

DANR: Discrepancy-aware Network Regularization

Hongyuan You, Furkan Kocayusufoglu, Ambuj K. Singh*

Abstract

Network regularization is an effective tool for incorporating structural prior knowledge to learn coherent models over networks, and has yielded provably accurate estimates in applications ranging from spatial economics to neuroimaging studies. Recently, there has been an increasing interest in extending network regularization to the spatio-temporal case to accommodate the evolution of networks. However, in both static and spatio-temporal cases, missing or corrupted edge weights can compromise the ability of network regularization to discover desired solutions. To address these gaps, we propose a novel approach—*discrepancy-aware network regularization* (DANR)—that is robust to inadequate regularizations and effectively captures model evolution and structural changes over spatio-temporal networks. We develop a distributed and scalable algorithm based on alternating direction method of multipliers (ADMM) to solve the proposed problem with guaranteed convergence to global optimum solutions. Experimental results on both synthetic and real-world networks demonstrate that our approach achieves improved performance on various tasks, and enables interpretation of model changes in evolving networks.

1 INTRODUCTION

Network regularization is a fundamental approach to encode and incorporate general relationships among variables, which are represented as nodes and linked together by weighted edges that describe their local proximity. A significant amount of effort has been devoted to developing successful regularizations [22, 8, 25] that take advantage of prior knowledge on network structures to enhance estimation performance for a variety of application settings, including image denoising [19] and genomic data analysis [15].

We consider the general setting of network regularization, now proposed as a convex optimization problem defined on an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with node set \mathcal{V} and edge set \mathcal{E} :

$$(1.1) \quad \min \sum_{i \in \mathcal{V}} f_i(\mathbf{x}_i) + \lambda \sum_{(j,k) \in \mathcal{E}} \omega_{jk} \cdot g(\mathbf{x}_j, \mathbf{x}_k),$$

where on each node i , we intend to learn a local model $\mathbf{x}_i \in \mathbb{R}^d$, which in addition to minimizing a predefined

convex loss function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, carries certain relations to its neighbors, defined by the function $g : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$.

In this paper, we focus on the particular setting where $g(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$, also known as a *sum-of-norms* (SON) regularizer. It assumes that the network is composed of multiple clusters, suggesting that all nodes within a cluster share the same consensus model ($\mathbf{x}_i = \dots = \mathbf{x}_j$). As the observed data on each node may be sparse, the SON regularizer allows nodes to “borrow” observations from their neighbors to improve their own models, as well as to determine the network cluster to which they belong. The SON objective was first introduced in [13] for convex clustering problems and used off-the-shelf sequential convex solvers. Chi & Lange [3] adapt SON clustering to incorporate similarity weights with an arbitrary norm and solve the problem by both alternating direction method of multipliers (ADMM) or alternating minimization algorithm (AMA) in a parallel manner. More recently, Hallac et al. [8] point out that weighted SON (named as Network Lasso) allows for simultaneous clustering and optimization on graphs, and is highly suitable for a broad class of large-scale network problems.

In network regularization, a static weight assigned to an edge determines how strictly the difference between models on the corresponding nodes is being penalized, relative to the other node pairs in the network. However, unlike few scenarios in which network information is explicitly given, edge weights are usually unavailable or even infeasible to obtain in most real-world networks (e.g., gene regulatory networks [17]). Moreover, due to possible measurement errors and inaccurate prior knowledge, the assigned edge weights don’t necessarily align with the underlying clustering structure of networks as shown in [1, 10]. As a result, heavily relying on the edge weights in determining how much penalty should be applied to neighboring models may contaminate the discovered solutions.

Similar intuition applies when we extend to the temporal setting. Models on nodes of temporal networks usually change at the level of groups over time. However, some groups exhibit different evolution patterns than others [14]. Moreover, the grouping structures themselves may evolve with time [12]. As static regularization cannot capture such temporal evolution, a direct solution is to induce a time-varying local consensus by employing the SON objective on both spatial and temporal directions, but we face an more drastic problem: how to assign weights between snapshots.

*Department of Computer Science, University of California, Santa Barbara ({hyou, furkan, ambuj}@ucsb.edu).

Consequently, a framework that is robust to missing or corrupted edge weights in network-based regularization and adapts to spatio-temporal settings is desired. Note that without the assigned network weights, the setting in Eq. (1.1) degrades to simply applying isotropic regularization for all the edges. In this work, we propose a generic formulation, called the *discrepancy-aware network regularization* (DANR), which deploys a suitable amount of anisotropic network regularization in both spatial and temporal aspects. DANR infers models at each node per timestamp and can learn evolution of models and transitions of network structures over time. We develop an ADMM-based algorithm that adopts an efficient and distributed iterative scheme to solve problems formulated by the DANR, and show that the proposed solution obtains guaranteed convergence towards global optimal solutions. By applying to both synthetic and real-world datasets, we demonstrate the effectiveness of the proposed approach on various network problems.

2 Problem Setting

2.1 Discrepancy-aware Network Regularization Since existing network-based regularizers rely on known weights on edges, the corresponding solutions get misled when the edge weights are erroneous or unknown. Directly learning the unknown edge weights by using the same regularizer as Eq. (1.1) would lead to an optimization problem:

$$(2.2) \quad \min_{\mathbf{x}, \boldsymbol{\omega}} \sum_{i \in \mathcal{V}} f_i(\mathbf{x}_i) + \lambda \sum_{(j,k) \in \mathcal{E}} \omega_{jk} \|\mathbf{x}_j - \mathbf{x}_k\|_2$$

which yields the trivial all-zero solution of $\boldsymbol{\omega}$ and thus becomes unsatisfactory. Instead of imposing more sophisticated model-based or problem-dependent regularization as suggested in [27], we consider an alternative formulation that explicitly accounts for discrepancies between models on adjacent nodes:

$$(2.3) \quad \min_{\mathbf{x}, \boldsymbol{\alpha}} \sum_{i \in \mathcal{V}} f_i(\mathbf{x}_i) + \lambda \cdot \mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha})$$

$$\mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha}) = \mu \sum_{(j,k) \in \mathcal{E}} \omega_{jk} \|\mathbf{x}_j + \boldsymbol{\alpha}_{jk} - \mathbf{x}_k\|_2 + (1-\mu) \|\boldsymbol{\alpha}\|_{1,p}$$

where we denote $\|\boldsymbol{\alpha}\|_{1,p} = \sum_{(j,k) \in \mathcal{E}} \|\boldsymbol{\alpha}_{jk}\|_p$ as the sum of p -norms of all $\boldsymbol{\alpha}_{jk}$'s. We set a penalty parameter $\lambda \geq 0$ to control the overall strength of network regularization, and a portion parameter $0 < \mu < 1$ to control the emphasis between the two terms in $\mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha})$. We name this formulation in Eq. (2.3) the *discrepancy-aware network regularization* (DANR).

In the first term of \mathcal{R}_S , we define a *discrepancy-buffering (DB) variable* $\boldsymbol{\alpha}_{jk} \in \mathbb{R}^d$ to denote the preserved discrepancy between local model parameters \mathbf{x}_i and \mathbf{x}_j . More specifically, when an edge weight ω_{jk} is given but potentially imprecise or otherwise corrupted, $\boldsymbol{\alpha}_{jk}$ would compensate for abnormally large differences between models

\mathbf{x}_j and \mathbf{x}_k to reduce the magnitude of ω_{jk} -weighted edge penalty term in \mathcal{R}_S . The DB variable provides additional flexibility to its associated models, allowing them to stay adequately close to their own local solutions w.r.t. minimizing local loss functions, and avoid over-penalized consensus. When all ω_{jk} 's are not given, $\boldsymbol{\alpha}_{jk}$'s enable solving Eq. 2.3 under anisotropic regularizations even with homogeneous weights.

