This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tran, Nguyen; Ambos, Henrik; Jung, Alexander

Classifying Partially Labeled Networked Data VIA Logistic Network Lasso

Published in:

2020 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2020 - Proceedings

DOI:

[10.1109/ICASSP40776.2020.9054408](https://doi.org/10.1109/ICASSP40776.2020.9054408)

Published: 01/05/2020

Document Version

Peer reviewed version

Please cite the original version:

Tran, N., Ambos, H., & Jung, A. (2020). Classifying Partially Labeled Networked Data VIA Logistic Network Lasso. In *2020 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2020 - Proceedings* (pp. 3832-3836). [9054408] (ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings; Vol. 2020-May). IEEE. <https://doi.org/10.1109/ICASSP40776.2020.9054408>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

CLASSIFYING PARTIALLY LABELED NETWORKED DATA VIA LOGISTIC NETWORK LASSO

Nguyen Tran, Henrik Ambos and Alexander Jung

Department of Computer Science, Aalto University, Finland; firstname.lastname(at)aalto.fi

ABSTRACT

We apply the network Lasso to classify partially labeled data points which are characterized by high-dimensional feature vectors. In order to learn an accurate classifier from limited amounts of labeled data, we borrow statistical strength, via an intrinsic network structure, across the dataset. The resulting logistic network Lasso amounts to a regularized empirical risk minimization problem using the total variation of a classifier as a regularizer. This minimization problem is a non-smooth convex optimization problem which we solve using a primal-dual splitting method. This method is appealing for big data applications as it can be implemented as a highly scalable message passing algorithm.

1. INTRODUCTION

The *least absolute shrinkage and selection operator* (Lasso) has been extended to networked data recently. This extension, coined the “network Lasso” (nLasso), allows efficient processing of massive datasets using convex optimization methods [1].

Most of the existing work on nLasso-based methods focuses on predicting numeric labels (or target variables) within regression problems [1–8]. In contrast, we apply nLasso to binary classification problems which assign binary-valued labels to data points [9–12].

In order to learn a classifier from partially labeled networked data, we minimize the logistic loss incurred on a training set constituted by few labeled nodes. Moreover, we aim at learning classifiers which conform to the intrinsic network structure of the data. In particular, we require classifiers to be approximately constant over well-connected subsets (clusters) of data points. This cluster assumption lends naturally to regularized empirical risk minimization with the total variation of the classifier as regularization term [13]. We solve this non-smooth convex optimization problem by applying the primal-dual method proposed in [14, 15].

The proposed classification method extends the toolbox for semi-supervised classification in networked data [12, 16–20]. In contrast to label propagation (LP), which is based on the squared error loss, we use the logistic loss which is more suitable for classification problems. Another important

difference between LP methods and nLasso is the different choice of regularizer. Indeed, LP uses the Laplacian quadratic form while the nLasso uses total variation for regularization.

Using a (probabilistic) stochastic block model for networked data, a semi-supervised classification method is obtained as an instance of belief propagation method for inference in graphical models [17]. In contrast, we assume the data (network-) structure as fixed and known. The proposed method provides a statistically well-founded alternative to graph-cut methods [18–20]. While our approach is based on a convex optimization (allowing for highly scalable implementation), graph-cuts is based on combinatorial optimization which makes scaling them to large datasets more challenging. Moreover, while graph-cut methods apply only to data which is characterized by network structure and labels, our method allows to exploit additional information provided by feature vectors of data points.

Contribution: Our main contributions are: (i) We present a novel implementation of logistic network Lasso by applying a primal-dual method. This method can be implemented as highly scalable message passing on the network structure underlying the data. (ii) We prove the convergence of this primal-dual method and (iii) verify its performance on synthetic classification problems in chain and grid-structured data.

Notation: Boldface lowercase (uppercase) letters denote vectors (matrices). We denote $\mathbf{x}^T$ the transpose of vector $\mathbf{x}$. The ℓ_2 -norm of a vector $\mathbf{x}$ is $\|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{x}}$. The convex conjugate of a function f is defined as $f^*(\mathbf{y}) = \sup_{\mathbf{x}} (\mathbf{y}^T \mathbf{x} - f(\mathbf{x}))$. We also need the sigmoid function $\sigma(z) := 1/(1 + \exp(-z))$.

