on a parametric form of the conditional probability, where both the mean and the standard deviation of \mathbf{z} are functions of \mathbf{x} dependent on the parameters ϕ . Following the CVAE framework, we define the distribution of our variable of interest as

$$p_\theta(\mathbf{Y}|\mathbf{x}, \mathbf{z}) = \mathcal{N}(g_\theta(\mathbf{x}, \mathbf{z}), \sigma_y^2 \mathbf{I}). \quad (2)$$

Under our probabilistic model, (2) reveals that if we have access to the true latent variable \mathbf{z} in combination with our observation \mathbf{x} , then we can apply a parametric *deprojection* function g_θ to get the expected value of the temporal evolution \mathbf{Y} , where σ_y^2 is a common noise variance for all entries. From the two conditional probabilities introduced, it follows that we can compute the distribution of interest as

$$p_{\theta, \phi}(\mathbf{Y}|\mathbf{x}) = \int_{\mathbf{z}} p_\theta(\mathbf{Y}|\mathbf{x}, \mathbf{z}) p_\phi(\mathbf{z}|\mathbf{x}) d\mathbf{z}. \quad (3)$$

However, solving the integral in (3) can become intractable even for fairly simple parametrizations of p_θ and p_ϕ . Thus, determining the parameters θ and ϕ that maximize the likelihood of the observed M pairs (\mathbf{x}, \mathbf{Y}) is a challenging endeavor in general. Instead, we follow the well-accepted route of variational inference and establish a tractable loss inspired by the evidence lower bound (see [24] for details) that can be optimized via stochastic gradient descent. To present this loss, we

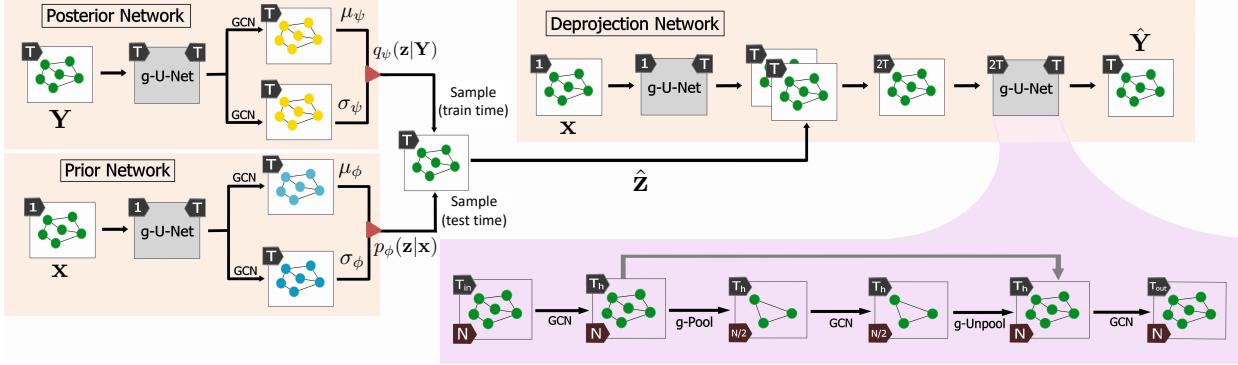


Fig. 1. Overall view of our proposed deep demixing (DDmix) architecture based on graph neural networks for the temporal reconstruction of epidemics. The three subnetworks – prior, posterior, and deprojection – parametrize key functions in the assumed probabilistic model.

first need to introduce a parametrization for the posterior probability of \mathbf{z} given \mathbf{Y} , namely $q_\psi(\mathbf{z}|\mathbf{Y}) = \mathcal{N}(\mu_\psi(\mathbf{Y}), \sigma_\psi^2(\mathbf{Y}))$. We then consider the following compound loss

$$\mathcal{L}_{\theta,\phi,\psi}(\mathbf{x}, \mathbf{Y}) = L_1(q_\psi(\mathbf{z}|\mathbf{Y}), p_\phi(\mathbf{z}|\mathbf{x})) + \eta_1 L_2(g_\theta(\mathbf{x}, \hat{\mathbf{z}}), \mathbf{Y}) + \eta_2 R_1(\theta, \phi, \psi) + \eta_3 R_2(g_\theta(\mathbf{x}, \hat{\mathbf{z}})), \quad (4)$$