The second term of \mathcal{R}_S is the $\ell_{1,p}$ -regularizer (with $1 < p < \infty$) [23] that ensures sparsity at the group level. Note that we regard each variable vector $\boldsymbol{\alpha}_{jk}$ as a group ($|\mathcal{E}|$ groups in total). The first key intuition here is that the regularization on $\boldsymbol{\alpha}$ helps to exclude trivial solutions, that is $\boldsymbol{\alpha}_{jk} = \mathbf{x}_k - \mathbf{x}_j$. Second, it allows us to identify a succinct set of non-zero vectors $\boldsymbol{\alpha}_{jk}$'s, compensating for possible intrinsic discrepancies between two adjacent nodes.

To sum up, discrepancy-buffering variable $\boldsymbol{\alpha}_{jk}$'s are designed to elaborately adjust network regularization strength on all edges, and thus reduce negative effects of the un-squared norm regularizer. We will show in later sections that this modified formulation remains convex and tractable via parallel optimization algorithms on large-scale networks.

2.2 Distributed ADMM-based Solution In this section, we propose an ADMM-based algorithm for solving the ST-DANR problem in Eq. (3.12) and present the convergence and complexity of the proposed algorithm. The algorithm can be easily adapted to the spatial-only DANR problem in Eq. (2.3) by omitting temporal-related updates.

The ADMM method was originally derived in [6] and has been reformulated in many contexts including optimal control and image processing [20]. The method can be considered as combining augmented Lagrangian methods and the method of multipliers [2]. It aims to solve optimization problems with two-block separable convex objectives in the following form of

$$(2.4) \quad \min_{\mathbf{x}, \mathbf{z}} f(\mathbf{x}) + g(\mathbf{z}) \quad \text{s.t. } A\mathbf{x} + B\mathbf{z} = \mathbf{c}$$

Observe that our proposed objective in Eq. (3.12) has a separable convex objective function: it can be reorganized into two-block separable convex objectives as in Eq. (2.4), for which ADMM methods guarantee convergence to global optimal solutions [5].

More precisely, to fit our problem into the ADMM framework, we first define $\mathbf{x} = \{\mathbf{x}_j\}_{j \in \mathcal{V}}$ for model parameter vectors on the given undirected network \mathcal{G} , and define consensus variables $\mathbf{u} = \{\mathbf{u}_{jk}, \mathbf{u}_{kj}\}_{(j,k) \in \mathcal{E}}$ as copies of \mathbf{x} in the spatial penalty term $\mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha})$. Then we rewrite Eq. (3.12) as follows:

$$(2.5) \quad \begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \boldsymbol{\alpha}} \quad & \sum_{j \in \mathcal{V}} f_j(\mathbf{x}_j) + \lambda_1 (1 - \mu_1) \|\boldsymbol{\alpha}\|_{1,p} \\ & + \lambda_1 \mu_1 \sum_{(j,k) \in \mathcal{E}} \omega_{jk} \|\mathbf{u}_{jk} + \boldsymbol{\alpha}_{jk} - \mathbf{u}_{kj}\|_2 \\ \text{s.t.} \quad & [\mathbf{u}_{jk}, \mathbf{u}_{kj}] = [\mathbf{x}_j, \mathbf{x}_k], \quad (j, k) \in \mathcal{E} \end{aligned}$$

in which the first term corresponds to the f block in Eq. (2.4), and the remaining terms correspond to the g block. The equality constraints on Eq. (3.12) are used to force consensus between variables \mathbf{x} and \mathbf{u} . Next, we derive the augmented Lagrangian of Eq. (2.5):

$$(2.6) \quad \mathcal{L}_{\rho_1}(\mathbf{x}, \mathbf{u}, \boldsymbol{\alpha}, \boldsymbol{\delta}) = \sum_{j \in V} f_j(\mathbf{x}_j) + \lambda_1(1 - \mu_1) \|\boldsymbol{\alpha}\|_{1,p} + \sum_{(j,k) \in E} \left(\lambda_1 \mu_1 \omega_{jk} \|\mathbf{u}_{jk} + \boldsymbol{\alpha}_{jk} - \mathbf{u}_{kj}\|_2 + \frac{\rho_1}{2} \|\mathbf{x}_j - \mathbf{u}_{jk} + \boldsymbol{\delta}_{jk}^u\|_2^2 - \frac{\rho_1}{2} \|\boldsymbol{\delta}_{jk}^u\|_2^2 + \frac{\rho_1}{2} \|\mathbf{x}_k - \mathbf{u}_{kj} + \boldsymbol{\delta}_{kj}^u\|_2^2 - \frac{\rho_1}{2} \|\boldsymbol{\delta}_{kj}^u\|_2^2 \right)$$

where $\boldsymbol{\delta} = [\{\boldsymbol{\delta}_{jk}^u, \boldsymbol{\delta}_{kj}^u\}_{(j,k) \in E}]$ are scaled dual variables for each equality constraint on the elements of \mathbf{u} . The parameters $\rho_1 > 0$ penalize the violation of equality constraints in the spatial and temporal domain [18] respectively. The iterative scheme of ADMM under the above setting can be written as follows, with l denoting the iteration index:

$$(2.7) \quad \left. \begin{aligned} \mathbf{x}^{(l+1)} &= \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{L}_{\rho_1}(\mathbf{x}, (\mathbf{u}, \boldsymbol{\alpha}, \boldsymbol{\delta})^{(l)}) \\ (\mathbf{u}, \boldsymbol{\alpha})^{(l+1)} &= \underset{\mathbf{u}, \boldsymbol{\alpha}}{\operatorname{argmin}} \mathcal{L}_{\rho_1}(\mathbf{x}^{(l+1)}, \mathbf{u}, \boldsymbol{\alpha}, \boldsymbol{\delta}^{(l)}) \\ \boldsymbol{\delta}^{(l+1)} &= \boldsymbol{\delta}^{(l)} + (\tilde{\mathbf{x}}^{(l+1)} - \mathbf{x}^{(l+1)}) \end{aligned} \right\}$$

where $\tilde{\mathbf{x}} = [\{\mathbf{x}_j, \mathbf{x}_k\}_{(j,k) \in E}]$ is composed of replicated elements in \mathbf{x} , and thus has a one-to-one correspondence with elements in \mathbf{u} .

Because the augmented Lagrangian in Eq. (2.6) has a separable structure as well, we can further split the optimization above over each univariate element in \mathbf{x} and \mathbf{z} . Next, we provide details for each ADMM update step.