2. PROBLEM FORMULATION

We consider networked data that is represented by an undirected weighted graph (the “empirical graph”) $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$. A particular node $i \in \mathcal{V} = \{1, \dots, N\}$ of the graph represents an individual data point (such as a document, or a social network user profile).¹ Two different data points $i, j \in \mathcal{V}$ are connected by an undirected edge $\{i, j\} \in \mathcal{E}$ if they are consid-

¹With a slight abuse of notation, we refer by $i \in \mathcal{V}$ to a node of the empirical graph as well as the data point which is represented by that node.

ered similar (such as documents authored by the same person or social network profiles of befriended users). For ease of notation, we denote the edge set $\mathcal{E}$ by $\{1, \dots, E := |\mathcal{E}|\}$.

Each edge $e = \{i, j\} \in \mathcal{E}$ is endowed with a positive weight $A_e = A_{ij} > 0$ which quantifies the amount of similarity between data points $i, j \in \mathcal{V}$. The neighborhood of a node $i \in \mathcal{V}$ is $\mathcal{N}(i) := \{j : \{i, j\} \in \mathcal{E}\}$.

Beside the network structure, datasets convey additional information in the form of features $\mathbf{x}^{(i)} \in \mathbb{R}^d$ and labels $y^{(i)} \in \{-1, 1\}$ associated with each data point $i \in \mathcal{V}$. In what follows, we assume the features to be normalized such that $\|\mathbf{x}^{(i)}\| = 1$ for each data points $i \in \mathcal{V}$. While features are typically available for each data point $i \in \mathcal{V}$, labels are costly to acquire and available only for data points in a small training set $\mathcal{M} = \{i_1, \dots, i_M\}$ containing M labeled data points.

We model the labels $y^{(i)}$ of the data points $i \in \mathcal{V}$ as independent random variables with (unknown) probabilities

$$p^{(i)} := \text{Prob}\{y^{(i)} = 1\} = \frac{1}{1 + \exp(-(\mathbf{w}^{(i)})^T \mathbf{x}^{(i)})}. \quad (1)$$

The probabilities $\{p^{(i)}\}_{i \in \mathcal{V}}$ are parametrized by some (unknown) weight vectors $\mathbf{w}^{(i)}$. Our goal is to develop a method for learning an accurate estimate $\hat{\mathbf{w}}^{(i)}$ of the weight vector $\mathbf{w}^{(i)}$. Given the estimate $\hat{\mathbf{w}}^{(i)}$, we can compute an estimate $\hat{p}^{(i)}$ for the probability $p^{(i)}$ by replacing $\mathbf{w}^{(i)}$ with $\hat{\mathbf{w}}^{(i)}$ in (1).

We interpret the weight vectors as the values of a graph signal $\mathbf{w} : \mathcal{V} \rightarrow \mathbb{R}^d$ assigning each node $i \in \mathcal{V}$ of the empirical graph $\mathcal{G}$ the vector $\mathbf{w}^{(i)} \in \mathbb{R}^d$. The set of all vector-valued graph signals is denoted $\mathcal{C} := \{\mathbf{w} : \mathcal{V} \rightarrow \mathbb{R}^d : i \mapsto \mathbf{w}^{(i)}\}$.