consisting of two fitting terms $L_1(\cdot)$ and $L_2(\cdot)$, two regularization terms $R_1(\cdot)$ and $R_2(\cdot)$, and their relative importance given by the scalar weights $\{\eta_i\}_{i=1}^3$. In our implementation, we select the loss $L_1(\cdot)$ as the KL divergence between the argument distributions which ensures that, during testing, our draws from $p_\phi(\mathbf{z}|\mathbf{x})$ will be close to the more informative draws from $q_\psi(\mathbf{z}|\mathbf{Y})$. For the reconstruction loss $L_2(\cdot)$, we recall that the entries of \mathbf{Y} are binary – either infected or not – and select an entry-wise binary cross entropy loss. In terms of regularization, we implement in $R_1(\cdot)$ an ℓ_2 penalty on all our parameters. More interestingly, in $R_2(\cdot)$ we incorporate the knowledge that the evolution of the epidemic should be local in \mathcal{G} . To be precise, we define $R_2(g_\theta(\mathbf{x}, \hat{\mathbf{z}})) = \|\sum_{t=2}^T [\hat{\mathbf{y}}^{(t)} - (\mathbf{A} + \mathbf{I})\hat{\mathbf{y}}^{(t-1)}]_+\|_1$, where $[\cdot]_+$ denotes the positive projection and $\hat{\mathbf{y}}^{(t)}$ is the t -th column of $\hat{\mathbf{Y}} = g_\theta(\mathbf{x}, \hat{\mathbf{z}})$. Intuitively, in $R_2(\cdot)$ we penalize the appearance of an infected node at time t when neither itself nor any of its neighbors was infected at time $t-1$. Finally, notice that the loss in (4) is actually a random quantity since the latent variable $\hat{\mathbf{z}}$ used in the computation of $L_2(\cdot)$ and $R_2(\cdot)$ is drawn from $q_\psi(\mathbf{z}|\mathbf{Y})$. However, we can still optimize over it using the reparametrization trick [24].

We implement the parametric functions in our model – g_θ and the mean and standard deviations of p_ϕ and q_ψ – using graph neural networks. Our overall deep demixing architecture (DDmix) is illustrated in Fig 1. An essential building block of our model is the Graph U-Net (g-U-Net) [25], which consists of down-sampling and up-sampling graph convolutional network (GCN) layers [5], alongside graph pooling layers (gPool and gUnpool). More precisely, the output of a generic GCN layer is given by $\mathbf{H} = \text{ReLU}(\tilde{\mathbf{A}}\mathbf{H}_0\mathbf{W})$, where \mathbf{H}_0 is the input to the layer, $\tilde{\mathbf{A}}$ is a normalized adjacency matrix, and \mathbf{W} is a matrix of trainable weights for that layer. The gPool layer performs a global max-pooling operation on graph data by introducing a trainable projection vector that maps the features of a given node onto a scalar; see [25] for details.

On the other hand, the gUnpool layer restores the original graph structure by performing the inverse operation to the corresponding gPool layer. In our implementation, the four g-U-Nets included in Fig. 1 have exactly one pair of pooling and unpooling layers, where the pooling operation reduces in half the number of nodes in the graph. For each of the g-U-Net blocks in Fig. 1, the number of input and output node features can be seen at the top left and right corners, respectively.

The prior and posterior networks are used to respectively parametrize $p_\phi(\mathbf{z}|\mathbf{x})$ and $q_\psi(\mathbf{z}|\mathbf{Y})$. The prior network takes as input the aggregated graph signal \mathbf{x} with a single feature per node and passes it through the g-U-Net block while increasing the number of features per node to T . Two parallel GCN layers are then applied to the output of g-U-Net, resulting in the mean μ_ϕ and standard deviation σ_ϕ of the Gaussian distribution p_ϕ . We design the posterior network analogously, with the only difference being that the number of features per node in the input signal \mathbf{Y} is T , corresponding to the complete, uncollapsed time series. The obtained parameters are then used to sample from the corresponding distribution (denoted by the red triangles in Fig. 1), where the posterior network is used during training and the prior network during testing.