\mathbf{x} -Update. In the first update step of our ADMM updating scheme, we can decompose the problem of minimizing \mathbf{x} into separately minimizing \mathbf{x}_j for each node j :

$$\mathbf{x}_j^{(l+1)} = \underset{\mathbf{x}_j}{\operatorname{argmin}} \left(f_j(\mathbf{x}_j) + \frac{\rho_1}{2} \sum_{k \in N_j} \|\mathbf{x}_j - \mathbf{u}_{jk}^{(l)} + \boldsymbol{\delta}_{jk}^{u(l)}\|_2^2 \right)$$

As above equation suggests, \mathbf{x}_j on node j first receives local information from the corresponding consensus variables belonging to all of its neighbors $k \in N_j$, then updates its value to minimize the loss function and remain close to neighbouring consensus variables. Since all remaining regularizations are quadratic, the \mathbf{x} -update problem can be efficiently solved whenever the loss function f_j has certain properties, such as strong convexity.

$(\mathbf{u}, \boldsymbol{\alpha})$ -Update. We can further decompose the $(\mathbf{z}, \boldsymbol{\alpha})$ -update step in Eq. (2.7) into subproblems on each edge (and its related variables $\mathbf{u}_{jk}, \mathbf{u}_{kj}, \boldsymbol{\alpha}_{jk}$) as follows, which can be solved in parallel:

$$(2.8) \quad (\mathbf{u}_{jk}^{(l+1)}, \mathbf{u}_{kj}^{(l+1)}, \boldsymbol{\alpha}_{jk}^{(l+1)}) = \underset{\mathbf{u}_{jk}, \mathbf{u}_{kj}, \boldsymbol{\alpha}_{jk}}{\operatorname{argmin}} \left(\lambda_1(1 - \mu_1) \|\boldsymbol{\alpha}_{jk}\|_p + \lambda_1 \mu_1 \omega_{jk} \|\mathbf{u}_{jk} + \boldsymbol{\alpha}_{jk} - \mathbf{u}_{kj}\|_2 + \frac{\rho_1}{2} \|\mathbf{x}_j^{(l+1)} - \mathbf{u}_{jk} + \boldsymbol{\delta}_{jk}^{u(l)}\|_2^2 + \frac{\rho_1}{2} \|\mathbf{x}_k^{(l+1)} - \mathbf{u}_{kj} + \boldsymbol{\delta}_{kj}^{u(l)}\|_2^2 \right)$$

Notice that the sum of convex functions which are defined on different sets of variables preserves the convexity. To minimize the objective with $\mathbf{u}_{jk}, \mathbf{u}_{kj}$ and $\boldsymbol{\alpha}_{jk}$ simultaneously, the convexity of Eq. (2.8) motivates us to adopt an alternating descent algorithm, which minimizes each component iteratively with respect to $(\mathbf{u}_{jk}, \mathbf{u}_{kj})$ and $\boldsymbol{\alpha}_{jk}$ while holding the other fixed. In detail, we solve Eq. (2.9) and Eq. (2.10) iteratively until convergence is achieved:

$$(2.9) \quad (\mathbf{u}_{jk}^{(l'+1)}, \mathbf{u}_{kj}^{(l'+1)}) = \underset{\mathbf{u}_{jk}, \mathbf{u}_{kj}}{\operatorname{argmin}} \left(\lambda_1 \mu_1 \omega_{jk} \|\mathbf{u}_{jk} + \boldsymbol{\alpha}_{jk}^{(l')}\|_2 - \|\mathbf{u}_{kj}\|_2 + \frac{\rho_1}{2} \|\mathbf{x}_j^{(l+1)} - \mathbf{u}_{jk} + \boldsymbol{\delta}_{jk}^{u(l)}\|_2^2 + \frac{\rho_1}{2} \|\mathbf{x}_k^{(l+1)} - \mathbf{u}_{kj} + \boldsymbol{\delta}_{kj}^{u(l)}\|_2^2 \right)$$

$$(2.10) \quad \boldsymbol{\alpha}_{jk}^{(l'+1)} = \underset{\boldsymbol{\alpha}_{jk}}{\operatorname{argmin}} \left((1 - \mu_1) \|\boldsymbol{\alpha}_{jk}\|_p + \mu_1 \omega_{jk} \|\mathbf{u}_{jk}^{(l'+1)} + \boldsymbol{\alpha}_{jk} - \mathbf{u}_{kj}^{(l'+1)}\|_2 \right)$$

where l' denoting the iteration index of alternating descent. Further, we can utilize the analytic solution to speed up the calculation of \mathbf{u}_{jk} and \mathbf{u}_{kj} :

LEMMA 2.1. *Problem (2.9) has a closed-form solution:*

$$\left. \begin{aligned} \mathbf{u}_{jk}^{(l'+1)} &= (1 - \theta) \mathbf{a} + \theta \mathbf{b} - \theta \boldsymbol{\alpha}_{jk}^{(l')} \\ \mathbf{u}_{kj}^{(l'+1)} &= \theta \mathbf{a} + (1 - \theta) \mathbf{b} + \theta \boldsymbol{\alpha}_{jk}^{(l')} \end{aligned} \right\}$$

where we denote $\mathbf{a} = \mathbf{x}_j^{(l'+1)} + \boldsymbol{\delta}_{jk}^{u(l')}, \mathbf{b} = \mathbf{x}_k^{(l'+1)} + \boldsymbol{\delta}_{kj}^{u(l')}, c = \lambda_1(1 - \mu_1) \omega_{jk}$, and $\theta = c / (\rho_1 \|\mathbf{a} - \mathbf{b} + \boldsymbol{\alpha}_{jk}^{(l')}\|_2)$ for simplification. Details in Appendix B.

$\boldsymbol{\delta}$ -Update. In the last step of our ADMM updating scheme, we have fully independent update rules for each scaled dual variable $\boldsymbol{\delta}_{jk}^u$ as follows:

$$(2.11) \quad \boldsymbol{\delta}_{jk}^{u(l+1)} = \boldsymbol{\delta}_{jk}^{u(l)} + (\mathbf{x}_j^{(l+1)} - \mathbf{u}_{jk}^{(l+1)})$$

Finally, we present the pseudo-code of our DANR approach in Appendix D.

Stopping Criterion and Global Convergence. For our ADMM iterative scheme in Eq. (2.7), we use the norm of primal residual $\mathbf{r}^{(l)} = \tilde{\mathbf{x}}^{(l)} - \mathbf{u}^{(l)}$ and dual residual $\mathbf{s}^{(l)} = \rho_1(\mathbf{u}^{(l)} - \mathbf{u}^{(l+1)})$ as the termination measure. The optimality condition [2] of ADMM shows that if both residuals are small then the objective suboptimality must be small, and thus suggests $\|\mathbf{r}^{(l)}\|_2 \leq \epsilon_1 \wedge \|\mathbf{s}^{(l)}\|_2 \leq \epsilon_2$ as a reasonable stopping criterion. Convex subproblems in \mathbf{x} -update and $(\mathbf{u}, \boldsymbol{\alpha})$ -update need iterative methods to solve as well. The stopping criterion for these subproblems is naturally to keep iteration differences below thresholds, *i.e.* in $(\mathbf{u}, \boldsymbol{\alpha})$ -update, we require $\|\Delta \mathbf{u}^{(l')}\|_2 \leq \epsilon'$ and $\|\Delta \boldsymbol{\alpha}^{(l')}\|_2 \leq \epsilon'$.