Each graph signal $\hat{\mathbf{w}} \in \mathcal{C}$ defines a classifier which maps a node with features $\mathbf{x}^{(i)}$ to the predicted label

$$\hat{y}^{(i)} = \begin{cases} 1 & \text{if } (\hat{\mathbf{w}}^{(i)})^T \mathbf{x}^{(i)} > 0 \\ -1 & \text{otherwise.} \end{cases} \quad (2)$$

Given partially labeled networked data, we aim at learning a classifier $\hat{\mathbf{w}} \in \mathcal{C}$ which agrees with the labels $y^{(i)}$ of labeled data points in the training set $\mathcal{M}$. In particular, we aim at learning a classifier having a small training error

$$\hat{E}(\hat{\mathbf{w}}) := (1/M) \sum_{i \in \mathcal{M}} \ell((\hat{\mathbf{w}}^{(i)})^T \tilde{\mathbf{x}}^{(i)}) \quad (3)$$

with $\tilde{\mathbf{x}}^{(i)} := y^{(i)} \mathbf{x}^{(i)}$ and the logistic loss

$$\ell(z) := \log(1 + \exp(-z)) = -\log(\sigma(z)). \quad (4)$$

3. LOGISTIC NETWORK LASSO

The criterion (3) by itself is not enough for guiding the learning of a classifier $\mathbf{w}$ since (3) completely ignores the weights $\hat{\mathbf{w}}^{(i)}$ at unlabeled nodes $i \in \mathcal{V} \setminus \mathcal{M}$. Therefore, we need to impose some additional structure on the classifier $\hat{\mathbf{w}}$. In particular, any reasonable classifier $\hat{\mathbf{w}}$ should conform with the *cluster structure* of the empirical graph $\mathcal{G}$ [21].

We measure the extend of a classifier $\hat{\mathbf{w}} \in \mathcal{C}$ conforming with the cluster structure of $\mathcal{G}$ by the total variation (TV)

$$\|\mathbf{w}\|_{\text{TV}} := \sum_{\{i,j\} \in \mathcal{E}} A_{ij} \|\mathbf{w}^{(j)} - \mathbf{w}^{(i)}\|. \quad (5)$$

A classifier $\hat{\mathbf{w}} \in \mathcal{C}$ has small TV if the weights $\hat{\mathbf{w}}^{(i)}$ are approximately constant over well connected subsets (clusters) of nodes.

We are led quite naturally to learning a classifier $\hat{\mathbf{w}}$ via the *regularized empirical risk minimization* (ERM)

$$\hat{\mathbf{w}} \in \underset{\mathbf{w} \in \mathcal{C}}{\text{argmin}} \hat{E}(\mathbf{w}) + \lambda \|\mathbf{w}\|_{\text{TV}}. \quad (6)$$

We refer to (6) as the logistic nLasso (lnLasso) problem. The parameter λ in (6) allows to trade-off small TV $\|\hat{\mathbf{w}}\|_{\text{TV}}$ against small error $\hat{E}(\hat{\mathbf{w}})$ (cf. (3)). The choice of λ can be guided by cross validation [22].

Note that lnLasso (6) does not enforce directly the labels $y^{(i)}$ to be clustered. Instead, it requires the classifier $\hat{\mathbf{w}}$, which parametrizes the probability distributed of the labels $y^{(i)}$ (see (1)), to be clustered.

It will be convenient to reformulate (6) using vector notation. We represent a graph signal $\mathbf{w} \in \mathcal{C}$ as the vector

$$\mathbf{w} = ((\mathbf{w}^{(1)})^T, \dots, (\mathbf{w}^{(N)})^T)^T \in \mathbb{R}^{dN}. \quad (7)$$

Define a partitioned matrix $\mathbf{D} \in \mathbb{R}^{(dE) \times (dN)}$ block-wise as

$$\mathbf{D}_{e,i} = \begin{cases} A_{ij} \mathbf{I}_d & e = \{i, j\}, i < j \\ -A_{ij} \mathbf{I}_d & e = \{i, j\}, i > j \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad (8)$$

where $\mathbf{I}_d \in \mathbb{R}^{d \times d}$ is the identity matrix. The term $A_{ij}(\mathbf{w}^{(i)} - \mathbf{w}^{(j)})$ in (5) is the e -th block of $\mathbf{D}\mathbf{w}$. Using (7) and (8), we can reformulate the lnLasso (6) as

$$\hat{\mathbf{w}} \in \underset{\mathbf{w} \in \mathbb{R}^{dN}}{\text{argmin}} h(\mathbf{w}) + g(\mathbf{D}\mathbf{w}), \quad (9)$$

with

$$h(\mathbf{w}) = \hat{E}(\mathbf{w}) \text{ and } g(\mathbf{u}) := \lambda \sum_{e=1}^E \|\mathbf{u}^{(e)}\| \quad (10)$$

with stacked vector $\mathbf{u} = (\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(E)}) \in \mathbb{R}^{dE}$.

