Networked Federated Multi-Task Learning

Y. SarcheshmehPour

Department of Computer Science
Aalto University

Y. Tian

Department of Computer Science
Aalto University

L. Zhang

Department of Computer Science
Aalto University

A. Jung

Department of Computer Science
Aalto University

Abstract

Many important application domains generate distributed collections of heterogeneous local datasets. These local datasets are often related via an intrinsic network structure that arises from domain-specific notions of similarity between local datasets. Different notions of similarity are induced by spatio-temporal proximity, statistical dependencies or functional relations. We use this network structure to adaptively pool similar local datasets into nearly homogenous training sets for learning tailored models. Our main conceptual contribution is to formulate networked federated learning using the concept of generalized total variation (GTV) minimization as a regularizer. This formulation is highly flexible and can be combined with almost any parametric model including Lasso or deep neural networks. We unify and considerably extend some well-known approaches to federated multi-task learning. Our main algorithmic contribution is a novel federated learning algorithm which is well suited for distributed computing environments such as edge computing over wireless networks. This algorithm is robust against model misspecification and numerical errors arising from limited computational resources including processing time or wireless channel bandwidth. As our main technical contribution, we offer precise conditions on the local models as well on their network structure such that our algorithm learns nearly optimal local models. Our analysis reveals an interesting interplay between the (information-) geometry of local models and the (cluster-) geometry of their network.

1 Introduction

Many important application domains generate distributed collections of local datasets that are related via an intrinsic network structure [1]. Such a network structure can arise from spatio-temporal proximity, statistical dependencies or functional relations [2, 3, 4]. The network structure of data might also arise from a distributed computing infrastructure that generates the data [5]. In what follows we assume that the network structure of the data is known. The challenge of learning the network structure from raw data is beyond the scope of this paper (see Section 6 for future research directions).

Two timely application domains facing distributed networked data are the high-precision management of pandemics and the Internet of Things (IoT) [6, 7]. Here, local datasets are generated either by smartphones and wearables or by IoT devices [8, 9]. These local datasets are related via physical contact networks, social networks [10], co-morbidity networks [11], or communication networks [12].

Federated learning (FL) is a recent main thread of machine learning which studies collaborative learning from distributed data [13, 14, 15]. FL methods provide privacy protection as they do not require exchange of raw data that might be sensitive [16, 17]. Moreover, FL offers robustness against malicious data perturbation due to its intrinsic averaging or aggregation over large collections of (mostly benign) datasets [18].

1.1 Related Work

Similar to [19], we frame FL as a multi-task learning problem that is solved using regularized empirical risk minimization (RERM). For each local dataset we obtain a separate learning task that amounts to finding an optimal choice for the weights of a local model. These individual learning tasks are coupled by a known network structure which we represent as a weighted undirected “empirical graph” (see Section 2). In contrast to [20], we do not use a probabilistic model for the network structure of the individual tasks. To capture the intrinsic cluster structure of networked data, we construct the regularization term using the concept of generalized total variation (GTV) which unifies widely-used measures for the variation of node attributes (see Section 2.2).

Our approach unifies and considerably extends some recent approaches to distributed FL including the network Lasso (nLasso) [21, 22, 23, 24] and FL methods that use the graph Laplacian quadratic form as regularizer [19, 25]. While the methods in [22, 23, 24, 26, 27] are limited to generalized linear models, the proposed FL algorithm can be combined with non-linear parametrized models including graphical Lasso or deep neural networks [28, 29]. In contrast to [25], we do not require access to the entire networked data but only a small subset of local datasets.

The main algorithmic device behind our method is a primal-dual approach to solving large-scale convex optimization problems [30, 31, 32]. We obtain a novel networked FL algorithm as the direct application of this primal-dual approach to GTV minimization (see Section 3). This generic primal-dual approach is appealing for FL applications due to its scalability and robustness against modelling errors or faults in the computational infrastructure [33, 34, 35, 36]. The iterations of primal-dual methods typically involve smaller (atomic) optimization problems. Given finite computational resources these atomic optimization problems can only be solved to some non-zero optimization error. The effect of inexact updates in primal-dual methods has been analyzed recently [37].