LEMMA 2.2. *Our ADMM approach to solve the DANR problem is guaranteed to converge to the global optimum. Details in Appendix C.*

Computational Complexity. Let N_c denote the number of iterations that ADMM takes to achieve an approximate solution $\hat{\mathbf{x}}$ with an accuracy of ϵ_x . Based on the convergence analysis in [11], the time complexity scales as $O(1/\epsilon_x)$, which in our problem mostly depends on the properties of cost functions $f_{j,t}$'s. Assume that all convex subproblems in x -update and (u, α) -update are solved by general first-order gradient descent methods that take N_x and N_α iterations to converge respectively, the overall complexity of the algorithm is therefore $O(N_c(N_x|\mathcal{V}| + N_\alpha|\mathcal{E}| + |\mathcal{E}|))$.

3 Extension to Spatio-temporal Setting

The aforementioned shortcomings with pre-defined edge weights accentuate when we consider temporal networks since acquiring explicit temporal edge weights is usually not feasible. Thus, discrepancy-buffering variables are also crucial for the temporal setting.

Consider a temporal undirected network \mathcal{G} consisting of M sequential network snapshots $\{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_M\}$, where each network snapshot \mathcal{G}_t contains a node set \mathcal{V}_t and an edge set \mathcal{E}_t . We denote $\omega_{j,k}^t$ as the weight of spatial edge (j_t, k_t) between nodes j_t and k_t within snapshot \mathcal{G}_t , and $\omega_j^{t,t+1}$ for the weight of temporal edge (j_t, j_{t+1}) that links node j_t with j_{t+1} across snapshots \mathcal{G}_t and \mathcal{G}_{t+1} (shown in Fig.1). On each node j_t , a convex loss function $f_{j,t} : \mathbb{R}^d \rightarrow \mathbb{R}$ is given to measure the fitness of a local model parameterized by $\mathbf{x}_{j,t} \in \mathbb{R}^d$. With sparse observations for all snapshots, a straightforward task is to find jointly optimal models for every node and every timestamp under the regularization for both network topology and temporal evolution. We can propose an extension of the DANR formulation to spatio-temporal networks, referred as ST-DANR:

$$(3.12) \quad \min_{\mathbf{x}, \boldsymbol{\alpha}, \boldsymbol{\beta}} \sum_{j \in \mathcal{V}} \sum_{t=1}^M f_{j,t}(\mathbf{x}_{j,t}) + \lambda_1 \cdot \mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha}) + \lambda_2 \cdot \mathcal{R}_T(\mathbf{x}, \boldsymbol{\beta})$$

where we define the discrepancy-aware regularizers as:

$$(3.13) \quad \begin{aligned} \mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha}) &= \mu_1 \sum_{t=1}^M \sum_{(j_t, k_t) \in \mathcal{E}_t} (\omega_{j,k}^t \|\mathbf{x}_{j,t} + \boldsymbol{\alpha}_{j,k,t} - \mathbf{x}_{k,t}\|_2) \\ &+ (1 - \mu_1) \|\boldsymbol{\alpha}\|_{1,p} \quad (\text{spatial penalty term}) \end{aligned}$$

$$(3.14) \quad \begin{aligned} \mathcal{R}_T(\mathbf{x}, \boldsymbol{\beta}) &= \mu_2 \sum_{j \in \mathcal{V}} \sum_{t=1}^{M-1} (\omega_j^{t,t+1} \|\mathbf{x}_{j,t} + \boldsymbol{\beta}_{j,t} - \mathbf{x}_{j,t+1}\|_2) \\ &+ (1 - \mu_2) \|\boldsymbol{\beta}\|_{1,p} \quad (\text{temporal penalty term}) \end{aligned}$$

As we are interested in the heterogeneous evolution of nodal models across time, the chosen unsquared norms in spatial and temporal regularization terms $\mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha})$ and $\mathcal{R}_T(\mathbf{x}, \boldsymbol{\beta})$ enforce piecewise consensus in both spatial and temporal aspects, which indicates abrupt changes of regional models or sudden transitions in network structures at particular timestamps, and also implies valid persistent models in the remaining segments of time [9]. Beside the spatial DB variables $\boldsymbol{\alpha}_{j,k,t}$ in $\mathcal{R}_S(\mathbf{x}, \boldsymbol{\alpha})$, we define another set of temporal DB variables $\boldsymbol{\beta}_{j,t} \in \mathbb{R}^d$ in $\mathcal{R}_T(\mathbf{x}, \boldsymbol{\beta})$ to compensate for

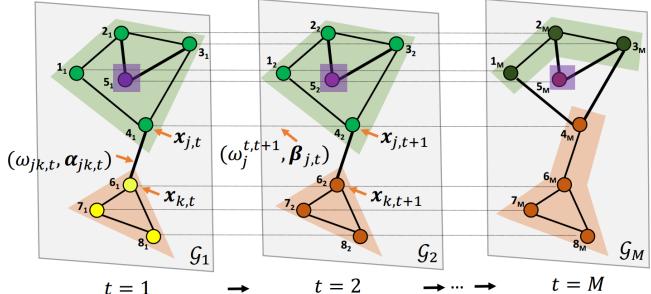


Figure 1: Overview of our problem setting in temporal networks.

inadequate regularizations caused by temporal fluctuations. For example, when modeling housing prices in a city, $\beta_{j,t}$ would tolerate anomalous short-term fluctuations in prices, while large values of $\alpha_{j,k,t}$ might be found near boundaries of disparate neighborhoods.

Adaptation to Streaming Data. In many applications, observed network snapshots arrive in a streaming fashion. Instead of recomputing all models on past snapshots whenever a new snapshot is observed, incorporating existing models to facilitate the new incoming snapshot is more efficient. Assume that we have learned models $\{\hat{\mathbf{x}}_{j,t}\}_{j \in \mathcal{V}}$ for the τ -th snapshot, we then read observations of the next m snapshots indexed by $\tau+1, \dots, \tau+m$ and attempt to learn models for them. To this intent, we deploy our ST-DANR formulation in Eq. (3.12) on a limited time interval $t = \tau, \dots, \tau+m$, and then integrate established model at $t = \tau$ by fixing $\{\mathbf{x}_{j,\tau}\}_{j \in \mathcal{V}} = \{\hat{\mathbf{x}}_{j,\tau}\}_{j \in \mathcal{V}}$ and solving the remaining variables $\{\mathbf{x}_{j,t}\}_{t=\tau+1, \dots, \tau+m}$ in the corresponding problem. Note that in our setting, temporal messages are only transferable through consecutive snapshots, meaning that fixing the τ -th snapshot is the same as fixing all past snapshots.

Lastly, we point out that distributed ADMM-based solution for DANR (§2.2) can be adapted to ST-DANR, for either batch formulation (Eq. (3.12)) or streaming-case formulation, which readers can refer to in Appendix E.

Temporal Evolutionary Patterns. In this work, we consider the unsquared edge penalty term $\|\mathbf{x}_{j,t} + \boldsymbol{\beta}_{j,t} - \mathbf{x}_{j,t+1}\|_2$ in \mathcal{R}_T since it enforces temporal reconstruction, indicating sharp transitions or change points at particular timestamps. Similar to the spatial case, there could be multiple available alternatives to enforce different types of evolutionary patterns [9]. For example, replacing the unsquared norm with squared norm form $\|\mathbf{x}_{j,t} + \boldsymbol{\beta}_{j,t} - \mathbf{x}_{j,t+1}\|_2^2$ will yield smoothly varying models along consecutive snapshots, and has been well studied in [24, 4].