4. PRIMAL-DUAL METHOD

The lnLasso (9) is a convex optimization problem with a non-smooth objective function which rules out the use of gradient descent methods [23]. However, the objective function is highly structured since it is the sum of a smooth convex function $h(\mathbf{w})$ and a non-smooth convex function $g(\mathbf{D}\mathbf{w})$, which can be optimized efficiently when considered separately. This

suggests to use a proximal splitting method [14, 24, 25] for solving (9). One particular such method is the preconditioned primal-dual method [26] which is based on reformulating the problem (9) as a saddle-point problem

$$\min_{\mathbf{w} \in \mathbb{R}^{dN}} \max_{\mathbf{u} \in \mathbb{R}^{dE}} \mathbf{u}^T \mathbf{D} \mathbf{w} + h(\mathbf{w}) - g^*(\mathbf{u}), \quad (11)$$

with the convex conjugate g^* of g [14].

Solutions $(\hat{\mathbf{w}}, \hat{\mathbf{u}})$ of (11) are characterized by [27, Thm 31.3]

$$\begin{aligned} -\mathbf{D}^T \hat{\mathbf{u}} &\in \partial h(\hat{\mathbf{w}}) \\ \mathbf{D} \hat{\mathbf{w}} &\in \partial g^*(\hat{\mathbf{u}}). \end{aligned} \quad (12)$$

This condition is, in turn, equivalent to

$$\begin{aligned} \hat{\mathbf{w}} - \mathbf{T} \mathbf{D}^T \hat{\mathbf{u}} &\in (\mathbf{I}_{dN} + \mathbf{T} \partial h)(\hat{\mathbf{w}}), \\ \hat{\mathbf{u}} + \mathbf{\Sigma} \mathbf{D} \hat{\mathbf{w}} &\in (\mathbf{I}_{dE} + \mathbf{\Sigma} \partial g^*)(\hat{\mathbf{u}}), \end{aligned} \quad (13)$$

with positive definite matrices $\mathbf{\Sigma} \in \mathbb{R}^{dE \times dE}$, $\mathbf{T} \in \mathbb{R}^{dN \times dN}$. The matrices $\mathbf{\Sigma}$, $\mathbf{T}$ are design parameters whose choice will be detailed below. The condition (13) lends naturally to the following coupled fixed point iterations [26]

$$\mathbf{w}_{k+1} = (\mathbf{I} + \mathbf{T} \partial h)^{-1}(\mathbf{w}_k - \mathbf{T} \mathbf{D}^T \mathbf{u}_k) \quad (14)$$

$$\mathbf{u}_{k+1} = (\mathbf{I} + \mathbf{\Sigma} \partial g^*)^{-1}(\mathbf{u}_k + \mathbf{\Sigma} \mathbf{D}(2\mathbf{w}_{k+1} - \mathbf{w}_k)). \quad (15)$$

The update (15) involves the resolvent operator

$$(\mathbf{I} + \mathbf{\Sigma} \partial g^*)^{-1}(\mathbf{v}) = \operatorname{argmin}_{\mathbf{v}' \in \mathbb{R}^{dE}} g^*(\mathbf{v}') + (1/2) \|\mathbf{v}' - \mathbf{v}\|_{\mathbf{\Sigma}^{-1}}^2, \quad (16)$$

where $\|\mathbf{v}\|_{\mathbf{\Sigma}} := \sqrt{\mathbf{v}^T \mathbf{\Sigma} \mathbf{v}}$. The convex conjugate g^* of g (see (10)) can be decomposed as $g^*(\mathbf{v}) = \sum_{e=1}^E g_2^*(\mathbf{v}^{(e)})$ with the convex conjugate g_2^* of the scaled ℓ_2 -norm $\lambda \|\cdot\|$. Moreover, since $\mathbf{\Sigma}$ is a block diagonal matrix, the e -th block of the resolvent operator $(\mathbf{I}_{dE} + \mathbf{\Sigma} \partial g^*)^{-1}(\mathbf{v})$ can be obtained by the Moreau decomposition as [28, Sec. 6.5]