A large body of existing work has studied computational and statistical aspects of RERM in multi-task learning (see, e.g., [38, 39, 40, 41]). Investigations of statistical properties of RERM have mainly focused on (group-) sparsity as the regularization principle [42, 43, 44]. However, we are not aware of studies about the interplay between estimation error and the network cluster structure in the context of FL. We close this gap by providing a precise characterization of network structures and available data such that RERM is successful (see Section 4).

Our approach relies on the clustering assumption that local datasets that form a network cluster have similar statistical properties and, in turn, similar optimizers for model weights. What sets our approach apart from existing methods for clustered FL [45, 46] is that we exploit a known network structure to pool local datasets. This pooling is guided by the network structure and aimed at obtaining sufficiently large training sets for learning the local model weights.

1.2 Contribution

We propose and study a novel FL algorithm that learns separate (parametrized) models for each local dataset within a collection of networked data. The networked data is represented by a given undirected empirical graph whose edges connect (statistically) similar datasets. Our method exploits the cluster structure of the empirical graph to adaptively pool datapoints for learning model weights.

This paper provides three main contributions.

- As our main conceptual contribution, we show that GTV minimization is a useful framework for federated multi-task learning in networked data. The GTV minimization framework unifies and significantly extends the nLasso and the MOCHA method.
- Our main methodological contribution is a novel family of distributed FL algorithms (see Algorithm 1). This family is obtained by solving GTV minimization with a primal-dual method for structured optimization problems [34]. Each member of this family corresponds

Figure 1: We represent networked data and corresponding models using an undirected empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node $i \in \mathcal{V}$ of the graph carries a local dataset $\mathcal{X}^{(i)}$ and model weights $\mathbf{w}^{(i)}$ which are scored using a local loss function $\mathcal{L}^{(i)}(\mathcal{X}^{(i)}; \mathbf{w}^{(i)})$. Two nodes are connected by a weighted edge $\{i, j\}$ if they carry similar datasets. The amount of similarity is encoded in an edge weight $A_{i,j} > 0$ (indicated by varying thickness). We rely on a clustering assumption in that the optimal weight vectors for nodes in the same cluster $\mathcal{C} \subseteq \mathcal{V}$ are in the proximity of a cluster-wise optimal weight vector $\mathbf{w}^{(\mathcal{C})}$. Here, we indicate a partition of the empirical graph into three disjoint clusters $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$. Note that our FL method does not require the (typically unknown) partition but rather learns the partition based on the local datasets and network structure of $\mathcal{G}$.

to a particular choices for the local models and the choice of penalty function used to define the GTV. By exploiting the network structure of data, our algorithm can also handle partially observed datasets which is relevant for semi-supervised learning settings.

- Our main analytical contribution is an upper bound on the estimation error incurred by GTV-based RERM. This upper bound reveals sufficient conditions on the local datasets and network structure such that our method achieves the performance of an oracle-based method that perfectly knows the true cluster structure of the data network.

Notation. The identity matrix of size $n \times n$ is denoted $\mathbf{I}_n$, with the subscript omitted if the size n is clear from context. The Euclidean norm of a vector $\mathbf{w} = (w_1, \dots, w_n)^T$ is $\|\mathbf{w}\|_2 := \sqrt{\sum_{r=1}^n w_r^2}$ and the ℓ_1 norm $\|\mathbf{w}\|_1 := \sum_{r=1}^n |w_r|$. Given a positive semi-definite matrix $\mathbf{M}$, we define the norm $\|\mathbf{w}\|_{\mathbf{M}} := \sqrt{\mathbf{w}^T \mathbf{M} \mathbf{w}}$. It will be convenient to use the notation $(1/2\tau)$ instead of $(1/(2\tau))$. We will need the clipping function $\mathcal{T}^{(\gamma)}(\mathbf{w}) := \gamma \mathbf{w} / \|\mathbf{w}\|_2$ for $\|\mathbf{w}\|_2 \geq \gamma$ and $\mathcal{T}^{(\gamma)}(\mathbf{w}) := \mathbf{w}$ otherwise.