4 EXPERIMENTS

In this section, we present an experimental analysis of the proposed method (DANR) and evaluate its performance on classification and regression tasks. First, we employ a synthetic dataset (§4.1) and two housing price datasets (BAY and SAC in §4.2) to demonstrate the set of scenarios and applica-

tions that are particularly good fit for DANR. In addition to an improvement on the estimation task, DANR learns fine-grained neighborhood boundaries and reveals interesting insights into the network's heterogeneous structure. Lastly, in order to validate the efficacy of ST-DANR in the temporal setting, we conduct experiments with the geospatial and temporal database of Northeastern US lakes (§4.3), from which we aim to estimate the water quality of lakes over 10 years.

Baseline Algorithms. In spatial network scenarios (§4.1 & §4.2), we compare DANR against three baselines: network lasso (NL) [8], robust multi-task feature learning (rMTFL) [7], factorized multi-task learning (FORMULA) [26]. In spatial-temporal network scenarios (§4.3), we compare ST-DANR with two widely-applied temporal regularizers in the literature, which are detailed in §4.3.

Parameters Setting. For both NL and DANR, we tune λ parameters from 10^{-3} to 10^2 where $\lambda^{n+1} = 1.3\lambda^n$. Concerning the DANR, for each value of λ , we further tune μ parameters from 0.3 to 1, where $\mu^{n+1} = \mu^n + 0.02$. Lastly, we set $p = 3$. We follow the same strategy for both spatial (DANR) and spatio-temporal (ST-DANR) variants of the proposed method. For all penalty parameters in rMTFL and FORMULA, we tune them in the same way as λ in DANR. In addition, we vary the number of factorized base models in FORMULA from 1 to 50 to achieve its best performance. For each dataset, we standardize all features and response variables to zero mean and unit variance.

4.1 Node Classification with Limited Data

Following previous work [8], we first experiment with a synthetic network in which each node attempts to solve its own support vector machine (SVM) classifier, but only has a limited amount of data to learn from. We further split the nodes into clusters where each cluster has an underlying model, i.e., nodes in the same cluster share the same underlying model. Our aim is to learn accurate models for each node by leveraging their network connections and limited observations. Note that the neighbors with different underlying models provide misleading information to each other, which potentially harms the overall performance. Therefore, we later investigate the robustness of the DANR with a varying ratio of such malicious edges.

Synthetic Network Generation. We create a network of 100 nodes, split into 5 ground-truth communities C^1, C^2, \dots, C^5 . Each of these communities consists of 20 nodes. We denote the community of a node i with C_i . Next, we form the network connections by adding edges between nodes based on the following criterion: The probability of adding an edge between any node pair (i, j) is 0.5 if $C_i = C_j$, i.e., the edge connects nodes from the same community (*intra-community edge*). Otherwise, we set the probability as 0.02 if $C_i \neq C_j$, i.e., the edge connects nodes from different communities (*inter-community edge*).

The resulting synthetic network $G_0(\mathcal{V}, \mathcal{E})$ has 100 nodes and 574 edges, with 82% of them being intra-community edges. Each ground-truth community C^k has an underlying classifier model $\mathbf{X}^k \in R^{10}$, drawn independently from a normal distribution of zero mean and unit variance, where $k \in [1, \dots, 5]$. Next, we generate 5 random training example pairs (\mathbf{w}, y) per node, where $\mathbf{w} \in R^{10}$ denotes an observation and $y \in \{-1, 1\}$ denotes the ground-truth value to estimate. The ground-truth value y for each observation \mathbf{w} is then computed using the underlying classifier of the community that a node belongs to:

$$(4.15) \quad y = \text{sign}(\mathbf{w}^T \mathbf{X}^k + \eta)$$

where $\mathbf{X}^k \in R^{10}$ represents the corresponding weight of the classifier w.r.t. observations, while η is the noise. All observations and noises are drawn independently from a normal distribution for each data point. Note that 5 observations are insufficient to accurately estimate a model $\mathbf{x}_i \in R^{10}$. Hence for this setting, the network connections play an important role as nodes can “borrow” training examples from its neighbors to improve their own models.

Optimization Problem. The objective function f_i on each node is formulated by the soft margin SVM objective, where node i estimates its model (i.e. separating hyperplane) $\mathbf{x}_i \in R^{10}$ using its five training examples (\mathbf{w}, y) as follows:

$$\min \left(\frac{1}{2} \mathbf{x}_i^T \mathbf{x}_i + C \sum_{l=1}^5 \xi_l \right) \quad \text{s.t. } y^{(l)} (\mathbf{x}_i^T \mathbf{w}^{(l)}) \geq (1 - \xi_l), \quad \xi_l \geq 0$$

where ξ_l 's are slack variables that accounts for non-separable data. The constant C controls the trade-off between minimizing the training error and maximizing the margin. We set C to 0.75, which we empirically find to perform well.

Test Results. In following subsections, we first thoroughly inspect DANR and its prototype method NL on synthetic network G_0 , in order to show the benefits of introducing discrepancy-buffering variable α . Later in §4.1.2, we evaluate the performance and robustness of DANR and all three baseline methods under more noisy scenarios. To assess the performance, prediction accuracies are computed over a separate set of 1000 test pairs (10 per node). The test pairs are again randomly generated using Eq. (4.15).

4.1.1 Effect of Discrepancy-Buffering Variables.

We compare DANR with NL and report the prediction accuracy with varying λ values in Figure 2. Recall that the proposed method introduces the μ parameter which is coupled with the λ parameter. Therefore in order to have an adequate comparison between DANR and NL, we modify our λ parameter to be $\tilde{\lambda} = \lambda/\mu$. Doing so ensures that both methods apply the same amount of penalty to their network regularization terms (See Eq. 2.2 and 2.3). For each value of λ , we vary μ from 0.3 to 1 and report the highest accuracy achieved.

We first briefly mention the desired behaviors of the NL method and later accentuate its drawbacks. As shown by Figure 2, for small λ values, the problem reduces to solving the local optimization problem where each node

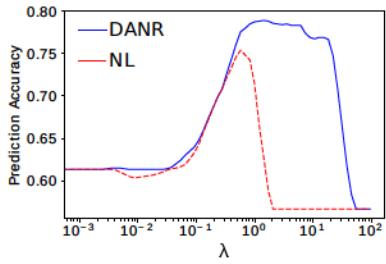


Figure 2: Prediction accuracy comparison with varying λ values.

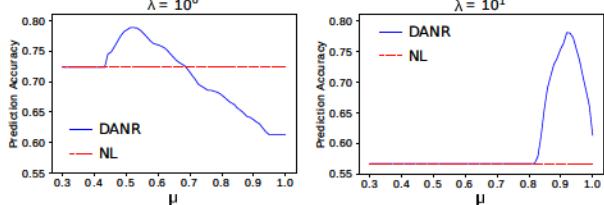


Figure 3: Effect of μ on the prediction accuracy, with a fixed λ .

estimates its model solely based on its limited observations. This achieves 61.3% accuracy on the test set. Furthermore, as λ increases, the NL penalty enforces nodes to form clusters, where nodes in the same cluster share the same model. This behavior firmly improves the test performance with prediction accuracy rising to 75.3%. Right after the peak performance is reached however, a slight increase in λ causes a sudden drop in the performance and the problem reduces to solving the global optimization problem over the entire network. Therefore when $\lambda > \lambda_{critical}$, the problem converges to a common model for all the nodes in the network, i.e., $x_i = x_j$ for $\forall i, j \in \mathcal{V}$. This further drops the prediction accuracy to 56.4% on the test set.