$$\begin{aligned} &((\mathbf{I}_{dE} + \mathbf{\Sigma} \partial g^*)^{-1}(\mathbf{v}))^{(e)} \\ &\stackrel{(16)}{=} \operatorname{argmin}_{\mathbf{v}' \in \mathbb{R}^d} g_2^*(\mathbf{v}') + (1/(2\sigma^{(e)})) \|\mathbf{v}' - \mathbf{v}^{(e)}\|^2 \\ &= \mathbf{v}^{(e)} - \sigma^{(e)} (\mathbf{I}_d + (\lambda/\sigma^{(e)}) \partial \|\cdot\|)^{-1}(\mathbf{v}^{(e)}/\sigma^{(e)}) \\ &= \begin{cases} \lambda \mathbf{v}^{(e)} / \|\mathbf{v}^{(e)}\| & \text{if } \|\mathbf{v}^{(e)}\| > \lambda \\ \mathbf{v}^{(e)} & \text{otherwise,} \end{cases} \end{aligned}$$

where $(a)_+ = \max\{a, 0\}$ for $a \in \mathbb{R}$.

The update (14) involves the resolvent operator $(\mathbf{I} + \mathbf{T} \partial h)^{-1}$ of h (see (3) and (10)), which does not have a closed-form solution. Choosing $\mathbf{T} = \operatorname{diag}\{\tau^{(i)} \mathbf{I}_d\}_{i=1}^N$, we can solve (14) approximately by a simple iterative method [29, Sec. 8.2]. Setting $\bar{\mathbf{w}} := \mathbf{w}_k - \mathbf{T} \mathbf{D}^T \mathbf{u}_k$, the update (14) becomes

$$\mathbf{w}_{k+1}^{(i)} := \operatorname{argmin}_{\tilde{\mathbf{w}} \in \mathbb{R}^d} 2\ell(\tilde{\mathbf{w}}^T \tilde{\mathbf{x}}^{(i)}) + (M/\tau^{(i)}) \|\tilde{\mathbf{w}} - \bar{\mathbf{w}}^{(i)}\|^2. \quad (17)$$

If the matrices $\mathbf{\Sigma}$ and $\mathbf{T}$ satisfy

$$\|\mathbf{\Sigma}^{1/2} \mathbf{D} \mathbf{T}^{1/2}\|^2 < 1, \quad (18)$$

the sequences obtained from iterating (14) and (15) converge to a saddle point of the problem (11) [26, Thm. 1]. The condition (18) is satisfied for the choice $\mathbf{\Sigma} = \{(1/(2A_e)) \mathbf{I}_d\}_{e \in \mathcal{E}}$ and $\{(\tau/d^{(i)}) \mathbf{I}_d\}_{i \in \mathcal{V}}$, with node degree $d^{(i)} = \sum_{j \neq i} A_{ij}$ and some $\tau < 1$ [26, Lem. 2].

Solving (17) is equivalent to the zero-gradient condition

$$-\tilde{\mathbf{x}}^{(i)} \sigma(-(\mathbf{w}^{(i)})^T \tilde{\mathbf{x}}^{(i)}) + (M/\tau^{(i)})(\mathbf{w}^{(i)} - \bar{\mathbf{w}}^{(i)}) = 0. \quad (19)$$

The solutions of (19) are fixed-points of the map

$$\Phi^{(i)}(\mathbf{u}) = \bar{\mathbf{w}}^{(i)} + (\tau^{(i)}/M) \tilde{\mathbf{x}}^{(i)} \sigma(-\mathbf{u}^T \tilde{\mathbf{x}}^{(i)}). \quad (20)$$