2 Problem Formulation

Networked data can be represented elegantly using an undirected “empirical” graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as illustrated in Figure 1. Every node $i \in \mathcal{V}$ of the empirical graph carries a separate local dataset $\mathcal{X}^{(i)}$. The local dataset $\mathcal{X}^{(i)}$ might be a single sensor measurement, an entire time series or even a whole collection of videos. Our approach is agnostic towards the details of the data representation such as if datapoints are labelled or unlabelled. A key aspect of our method is that it accesses $\mathcal{X}^{(i)}$ only indirectly via a local loss function (see Section 2.1). Another key aspect of our approach is that we use local datasets only at nodes in a (small) training set $\mathcal{M} = \{i_1, \dots, i_M\} \subseteq \mathcal{V}$. To compensate for not using local datasets outside $\mathcal{M}$ we exploit the cluster structure of $\mathcal{G}$. This is relevant for applications where accessing local datasets is computationally costly.

An undirected edge $\{i, j\} \in \mathcal{E}$ indicates that the corresponding local datasets $\mathcal{X}^{(i)}$ and $\mathcal{X}^{(j)}$ have similar statistical properties. The strength of the similarity is quantified by the edge weight $A_{ij} > 0$. The neighbourhood of a node $i \in \mathcal{V}$ is $\mathcal{N}_i := \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}\}$. It will be convenient to define the head and tail of an undirected edge $\{i, j\}$ as $e_+ := \min\{i, j\}$ and $e_- := \max\{i, j\}$, respectively.

In what follows, we will make use of two vector spaces that are naturally associated with an empirical graph $\mathcal{G}$. The space $\mathbb{R}^{\mathcal{V} \times n}$ of networked weight vectors consists of maps $\mathbf{w} : \mathcal{V} \rightarrow \mathbb{R}^n : i \mapsto \mathbf{w}^{(i)}$ that assign each node $i \in \mathcal{V}$ a vector $\mathbf{w}^{(i)} \in \mathbb{R}^n$. Another (in some sense dual) space $\mathbb{R}^{\mathcal{E} \times n}$ is obtained by all maps $\mathbf{u} : \mathcal{E} \rightarrow \mathbb{R}^n : e \mapsto \mathbf{u}^{(e)}$ that assign each edge $e \in \mathcal{E}$ a vector $\mathbf{u}^{(e)} \in \mathbb{R}^n$. These two spaces are linked via the block-incidence matrix $\mathbf{D} \in \mathbb{R}^{n|\mathcal{E}| \times n|\mathcal{V}|}$,

$$D_{e,i} = \mathbf{I} \text{ for } i = e_+, D_{e,i} = -\mathbf{I} \text{ for } i = e_- \text{ and } D_{e,i} = \mathbf{0} \text{ otherwise.} \quad (1)$$

Applying $\mathbf{D}$ to a given networked weight vector $\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}$ results in the networked dual weights $\mathbf{u} \in \mathbb{R}^{\mathcal{E} \times n}$ with $\mathbf{u}^{(e)} = \mathbf{w}^{(e_+)} - \mathbf{w}^{(e_-)}$.

2.1 Networked Models

Our goal is to learn, for each node $i \in \mathcal{V}$, the weights (parameters) $\mathbf{w}^{(i)}$ of a local model that is tailored to the local dataset $\mathcal{X}^{(i)}$. The usefulness of a particular choice for the weights $\mathbf{w}^{(i)}$ is measured by a loss function $\mathcal{L}(\mathcal{X}^{(i)}; \mathbf{w}^{(i)})$. In principle, our approach allows for the use of an arbitrary loss function. However, unless otherwise noted, we tacitly assume each loss function $\mathcal{L}(\mathcal{X}^{(i)}; \mathbf{w}^{(i)})$ to be a convex function of the weight vector $\mathbf{w}^{(i)}$ [47]. Specific choices for the local loss function, corresponding to some well-known parametric models, are discussed in Appendix B.