The observed rapid drop in performance implies that the NL method is highly sensitive to the λ parameter. There is only a narrow window of λ values that result in a proper clustering of the network. As a result, one needs to excessively tune the λ parameter around $\lambda_{critical}$ to find its optimal value. In addition, such clustering performance is also highly affected by the volume of “malicious” (inter-cluster) edges in the network, as shown in §4.1.2.

From Figure 2, we can observe that the DANR approach exhibits improved performance thanks to its discrepancy-buffering variables, α , which allow our regularization term to better exploit the good edges while providing additional tolerance towards bad edges. As a result, DANR achieves 79.1% accuracy on the test set, outperforming the best baseline method by 3.8%. Another important observation is that DANR provides a much wider window for selecting near-optimal λ than the baseline approach. That being said, we now analyze how the μ parameter couples with the λ parameter, and further present its effect on DANR’s performance. Figure 3 displays the prediction accuracy vs. μ , with a fixed $\lambda = 10^0$ and 10^1 . Intuitively, as λ increases, the optimal value of μ also increases. The reason behind is that as λ increases, the more non-zero discrepancy-buffering

variables (parameterized by $(1 - \mu)$) are needed to persist the accurate clustering formation since high values of λ forces global consensus over the network.

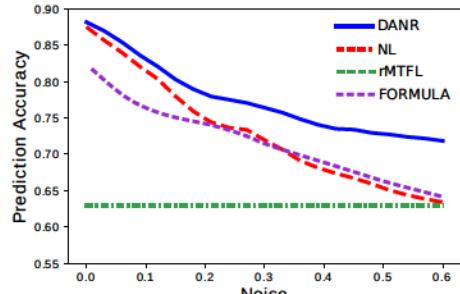


Figure 4: Prediction accuracy comparison with varying noise.

4.1.2 Robustness to Malicious Edges. Note that the ratio of inter-community edges in the earlier synthetic network G_0 is 18%. Here we use term *noise* for the ratio of such malicious edges in the network. Taking into account that the amount of noise is critical in learning accurate models for nodes in a network, we now investigate the efficacy of DANR w.r.t. varying noise. First, we remove all inter-community edges and later iteratively add malicious edges to the network. Then we solve the same problem with the noise varying from 0 to 0.6, and report prediction accuracies of all methods in Figure 4. As shown, the performance of DANR, NL and FORMULA drops in the absence of noise. Meanwhile the pure multitask method rMTFL maintains the same low accuracy, since it doesn’t utilize the network information. However, as we increase the noise, DANR starts to outperform NL and FORMULA, where the gap between the models’ performances becomes larger at a higher ratio of malicious edges. This confirms that compared to the baseline methods, the DANR exhibits more robust performance in noisy settings thanks to its discrepancy-buffering variables.

4.2 Spatial: Housing Rental Price Estimation In this section, our goal is to jointly (i) estimate the rental prices of houses by leveraging their geological location and the set of features; (ii) discover boundaries between neighborhoods. The intuition is that the houses in the same neighborhood often tend to have similar pricing models. However, such neighborhoods can have complex underlying structures (as described later in this section), which imposes additional challenges in learning accurate models.

Dataset. We experiment with two largely populated areas in Northern California: the Greater Sacramento Area (SAC) and the Bay Area (BAY). The anonymized data is provided by the property management software company *Appfolio Inc*, from which we further sample houses with at least one signed rental agreement during the year 2017. The resulting dataset covers 3849 houses in SAC and 1498 houses in BAY. Concerning the houses that have more than one rental agreement signed during 2017, we average the rental

prices listed in all the agreements. Each house holds the information about its location (latitude/longitude), number of bedrooms, number of bathrooms, square footage, and the rental price. We regard these areas (SAC and BAY) as two separate datasets for the remainder of this section. We also randomly split 20 % of the houses in each dataset for testing.

Network Construction. After excluding test houses from both datasets, we construct two networks (one for each dataset) based on the houses' locations. An undirected edge exists between node i and j , if at least one of them is in the set of 10 nearest neighbors of the other. (Note that node j being one of the 10 nearest neighbors of i doesn't necessarily imply that node i is also in the set of nearest neighbors of j .) In the resulting networks, the average degree of a node is 12.16 for SAC and 12.04 for BAY. Additionally, we construct two versions of these networks, *weighted* and *unweighted*. While the weighted network has edge weights inversely proportional to the distances between houses, the unweighted network ignores the proximity between houses, and thus weights on all the edges are 1.

Optimization Problem. The model at each house estimates its rental price by solving a linear regression problem. More specifically, at each node i we learn a 4-dimensional model $\mathbf{x}_i = [a_i, b_i, c_i, d_i]$. \mathbf{x} simply represents the coefficients of each feature, which later is used to estimate the rental price p_i as follows:

$$\begin{aligned} \bar{p}_i = a_i \cdot (\#bedrooms) + b_i \cdot (\#bathrooms) \\ + c_i \cdot (\text{squarefootage}) + d_i, \end{aligned}$$

where d_i is the bias term. The training objective is

$$\min_{\mathbf{x}, \alpha} \sum_{i \in \mathcal{V}} \|\bar{p}_i - p_i\|_2^2 + c \cdot \|\mathbf{x}_i\|_2^2 + \lambda \cdot \mathcal{R}_S(\mathbf{x}, \alpha)$$

where \mathcal{R}_S represents the proposed network regularization term (see Eq. 2.3), while c is the regularization weight to prevent over-fitting.

Test Results. Once training converges, we predict the rental prices on the test set. To do that, we connect each house in the test set to its 10 nearest neighbors in the training set. We then infer the new model \mathbf{x}_j by taking the average of the models on its neighbors: $\mathbf{x}_j = (1/|N(j)|) \sum_{k \in N(j)} \mathbf{x}_k$. The model \mathbf{x}_j is further used to estimate the rental price of

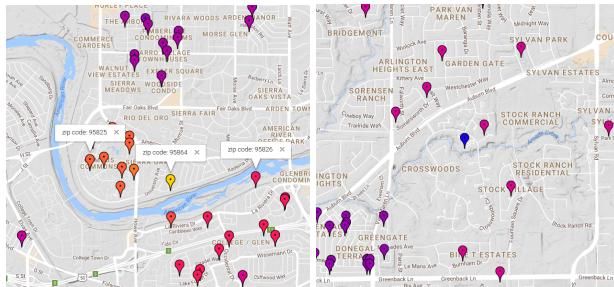


Figure 5: Examples of complex neighborhood structures captured by the DANR. In left, the house shown in yellow resides at the border of three different area codes. In right, the creek side house (colored in blue) differs from all of its neighbors, possibly due to its appealing location.

the corresponding house. Alternatively, one can also infer \mathbf{x}_j by solving $\min_{\mathbf{x}_j} \sum_{k \in N(j)} \|\mathbf{x}_j - \mathbf{x}_k^*\|_2$, while keeping the models on neighbors fixed. However, we empirically find that simply averaging the neighbors' models performs better for both methods in this particular setting. We compute Mean Squared Error (MSE) over test nodes to evaluate the performance.