Finally, the deprojection network takes as input both the observed \mathbf{x} and the drawn latent variable $\hat{\mathbf{z}}$ and seeks to output a good estimate $\hat{\mathbf{Y}} = g_\theta(\mathbf{x}, \hat{\mathbf{z}})$ of the temporal evolution. In our implementation, we first pass \mathbf{x} through a g-U-Net block while increasing the number of node features to T . The output of this block is concatenated with $\hat{\mathbf{z}}$, resulting in a graph signal with $2T$ features per node. This signal is then passed through our last g-U-Net block that combines the information in \mathbf{x} and $\hat{\mathbf{z}}$ and reduces the number of node features to T to match the length of the time series being estimated.

IV. NUMERICAL EXPERIMENTS

Through synthetic and real-world graphs, we illustrate the behavior of DDMix in diverse settings.¹ In evaluating its performance, we compare DDMix with the following baselines:

- i) **MLP:** A multi-layer perceptron that takes \mathbf{x} as input, propagates it through 3 hidden layers with $\{\frac{1}{4}NT, \frac{1}{4}NT, NT\}$

¹Code to replicate the numerical experiments here presented can be found at https://github.com/gojko54/Deep_demixing.

TABLE I
MSE OF PREDICTIONS ON UNSEEN GRAPHS WITH DIFFERENT GRAPH DENSITY FOR TWO DIFFERENT LENGTHS OF THE EPIDEMICS.

Graph density	Baseline		Denser		Sparsier	
Time steps (T)	10	20	10	20	10	20
MLP	.217	.233	.289	.268	.254	.254
CNN-nodes	.106	.188	.270	.239	.177	.232
CNN-time	.143	.194	.267	.240	.168	.226
DDmix	.101	.177	.160	.197	.075	.216

neurons and ReLU activations. The output layer of size NT estimates the vertical concatenation of the columns $\mathbf{y}^{(t)}$ of \mathbf{Y} .

ii) **CNN-nodes:** A convolutional neural network (CNN) that takes \mathbf{x} as input and performs one-dimensional depthwise convolutions with kernels of size 3, and both stride and padding equal to 1. No pooling is performed and the number of channels gradually increases from 1 to T ($1, T/4, T/2, T$) so that the $N \times T$ two-dimensional output estimates $\hat{\mathbf{Y}}$.

iii) **CNN-time:** A CNN that performs one-dimensional transposed convolutions [26] over the temporal dimension with fractional strides. For each node i , CNN-time takes x_i as one-dimensional input and after 6 blocks of convolution, batch normalization and elementwise ReLU activation, it computes a T -dimensional output estimating $[y_i^{(1)}, \dots, y_i^{(T)}]$.

It should be noted that, although data-driven, the three conventional machine learning baselines considered are graph agnostic. More precisely, MLP incorporates fully-connected layers, thus overlooking the notion of locality in \mathcal{G} . CNN-nodes, relying on convolutional filters, assumes a notion of locality inherited by the indexing of the nodes, which need not align with the true notion of locality driving the underlying epidemic. Finally, CNN-time ignores the effect of interconnections between nodes and seeks to solve the temporal reconstruction for each node independently. By contrast, DDMix explicitly incorporates the graph structure in its architecture redoubling in higher performance and enhancing its generalizability.