We define networked (model) weights as a map $\mathbf{w} : \mathcal{V} \rightarrow \mathbb{R}^n : i \mapsto \mathbf{w}^{(i)}$ that assigns each node $i \in \mathcal{V}$ the a local weight vectors $\mathbf{w}^{(i)}$. The space of all such networked weights is denoted $\mathbb{R}^{\mathcal{V} \times n}$. To evaluate the quality of given networked weights $\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}$, we compute the training error

$$f(\mathbf{w}) := \sum_{i \in \mathcal{M}} \mathcal{L}^{(i)}(\mathbf{w}^{(i)}). \quad (2)$$

Here we used the shorthand $\mathcal{L}^{(i)}(\mathbf{w}^{(i)}) := \mathcal{L}(\mathcal{X}^{(i)}; \mathbf{w}^{(i)})$. Unless noted otherwise, we assume that $\mathcal{L}^{(i)}(\cdot)$, for $i \in \mathcal{V}$ are convex functions. Additional restrictions are placed on $\mathcal{L}^{(i)}(\cdot)$ in Section 4 to facilitate the analysis of statistical properties of our proposed algorithm. This algorithm aims at learning networked weights $\mathbf{w}$ such that the training error (2) is small.

The criterion (2) alone is insufficient to guide the learning of the weights $\mathbf{w}^{(i)}$ since it completely ignores the weights $\mathbf{w}^{(i)}$ at nodes $i \in \mathcal{V} \setminus \mathcal{M}$ outside the training set. We therefore need to impose some additional structure on the collection of weight vectors $\mathbf{w}^{(i)}$, for $i \in \mathcal{V}$. To this end, we exploit the network structure of the empirical graph $\mathcal{G}$. We assume that the empirical graph can be decomposed into few tight-knit clusters. We then require the weight vectors $\mathbf{w}^{(i)}$ to be approximately constant for all nodes $i \in \mathcal{V}$ belonging to the same cluster.

2.2 Generalized Total Variation Minimization

To learn the weights $\mathbf{w}^{(i)}$ of the local models, one for each local dataset $\mathcal{X}^{(i)}$, we assume that local datasets forming a tight-knit subset (or cluster) $\mathcal{C} \subseteq \mathcal{V}$ have similar statistical properties. It is therefore sensible to enforce the weight vectors $\mathbf{w}^{(i)}$ to be approximately constant for all nodes $i \in \mathcal{V}$ in the same cluster $\mathcal{C} \subseteq \mathcal{V}$. Thus, we learn a cluster-wise weight vector $\mathbf{w}^{(\mathcal{C})}$ by pooling local datasets $\mathcal{X}^{(i)}$ with $i \in \mathcal{C}$ into an intermediate training set.

Let us define the variation of networked weights $\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}$ using the map $\mathbf{u} : e \in \mathcal{E} \mapsto \mathbf{u}^{(e)} \in \mathbb{R}^n$ which assigns each edge $e \in \mathcal{E}$ the difference $\mathbf{u}^{(e)} := \mathbf{w}^{(e_+)} - \mathbf{w}^{(e_-)}$. Using the block-incidence matrix we can write more compactly $\mathbf{u} = \mathbf{D}\mathbf{w}$.

To enforce similar predictors $\mathbf{w}^{(i)} \approx \mathbf{w}^{(j)}$ for nodes $i, j \in \mathcal{V}$ in the same cluster, we require them to have a small generalized TV (GTV)

$$\|\mathbf{w}\|_{\text{GTV}} := \sum_{\{i,j\} \in \mathcal{E}} A_{ij} \phi(\mathbf{w}^{(j)} - \mathbf{w}^{(i)}). \quad (3)$$