Lemma 1. *The mapping $\Phi^{(i)}$ (20) is Lipschitz with constant $\beta_i = \tau^{(i)} \|\mathbf{x}^{(i)}\|^2 / M$.*

Proof. For any $a, b \in \mathbb{R}$,

$$|1/(1 + \exp(a)) - 1/(1 + \exp(b))| \leq |a - b|$$

which implies

$$|\sigma(-\mathbf{u}^T \tilde{\mathbf{x}}^{(i)}) - \sigma(-\mathbf{v}^T \tilde{\mathbf{x}}^{(i)})| \leq \|\mathbf{x}^{(i)}\| \|\mathbf{u} - \mathbf{v}\|,$$

and, in turn,

$$\begin{aligned} \|\Phi^{(i)}(\mathbf{u}) - \Phi^{(i)}(\mathbf{v})\| &\leq (\tau^{(i)} \|\tilde{\mathbf{x}}^{(i)}\| / M) \|\mathbf{x}^{(i)}\| \|\mathbf{u} - \mathbf{v}\| \\ &= \beta_i \|\mathbf{u} - \mathbf{v}\|. \end{aligned} \quad \square$$

We approximate the exact update (17) with

$$\hat{\mathbf{w}}_{k+1}^{(i)} = \underbrace{\Phi^{(i)} \circ \dots \circ \Phi^{(i)}}_{\lceil 2 \log(k) / \log(1/\beta_i) \rceil}(\bar{\mathbf{w}}^{(i)}). \quad (21)$$

According to [30, Thm. 1.48], for $\tau^{(i)} < M/\|\mathbf{x}^{(i)}\|^2$ the error incurred by replacing (17) with (21) satisfies

$$e_k = \|\hat{\mathbf{w}}_{k+1}^{(i)} - \mathbf{w}_{k+1}^{(i)}\| \leq 1/k^2. \quad (22)$$

Given the error bound (22), as can be verified using [15, Thm. 3.2], the sequences obtained by (14) and (15) when replacing the exact update (17) with (21) converge to a saddle-point of (11) and, in turn, a solution of lnLasso (9).

5. NUMERICAL EXPERIMENTS

We assess the performance of lnLasso Alg. 1 on datasets whose empirical graph is a grid graph consisting $N = 20 \times 20 = 400$ nodes. The graph is partitioned into 4 clusters $\mathcal{C}_r$, which are grid graphs of size 10×10 . The edge weights $A_{ij} = 100$ if nodes i and j belong to the same cluster and $A_{ij} = 1$ otherwise.

Algorithm 1 lnLasso via primal-dual method

Input: $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, $\{\mathbf{x}^{(i)}\}_{i \in \mathcal{V}}$, $\mathcal{M}$, $\{y^{(i)}\}_{i \in \mathcal{M}}$, λ , $\mathbf{D}$, $\Sigma = \text{diag}\{\sigma^{(e)} = 1/(2A_e)\mathbf{I}_d\}_{e=1}^E$, $\mathbf{T} = \text{diag}\{\tau^{(i)} = 0.9/d^{(i)}\mathbf{I}_d\}_{i \in \mathcal{V}}$, $\beta_i = \tau^{(i)}/|\mathcal{M}|$

Initialize: $k := 0$, $\hat{\mathbf{w}}_0 := 0$, $\hat{\mathbf{u}}_0 := 0$

1: **repeat**

2: $\hat{\mathbf{w}}_{k+1} := \hat{\mathbf{w}}_k - \mathbf{T}\mathbf{D}^T\hat{\mathbf{u}}_k$

3: **for** each labeled node $i \in \mathcal{M}$ **do**

4: $\hat{\mathbf{w}}_{k+1}^{(i)} := \underbrace{\Phi^{(i)} \circ \dots \circ \Phi^{(i)}}_{\lceil 2 \log(k) / \log(1/\beta_i) \rceil} (\hat{\mathbf{w}}_{k+1}^{(i)})$

5: **end for**

6: $\bar{\mathbf{u}} := \mathbf{u}_k + \Sigma \mathbf{D}(2\hat{\mathbf{w}}_{k+1} - \hat{\mathbf{w}}_k)$

7: $\hat{\mathbf{u}}_{k+1}^{(e)} = \bar{\mathbf{u}}^{(e)} - \left(1 - \frac{\lambda}{\|\bar{\mathbf{u}}^{(e)}\|}\right)_+ \bar{\mathbf{u}}^{(e)}$ for $e \in \mathcal{E}$

8: $k := k + 1$

9: **until** stopping criterion is satisfied

Output: $(\hat{\mathbf{w}}_k, \hat{\mathbf{u}}_k)$.