For our synthetic graphs, we use a random geometric graph generator that places N nodes uniformly at random in the unit cube. An edge is inserted between two nodes if their Euclidean distance is less than or equal to d_r . For generating graphs of size $N \in \{100, 250, 500, 1000\}$ nodes, we use parameters $d_r \in \{0.25, 0.15, 0.1, 0.075\}$, respectively. SIRS parameters (cf. Section II) are set to $\beta = 0.15$, $\delta = 0.1$ and $\gamma = 0.01$. We set $\eta_1 = \eta_3 = 1$ and $\eta_2 = 10^{-6}$ in the loss (4). We use the Adam optimizer with a fixed learning rate of 10^{-2} , $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The models are trained in mini-batches of size 4, and a maximum of 50 epochs with auto-stop mechanism in case the validation loss stops dropping for 5 epochs. Unless otherwise stated, we train the models with 4500 realizations of $T=20$ step data on a random graph of size $N=100$ nodes and use as figure of merit the mean square error (MSE) of the reconstruction given by $\frac{1}{NT} \|\mathbf{Y} - \hat{\mathbf{Y}}\|_F^2$.

Varying the density of the graph. In this experiment, we test the candidate models using data randomly generated on three different graphs with the same number of nodes ($N=100$) as in the training graph but different underlying topology. To be specific, their graph densities are respectively the same (d_r), denser ($1.2d_r$) and sparser ($0.7d_r$) compared to the

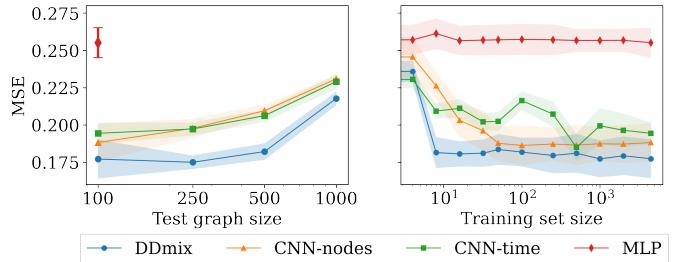


Fig. 2. Comparison of models by their temporal reconstruction performance (MSE) against test graph size and training set size. DDMix performs favorably in terms of robustness and generalizability.

density of the training graph. For each test graph, we test the temporal reconstruction of $T \in \{10, 20\}$ consecutive steps. Table I summarizes the models' performance revealing that DDMix significantly outperforms the baseline methods. This observation aligns well with our hypothesis: Due to their lack of dependence on the underlying graph structure, non-graph-aware methods are unable to adjust their outputs to topologies unseen during training. However, DDMix incorporates the adjacency matrix in its architecture, thus, it is more robust to topological perturbations. Indeed, provided identical collapsed signals on differently structured graphs, CNN and MLP based models would always output identical reconstructions.

Varying the size of the graph. We also tested the models with data generated on graphs of different sizes. We fix the number of steps to be $T = 20$ and apply models trained on 100-node data to graphs of size $N \in \{100, 250, 500, 1000\}$. As one might expect, the MLP model is unable to handle inputs different from $N = 100$. The variation of MSE for the other methods as a function of test graph size is depicted in Fig. 2 (left). It can be seen that DDMix consistently outperforms the competing approaches. More importantly, DDMix is more robust to changes in the growing size of the test graphs, leading to increased performance gaps for $N \in \{250, 500\}$.

Varying the size of the training set. Finally, we are interested in the minimal sufficient training set for each model to achieve high accuracy and good generalizability. In Fig. 2 (right) we illustrate how MSE generally drops as more data (from 0 to 4500 samples) are used for training each model. Initially, all models perform poorly with 4 or fewer training samples. Starting at 8 samples, the models (except for MLP) begin to show improved performance and saturate at around 1000 samples. MLP, having the largest number of parameters and not exploiting any locality, cannot effectively learn in this setting. Most interestingly, DDMix is able to generalize fairly well with only 8 independent training samples, significantly faster than the competing approaches. The fact that DDMix learns to recover collapsed signals much easier indicates that our proposed architecture has successfully imposed the right implicit bias by incorporating the graph topology.

Reconstructing an epidemic in a primary school. We study the identification of epidemic sources in a real-world school network. Using a public dataset [27], [28] containing 2 days of temporal contact among 232 children from 10 school classes, we construct 2 daily contact graphs, one for training ($N = 226$)

and one for testing ($N=228$). The graphs are undirected and unweighted, and an edge exists between a pair of children if they had more than 5 face-to-face contacts per day. During training, a subset of 4 classes is dropped in order to speed up the process and provide a better setting to test generalizability. During testing, the graph has all possible classes and children. Synthetic epidemic data contains $T = 20$ days of SIRS epidemic evolution with random initial infected locations. We simulate 1000 samples for training and 1000 for testing.