We define the GTV for an arbitrary edge set $\mathcal{S} \subseteq \mathcal{E}$ as $\|\mathbf{w}\|_{\mathcal{S}} := \sum_{\{i,j\} \in \mathcal{S}} A_{ij} \phi(\mathbf{w}^{(j)} - \mathbf{w}^{(i)})$. The GTV (3) uses a non-negative penalty function $\phi(\mathbf{v}) \in \mathbb{R}_+$ which is typically increasing with increasing norm $\|\mathbf{v}\|$ of the argument $\mathbf{v} \in \mathbb{R}^n$. Our approach allows for different choices for the penalty function $\phi(\cdot)$. Two specific choices are $\phi(\mathbf{v}) := \|\mathbf{v}\|_2$, which is used by nLasso [21], and $\phi(\mathbf{v}) := (1/2)\|\mathbf{v}\|_2^2$ which is used by ‘‘MOCHA’’ [19]. Another recent FL method for networked data uses the choice $\phi(\mathbf{v}) := \|\mathbf{v}\|_1$ [48].

Enforcing a small GTV (3), requires weights $\mathbf{w}^{(i)}$ to change only over few edges $e \in \mathcal{E}$ with small weight A_e . To balance between training error (2) and GTV (3), we solve

$$\hat{\mathbf{w}} \in \arg \min_{\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}} f(\mathbf{w}) + \lambda \|\mathbf{w}\|_{\text{GTV}}. \quad (4)$$

Note that GTV minimization (4) is an instance of RERM with the GTV as the regularizer. The empirical risk incurred by networked weights $\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}$ is measured by the training error $f(\mathbf{w})$ (2).

We will solve (4) using a primal-dual method lending to our FL algorithm in Section 3. The statistical properties of the solutions to (4) are the subject of Section 4.

The regularization parameter $\lambda > 0$ in (4) allows to trade a small GTV $\|\hat{\mathbf{w}}\|_{\text{GTV}}$ against a small training error $f(\hat{\mathbf{w}})$ (2). The choice of λ can be guided by cross validation [49] or by our analysis of the solutions of (4) in Section 4. Loosely speaking, increasing the value of λ results in the solutions of (4) becoming increasingly clustered, i.e., the weight vectors $\hat{\mathbf{w}}^{(i)}$ become constant over increasingly larger subsets of nodes. Choosing λ larger than some critical value, that depends on the local datasets and network structure of $\mathcal{G}$, results in constant weight vectors $\hat{\mathbf{w}}^{(i)}$ for all nodes $i \in \mathcal{V}$.

GTV minimization (4) unifies and extends some well-known methods for distributed optimization and learning. The nLasso [21] is obtained from (4) for the choice $\phi(\mathbf{v}) := \|\mathbf{v}\|_2$. The FL algorithm MOCHA [19] is obtained from (4) for the choice $\phi(\mathbf{v}) := \|\mathbf{v}\|_2^2$. Another special case of (4), obtained for the choice $\phi(\mathbf{v}) := \|\mathbf{v}\|_1$, has been studied recently [48].

Different choices for the penalty function $\phi(\cdot)$ offer different trade-offs between computational complexity and statistical properties of the resulting FL algorithm. Non-smooth TV minimization ($\phi(\mathbf{u}) = \|\mathbf{u}\|_2$) is computationally more challenging than methods using graph Laplacian quadratic form ($\phi(\mathbf{u}) = \|\mathbf{u}\|_2^2$). Statistically, non-smooth TV minimization can be more accurate on network structures that are challenging for approaches relying on graph Laplacian quadratic form [50, 51].

3 Networked Federated Multi-Task Learning Algorithm

We now discuss our FL algorithm that is obtained by applying an established primal-dual approach to solve GTV minimization (4). Let us first rewrite (4) more compactly (see (1) and (3)),

$$\hat{\mathbf{w}} \in \arg \min_{\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}} f(\mathbf{w}) + g(\mathbf{D}\mathbf{w}) \text{ with } g(\mathbf{u}) := \lambda \sum_{e \in \mathcal{E}} A_e \phi(\mathbf{u}^{(e)}). \quad (5)$$

The objective function in of the GTV minimization problem (5) consists of two components. The first component is the training error $f(\mathbf{w})$ (2), which depends on the networked weights $\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}$. The second component is the scaled GTV $g(\mathbf{u})$ (3) which depends on the dual weights $\mathbf{u} = \mathbf{D}\mathbf{w}$, with $\mathbf{u}^{(e)} = \mathbf{w}^{(e+)} - \mathbf{w}^{(e-)}$, for $e \in \mathcal{E}$.