Method	BAY	SAC
Local estimation ($\lambda = 0$)	0.5984	0.6250
Global estimation ($\lambda > \lambda_{critical}$)	0.4951	0.5403
rMTFL	0.4774	0.4115
FORMULA (unweighted)	0.4446	0.3503
FORMULA (weighted)	0.4181	0.3379
Network Lasso (unweighted)	0.4392	0.3273
Network Lasso (weighted)	0.4173	0.3022
DANR (unweighted)	0.4106	0.2978

Table 1: MSE for housing rental price prediction on test set.

Table 1 displays the test results of eight different settings on both datasets. As shown, local & global estimations and rMTFL method produce high errors for both datasets. We further apply FORMULA and Network Lasso (NL) methods to both weighted and unweighted versions of the networks, while the proposed method is only applied to the unweighted version. Notice that the network weights are irrelevant for the local & global estimation settings and the rMTFL method. Intuitively, both FORMULA and NL performs better on the weighted setting compared to the unweighted setting for both datasets. As shown, the rental price estimation errors achieved by the NL are 0.4173 (weighted) vs. 0.4392 (unweighted) for BAY and 0.3022 (weighted) vs. 0.3273 (unweighted) for SAC. These results suggest that the pre-defined weights on these networks help to learn more accurate models on houses.

However, we further argue that although such pre-defined weights improve the overall clustering performance, they don't account for more complex clustering scenarios, e.g., two nearby houses falling into different school districts or some houses having higher values compared to their neighbors due to geography, e.g., having a view of the city. That being said, DANR outperforms all the other baselines and achieves the smallest errors for both datasets; 0.4106 for BAY and 0.2978 for SAC. Especially, the DANR (unweighted) even outperforms the weighted version of the baseline approaches by a notable margin. This reveals that the DANR indeed accounts for such heterogeneities in data and provides enhanced clustering of the network.

Figure 5 shows examples of two complex scenarios that are captured by the DANR from the SAC network. We use a color code to represent the clusters in the network, where the same colored houses (i, j) are in consensus on their models ($\mathbf{x}_i = \mathbf{x}_j$). In the left subfigure, the house shown in yellow uses a different model than all of its neighbors. Interestingly,

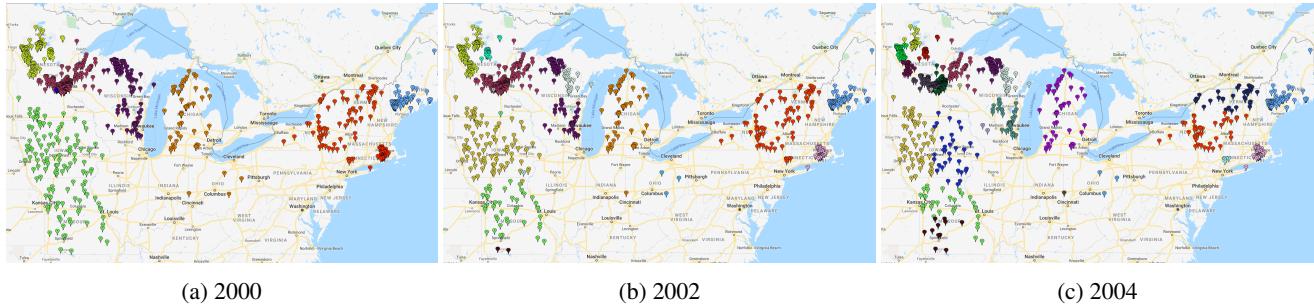


Figure 6: Evolution of clustering captured by the ST-DANR over years.

it resides at the border of three different area codes. The area code for this house is 95864, while the area code on its west is 95825 and 95826 on its south-east. Additionally, we observe similar heterogeneous behaviors in some houses that are near creeks, lakes, and rivers. As an example, Figure 5 (right) displays a creek side house (colored in blue), which again uses a different underlying model from its neighbors.

4.3 Spatio-Temporal: Water Quality Estimation of Northeastern US Lakes We now evaluate our method in spatio-temporal setting, where the aim is to dynamically estimate the water quality of Northeastern US lakes over years. We follow the same procedure in §4.2, but have an additional temporal regularization term in our objective that allows nodes to obtain signals from past snapshots of the network, along with their neighbors in the current snapshot.

Dataset and Network. We experiment with the geospatial and temporal dataset of lakes in 17 states in the Northeastern United States, called LAGOS [21]. The dataset holds extensive information about the physical characteristics, various nutrient concentration measurements (water quality), ecological context (surrounding land use/cover, climate etc.), and location of lakes; from which we select the following available set of features: $\{\text{max depth, surface area, elevation, annual mean temperature, \% of agricultural land, \% of urban land, \% of forest land, \% of wetland}\}$. The water quality metric to estimate is the summerly mean *chlorophyll concentration* of the lakes. We represent the feature vector with $\mathbf{w} \in \mathbb{R}^8$ and the water quality score with $y \in \mathbb{R}$.

During our experiments, we consider a 10-year period between 2000 and 2010. Due to sparsity in the data, we allow 2 years range between the two consecutive snapshots of the network. This results in total of 1039 lakes with all the aforementioned features and the water quality measurements available for each of the selected years. After randomly splitting 20% of the lakes for testing, we build our network by using the latitude/longitude information of lakes, where each lake (node) is connected to 10 nearest lakes.

Optimization Problem. After constructing the network, we now aim to *dynamically* estimate the water quality ($y_{i,t}$) of lakes by using their feature vectors ($\mathbf{w}_{i,t}$). More formally, for each year $t \in [2000, 2002, \dots, 2010]$, we learn a model $\mathbf{x}_{i,t} \in \mathbb{R}^8$ per node by solving the following spatio-

temporal problem in a *streaming* fashion: (4.16)

$$\min_{\mathbf{x}, \alpha, \beta} \sum_{i \in \mathcal{V}} f(\mathbf{x}_{i,t}, \mathbf{w}_{i,t}, y_{i,t}) + \lambda_1 \cdot \mathcal{R}_S^t(\mathbf{x}, \alpha) + \lambda_2 \cdot \mathcal{R}_T^t(\mathbf{x}, \beta)$$

where f is the linear regression objective and \mathcal{R}_S is the DANR term applied on the spatial edges. \mathcal{R}_T is the temporal regularization term, for which we consider two formulations as described next.

Test results. For each year, we first solve the spatial problem (Eq. (4.16) without the temporal term \mathcal{R}_T) and report the Mean Squared Errors in Table 2. Note that the test nodes are again inferred by averaging the neighbors' models in the current snapshot. Later, in order to successfully assess the improvements gained by the temporal discrepancy-buffering variables (β), we solve the above spatio-temporal problem with three different temporal regularizers:

DANR+T-SON: DANR with temporal sum-of-norms regularizer where $\mathcal{R}_T^t(\mathbf{x}, \cdot) = \sum_{i \in \mathcal{V}} \|\mathbf{x}_{i,t} - \hat{\mathbf{x}}_{i,t-1}\|_2$

DANR+T-SOS: DANR with temporal sum-of-squares regularizer where $\mathcal{R}_T^t(\mathbf{x}, \cdot) = \sum_{i \in \mathcal{V}} \|\mathbf{x}_{i,t} - \hat{\mathbf{x}}_{i,t-1}\|_2^2$

ST-DANR: Spatio-temporal discrepancy-aware network regularizer where $\mathcal{R}_T^t(\mathbf{x}, \beta) = \mu_2 \cdot \sum_{i \in \mathcal{V}} \|\mathbf{x}_{i,t} - \hat{\mathbf{x}}_{i,t-1} + \beta_{i,t}\|_2 + (1 - \mu_2) \cdot \|\beta\|_1$

DANR+T-SON and DANR+T-SOS approaches simply apply two widely adopted temporal regularizers on the temporal edges (the sum-of-norms regularizer [16] and sum-of-squares regularizer [4, 24] respectively), while ST-DANR includes the discrepancy-buffering variables on both spatial and temporal edges. Recall that we solve the Eq. (4.16) in a streaming fashion for simplicity, i.e., each snapshot of the network solves the problem while holding the models learned on the previous snapshot (if available) fixed. Potentially, the performance can further be improved by allowing updates on the previous snapshots.