Based on this real-world interpersonal network, we evaluate the models by their ability to correctly locate the class where the epidemic started. Given a reconstructed \hat{Y} , we determine a top- k ranking of the source class as follows. First, we determine the first day for which a prediction larger than 0.5 was made for at least one node, i.e., the smallest t' for which there exists some i such that $\hat{y}_i^{(t')} > 0.5$. We then focus on day t' and rank the classes based on their largest prediction of infection, i.e., we rank each class C based on $\max_i \hat{y}_i^{(t')}$ for all nodes $i \in C$. We say that a top- k ranking yields accuracy equal to 1 if it contains the true source class and 0 otherwise. The average accuracy of DDMix is {25.9, 61.2, 82.9}% for the top-1, 3 and 5 predictions, respectively. This significantly outperforms CNN-nodes' {12.8, 36.8, 60.0}% and CNN-time's {11.6, 30.4, 53.5}% mean accuracies. Accurately tracing the source of epidemics is an increasingly important problem, thus motivating part of our future work.

V. CONCLUSIONS AND FUTURE WORK

We introduced DDMix, a novel graph CVAE architecture for temporal reconstruction of network epidemics from aggregated observations. Being agnostic to the model of epidemic spread, DDMix relies on the network topology and the observation of past epidemics to solve this inverse (temporal demixing) problem. Through numerical experiments, we showed that DDMix outperforms non-graph-aware learning techniques in accuracy as well as generalizability across graph distributions and sizes. Current and future research goals include: i) The use of DDMix for epidemic prediction and its evaluation on a variety of real epidemic processes, ii) The development of architectures geared towards the identification of key epidemic actors (such as sources or super-spreaders) without the need for full temporal reconstruction, iii) The study of other observation models where data is only observed at a subset of nodes or temporal data is randomly missing, and iv) The analysis of how the performance and generalization ability of DDMix depend on the network topology (e.g., its community structure) [29].