All nodes $i \in \mathcal{C}_r$ in a particular cluster $\mathcal{C}_r$ share the same weight vector $\mathbf{w}^{(r)} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ which is generated from a standard normal distribution, with $\sigma = 10$. The feature vectors $\mathbf{x}^{(i)} \in \mathbb{R}^3$ are generated i.i.d. using a uniform distribution over $[0, 1]^3$. The true node labels $y^{(i)}$ are drawn from the distribution (1). The training set is obtained by independently selecting each node with probability (labeling rate) p .

We apply Alg. 1 to obtain a classifier $\hat{\mathbf{w}}$ which allows to classify data points as $\hat{y}^{(i)} = \text{sign}((\hat{\mathbf{w}}^{(i)})^T \mathbf{x}^{(i)})$. In order to assess the performance of Alg. 1 we compute the accuracy within the unlabeled nodes, i.e., the ratio of the number of correct labels achieved by Alg. 1 for unlabeled nodes to the number of unlabeled nodes,

$$ACC := (1/(N - M)) |\{i : y^{(i)} = \hat{y}^{(i)}, i \notin \mathcal{M}\}|. \quad (23)$$

We compute the accuracy of Alg. 1 for different choices of $p \in \{0.1, \dots, 0.9\}$ and $\lambda \in \{10^{-5}, \dots, 10^{-1}\}$. For a pair of $\{\lambda, p\}$, we repeat the experiment 100 times and compute the average accuracy.

The average accuracies obtained for varying labeling rates p and lnLasso parameter λ is plotted in Fig. 1. As indicated in Fig. 1, the accuracy slightly increases for labeling rate p between 0.1 and 0.4 and does not improve for higher labeling rate p . This implies that the lnLasso yields reasonable accuracy even with low labeling rate.

The accuracies in Fig. 1 are restricted since the classifier assigns the labels to nodes based on (2) while the true labels are drawn from the probabilistic model (1). Indeed, we compute the average probability that the nodes are assigned to their true label (label with higher probability) using (1) and plot the result, called reference accuracy, in Fig. 1 as a horizontal dashed line. Fig. 1 shows that accuracies tends toward the reference accuracy.

Fig. 2 plots the accuracy as a function of the number of iterations used in Alg. 1 for varying regularizer λ but fixed labeling rate $p = 0.4$. Fig. 2 shows that the accuracies are close to the reference accuracy and do not improve after few iterations. This implies that lnLasso can yield reasonable accuracy after few numbers of iterations.

Fig. 1: The classification accuracy for grid-structured data.

Fig. 2: The classification accuracy of Alg. 1 for grid structured data and labeling rate $p = 0.4$.

Acknowledgement

We thank Sahel Sarcheshmehpour for helpful comments.

6. REFERENCES

- [1] D. Hallac, J. Leskovec, and S. Boyd, “Network lasso: Clustering and optimization in large graphs,” in *Proc. SIGKDD*, 2015, pp. 387–396.
- [2] A. Jung, N.T. Quang, and A. Mara, “When is network lasso accurate?,” *Frontiers in Appl. Math. and Stat.*, vol. 3, pp. 28, 2018.
- [3] M. Yamada, T. Koh, T. Iwata, J.S. Taylor, and S. Kaski, “Localized Lasso for High-Dimensional Regression,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, USA, 20–22 Apr 2017, vol. 54, pp. 325–333, PMLR.