As we can see from the Table 2, leveraging the temporal connections between the two consecutive snapshots significantly improves the performance. The DANR+T-SON outperforms the DANR by 8%-14%, while DANR+TSOS outperforms the DANR by 10%-15%. Moreover, the ST-DANR consistently outperforms both baselines for all the years. This confirms that the proposed formulation accounts for the heterogeneous nature of the temporal networks where some group of nodes exhibits different evolution patterns than the others.

Method	2000	2002	2004	2006	2008	2010
DANR	0.8479	0.9033	0.6384	0.6722	0.4556	0.4113
+ T-SON	N/A	0.8308	0.5744	0.6061	0.4109	0.3517
+ T-SOS	N/A	0.8045	0.5618	0.6045	0.3978	0.3476
ST-DANR	N/A	0.8037	0.5604	0.5866	0.3844	0.3311

Table 2: MSE for water quality estimation over years.

Figure 6 further displays the models (color-coded) learned on three consecutive snapshots, corresponding to years 2000, 2002 and 2004. In 2000, the formed clusters don't go much beyond the boundaries of the states, resulting in coarse clustering of the network. Yet some states such as Missouri and Iowa are grouped into one cluster (shown in green). This suggests that accurately estimating the chlorophyll concentration of lakes based on the selected set of features is indeed challenging. Specifically, the estimation problem depends on several other external factors that potentially affect the volume of plants and algae in lakes; cultural eutrophication, nutrient inputs from human activities to name a few [21]. However, as it can be seen from the models learned in 2002 and 2004, the clustering performance improves once we allow temporal regularization to synchronize models between two consecutive snapshots, which is analogous with the reduction in errors over years as reported in Table 2. Overall, we observe a split of clusters over time, e.g., in Wisconsin, Minnesota, Iowa and New England. Additionally, a dark brown cluster in south Missouri begins to appear in 2002 and further spreads north in 2004. This indicates that by leveraging the temporal connections, the ST-DANR learns improved models and clustering while allowing for heterogeneity in group level evolutions of nodes.

5 CONCLUSIONS

We propose discrepancy-aware network regularization (DANR) approach for spatio-temporal networks. By introducing discrepancy-buffering variables, the DANR automatically compensates for inaccurate prior knowledge encoded in edge weights, and enables modeling heterogeneous temporal patterns of network clusters. We develop a scalable algorithm based on alternating direction method of multipliers (ADMM) to solve the proposed problem, which enjoys analytic solutions for decomposed subproblems and employs a distributed update scheme on nodes and edges. Our experimental results show that DANR successfully captures structural changes over spatio-temporal networks and yields enhanced models by providing robustness towards missing or inaccurate edge weights.

6 ACKNOWLEDGMENTS

The work was supported in part by the National Science Foundation under grant IIS-1817046.

References

- [1] L. ANSELIN, *Issues in the specification and interpretation of spatial regression models*, Agricultural economics, 27 (2002).
- [2] S. BOYD ET AL., *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Foundations and Trends® in Machine Learning, 3 (2011).
- [3] E. C. CHI ET AL., *Splitting methods for convex clustering*, Journal of Computational and Graphical Statistics, 24 (2015).
- [4] X.-H. DANG ET AL., *Subnetwork mining with spatial and temporal smoothness*, in SDM, 2017.
- [5] J. ECKSTEIN ET AL., *On the douglasrachford splitting method and the proximal point algorithm for maximal monotone operators*, Mathematical Programming, 55 (1992).
- [6] D. GABAY AND B. MERCIER, *A dual algorithm for the solution of nonlinear variational problems via finite element approximation*, Computers & Mathematics with Applications.
- [7] P. GONG, J. YE, AND C. ZHANG, *Robust multi-task feature learning*, in ACM SIGKDD, 2012.
- [8] D. HALLAC ET AL., *Network lasso: Clustering and optimization in large graphs*, in ACM SIGKDD, 2015.
- [9] D. HALLAC, Y. PARK, S. BOYD, AND J. LESKOVEC, *Network inference via the time-varying graphical lasso*, 2017.
- [10] R. HARRIS, J. MOFFAT, AND V. KRAVTSOVA, *In search of 'w'*, Spatial Economic Analysis, 6 (2011), pp. 249–270.
- [11] B. HE AND X. YUAN, *On non-ergodic convergence rate of douglas–rachford alternating direction method of multipliers*, Numerische Mathematik, 130 (2015), pp. 567–577.
- [12] Q. HO ET AL., *Evolving cluster mixed-membership block-model for time-evolving networks*, in AISTATS, 2011.
- [13] T. D. HOCKING ET AL., *Clusterpath an algorithm for clustering using convex fusion penalties*, in ICML, 2011.
- [14] M. KIM AND J. LESKOVEC, *Nonparametric multi-group membership model for dynamic networks*, in NIPS, 2013.
- [15] C. LI AND H. LI, *Network-constrained regularization and variable selection for analysis of genomic data*, Bioinformatics, 24 (2008), pp. 1175–1182.
- [16] F. LINDSTEN ET AL., *Clustering using sum-of-norms regularization*, in Statistical Signal Processing Workshop, 2011.
- [17] Y.-Y. LIU, J.-J. SLOTINE, AND A.-L. BARABÁSI, *Controllability of complex networks*, Nature, 473 (2011), p. 167.
- [18] R. NISHIHARA ET AL., *A general analysis of the convergence of admm*, in ICML, 2015, pp. 343–352.
- [19] J. PANG ET AL., *Graph laplacian regularization for image denoising*, IEEE Transactions on Image Processing, (2017).
- [20] Y. PENG ET AL., *Robust alignment by sparse and low-rank decomposition for linearly correlated images*, PAMI, (2012).
- [21] P. A. SORANNO ET AL., *Lagos-ne*, GigaScience, 6 (2017).
- [22] R. TIBSHIRANI ET AL., *Sparsity and smoothness via the fused lasso*, Journal of the Royal Statistical Society, (2005).
- [23] J. E. VOGT AND V. ROTH, *A complete analysis of the $\ell_{1,p}$ group-lasso*, in ICML, 2012.
- [24] H. WANG, F. NIE, AND H. HUANG, *Low-rank tensor completion with spatio-temporal consistency*, in AAAI, 2014.
- [25] Y.-X. WANG, J. SHARPNACK, A. J. SMOLA, AND R. J. TIBSHIRANI, *Trend filtering on graphs*, in AISTATS, 2015.
- [26] J. XU ET AL., *Factorized multi-task learning for task discovery in personalized medical models*, in SDM, 2015.
- [27] H.-F. YU ET AL., *Temporal regularized matrix factorization for high-dimensional time series prediction*, in NIPS, 2016.