REFERENCES

- [1] J. D. Medaglia, W. Huang, S. Segarra, C. Olm, J. Gee, M. Grossman, A. Ribeiro, C. T. McMillan, and D. S. Bassett, "Brain network efficiency is influenced by the pathologic source of corticobasal syndrome," *Neurology*, vol. 89, no. 13, pp. 1373–1381, 2017.
- [2] M. Newman, D. Watts, and S. Strogatz, "Random graph models of social networks," *Nat. Academy of Sc. of the USA*, vol. 99 Suppl 1, pp. 2566–72, 2002.
- [3] E. Dogrusoz and S. Aksoy, "Modeling urban structures using graph-based spatial patterns," in *IEEE Int'l. Geoscience and Remote Sensing Symp.*, 2007, pp. 4826–4829.
- [4] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [5] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *CoRR*, vol. abs/1609.02907, 2016.
- [6] M. Zhang and Y. Chen, "Link prediction based on graph neural networks," in *Advances in Neural Inf. Process. Sys.*, 2018, pp. 5165–5175.
- [7] S. Segarra, G. Mateos, A. G. Marques, and A. Ribeiro, "Blind identification of graph filters," *IEEE Trans. Signal Process.*, vol. 65, no. 5, pp. 1146–1159, March 2017.
- [8] Y. Zhu, F. J. I. Garcia, A. G. Marques, and S. Segarra, "Estimating network processes via blind identification of multiple graph filters," *IEEE Trans. Signal Process.*, vol. 68, pp. 3049–3063, 2020.
- [9] I. Z. Kiss, J. C. Miller, and P. Simon, *Mathematics of epidemics on networks: from exact to approximate models*, Springer, 2017.
- [10] L. Zheng and L. Tang, "A node-based sirs epidemic model with infective media on complex networks," *Complexity*, vol. 2019, pp. 2849196, 2019.
- [11] B. S. Pujari and S. M. Shekhar, "Multi-city modeling of epidemics using spatial networks: Application to 2019-ncov (COVID-19) coronavirus in india," *medRxiv*, 2020.
- [12] O. Cliff, V. Sintchenko, T. Sorrell, K. Vadlamudi, N. McLean, and M. Prokopenko, "Network properties of salmonella epidemics," *Scientific Reports*, vol. 9, 2019.
- [13] P. Block, M. Hoffman, I. J. Raabe, J. B. Dowd, C. Rahal, R. Kashyap, and M. C. Mills, "Social network-based distancing strategies to flatten the COVID-19 curve in a post-lockdown world," *Nature Hum. Behav.*, vol. 4, no. 6, pp. 588–596, 2020.
- [14] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Advances in Neural Inf. Process. Sys.*, 2015, pp. 3483–3491.
- [15] S. Chen, R. Varma, A. Sandryhaila, and J. Kovacevic, "Discrete signal processing on graphs: Sampling theory," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6510–6523, Dec. 2015.
- [16] F. J. Iglesias, S. Segarra, S. Rey-Escudero, A. G. Marques, and D. Ramírez, "Demixing and blind deconvolution of graph-diffused sparse signals," in *IEEE Int'l. Conf. on Acoustics, Speech and Signal Process.*, 2018, pp. 4189–4193.
- [17] P. Rozenshtein, A. Gionis, B. A. Prakash, and J. Vreeken, "Reconstructing an epidemic over time," in *Intl Conf. on Knowledge Disc. and Data Mining*, New York, NY, USA, 2016, p. 1835–1844.
- [18] B. A. Prakash, J. Vreeken, and C. Faloutsos, "Spotting culprits in epidemics: How many and which ones?," in *IEEE Intl Conf. on Data Mining*, 2012, pp. 11–20.
- [19] T. Lappas, E. Terzi, D. Gunopulos, and H. Mannila, "Finding effectors in social networks," in *Intl Conf. on Knowledge Disc. and Data Mining*, 2010, p. 1059–1068, Association for Computing Machinery.
- [20] J. Lim, S. Ryu, J. W. Kim, and W. Y. Kim, "Molecular generative model based on conditional variational autoencoder for de novo molecular design," *Journal of Cheminformatics*, vol. 10, no. 1, pp. 31, 2018.
- [21] G. Balakrishnan, A. V. Dalca, A. Zhao, J. V. Guttag, F. Durand, and W. T. Freeman, "Visual deprojection: Probabilistic recovery of collapsed dimensions," in *IEEE/CVF Intl. Conf. on Comp. Vision (ICCV)*, October 2019, pp. 171–180.
- [22] T. N. Kipf and M. Welling, "Variational graph auto-encoders," *arXiv preprint arXiv:1611.07308*, 2016.
- [23] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in Neural Inf. Process. Sys.*, 2017, pp. 1024–1034.
- [24] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [25] H. Gao and S. Ji, "Graph u-nets," *arXiv preprint arXiv:1905.05178*, 2019.
- [26] V. Dumoulin and F. Visin, "A guide to convolution arithmetic for deep learning," *arXiv preprint arXiv:1603.07285*, 2016.
- [27] V. Gemmetto, A. Barrat, and C. Cattuto, "Mitigation of infectious disease at school: targeted class closure vs school closure," *BMC Infect. dis.*, vol. 14, no. 1, pp. 695, 2014.
- [28] J. Stehlé, N. Voisin, A. Barrat, C. Cattuto, L. Isella, J.-F. Pinton, M. Quaggiotto, W. Van den Broeck, C. Régis, B. Lina, et al., "High-resolution measurements of face-to-face contact patterns in a primary school," *PloS one*, vol. 6, no. 8, pp. e23176, 2011, Data available at <http://www.sociopatterns.org>.
- [29] A. Jung, A. O. Hero III, A. C. Mara, S. Jahromi, A. Heimowitz, and Y. C. Eldar, "Semi-supervised learning in network-structured data via total variation minimization," *IEEE Trans. Signal Process.*, vol. 67, no. 24, pp. 6256–6269, 2019.