- [4] A. Jung and M. Hulsebos, “The network nullspace property for compressed sensing of big data over networks,” *Front. Appl. Math. Stat.*, Apr. 2018.
- [5] S. Chen, A. Sandryhaila, J. M. F. Moura, and J. Kovačević, “Signal recovery on graphs: Variation minimization,” *IEEE Trans. Signal Processing*, vol. 63, no. 17, pp. 4609–4624, Sept. 2015.
- [6] A. Sandryhaila and J. M. F. Moura, “Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure,” *IEEE Signal Processing Magazine*, vol. 31, no. 5, pp. 80–90, Sept 2014.
- [7] A. Jung, “On the complexity of sparse label propagation,” *Front. Appl. Math. Stat.*, vol. 4, pp. 22, Jul. 2018.
- [8] A. Jung and N. Tran, “Localized linear regression in networked data,” *IEEE Signal Processing Letters*, vol. 26, no. 7, pp. 1090–1094, July 2019.
- [9] S. Bhagat, C. Graham, and S. Muthukrishnan, *Social network data analytics*, Springer, 2011.
- [10] L. Lovász, *Large Networks and Graph Limits*, American Mathematical Society, 2012.
- [11] S. Cui, A. Hero, Z.-Q. Luo, and J.M.F. Moura, Eds., *Big Data over Networks*, Cambridge Univ. Press, 2016.
- [12] Sofus A. Macskassy and Foster Provost, “Classification in networked data: A toolkit and a univariate case study,” *J. Mach. Learn. Res.*, vol. 8, pp. 935–983, May 2007.
- [13] V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer, 1999.
- [14] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging,” *J. Math. Imaging Vision*, vol. 40, no. 1, pp. 120–145, 2011.
- [15] L. Condat, “A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms,” *Journal of Optimization Theory and Applications*, vol. 158, no. 2, pp. 460–479, Aug 2013.
- [16] O. Chapelle, B. Schölkopf, and A. Zien, Eds., *Semi-Supervised Learning*, The MIT Press, Cambridge, Massachusetts, 2006.
- [17] Pan Zhang, Cristopher Moore, and Lenka Zdeborová, “Phase transitions in semisupervised clustering of sparse networks,” *Physical review. E, Statistical, nonlinear, and soft matter physics*, vol. 90 5-1, pp. 052802, 2014.
- [18] P. Ruusuvuori, T. Manninen, and H. Huttunen, “Image segmentation using sparse logistic regression with spatial prior,” in *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*, Aug. 2012, pp. 2253–2257.
- [19] R. Kchichian, S. Valette, and M. Desvignes, “Automatic multiorgan segmentation via multiscale registration and graph cut,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 12, pp. 2739–2749, Dec 2018.
- [20] Y. Boykov and V. Kolmogorov, “An experimental comparison of min-cut/max- flow algorithms for energy minimization in vision,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124–1137, Sept. 2004.
- [21] M. E. J. Newman, *Networks: An Introduction*, Oxford Univ. Press, 2010.
- [22] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics. Springer, New York, NY, USA, 2001.
- [23] A. Jung, “A fixed-point of view on gradient methods for big data,” *Frontiers in Applied Mathematics and Statistics*, vol. 3, pp. 18, 2017.
- [24] P. L. Combettes and J.-C. Pesquet, “Proximal Splitting Methods in Signal Processing,” *ArXiv e-prints*, Dec. 2009.
- [25] D. O’Connor and L. Vandenberghe, “Primal-dual decomposition by operator splitting and applications to image deblurring,” *SIAM Journal on Imaging Sciences*, vol. 7, no. 3, pp. 1724–1754, 2014.
- [26] T. Pock and A. Chambolle, “Diagonal preconditioning for first order primal-dual algorithms in convex optimization,” in *IEEE International Conference on Computer Vision (ICCV)*, Nov 2011, pp. 1762–1769.
- [27] R. T. Rockafellar, *Convex Analysis*, Princeton Univ. Press, Princeton, NJ, 1970.
- [28] N. Parikh and S. Boyd, “Proximal algorithms,” *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 123–231, 2013.
- [29] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers,” *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [30] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, New York, 2011.