We solve (5) jointly with the dual problem

$$\max_{\mathbf{u} \in \mathbb{R}^{\mathcal{E} \times n}} -g^*(\mathbf{u}) - f^*(-\mathbf{D}^T \mathbf{u}). \quad (6)$$

The objective function in (6) is composed of the convex conjugates [47]

$$g^*(\mathbf{u}) := \sup_{\mathbf{z} \in \mathbb{R}^{\mathcal{E} \times n}} \sum_{e \in \mathcal{E}} (\mathbf{u}^{(e)})^T \mathbf{z}^{(e)} - g(\mathbf{z}) \stackrel{(5)}{=} \sup_{\mathbf{z} \in \mathbb{R}^{\mathcal{E} \times n}} \sum_{e \in \mathcal{E}} (\mathbf{u}^{(e)})^T \mathbf{z}^{(e)} - \lambda A_e \phi(\mathbf{z}^{(e)}), \quad (7)$$

and

$$f^*(\mathbf{w}) := \sup_{\mathbf{z} \in \mathbb{R}^{\mathcal{V} \times n}} \sum_{i \in \mathcal{V}} (\mathbf{w}^{(i)})^T \mathbf{z}^{(i)} - f(\mathbf{z}) \stackrel{(2)}{=} \sup_{\mathbf{z} \in \mathbb{R}^{\mathcal{V} \times n}} \sum_{i \in \mathcal{V}} (\mathbf{w}^{(i)})^T \mathbf{z}^{(i)} - \sum_{i \in \mathcal{M}} \mathcal{L}(\mathbf{w}^{(i)}). \quad (8)$$

The domain of the dual problem (6) is the dual space $\mathbb{R}^{\mathcal{E} \times n}$ of maps $\mathbf{u} : \mathcal{E} \rightarrow \mathbb{R}^n$ that assign a separate dual vector $\mathbf{u}^{(e)}$ to each edge $e \in \mathcal{E}$.

The duality between (5) and (6) is made precise in [52, Ch. 31] (see also [32, Sec. 3.5]). The optimal values of both problems coincide [52, Cor. 31.2.1],

$$\min_{\mathbf{w} \in \mathbb{R}^{\mathcal{V} \times n}} f(\mathbf{w}) + g(\mathbf{D}\mathbf{w}) = \max_{\mathbf{u} \in \mathbb{R}^{\mathcal{E} \times n}} -g^*(\mathbf{u}) - f^*(-\mathbf{D}^T \mathbf{u}). \quad (9)$$

The strong duality (9) allows to bound the sub-optimality of given networked weights. Indeed, for any given dual variable $\hat{\mathbf{u}}$, the objective function value $-g^*(\hat{\mathbf{u}}) - f^*(-\mathbf{D}^T \hat{\mathbf{u}})$ is a lower bound for the optimal value of (5).

We obtain Algorithm 1 by applying a proximal point algorithm to jointly solve (5) and (6). The details of this application are provided in Appendix C. The key components of Algorithm 1 are node-wise and edge-wise updates in step 6 and 10. The results of these local computations are then propagated to the adjacent edges and nodes. Let us next discuss the node-wise and edge-wise updates in detail.

The node-wise updates in step 6 of Algorithm 1 use the primal update operator

$$\mathcal{PU}^{(i)}\{\mathbf{v}\} := \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \mathcal{L}^{(i)}(\mathbf{z}) + (1/2\tau_i)\|\mathbf{v} - \mathbf{z}\|^2. \quad (10)$$

The primal update operator (10) solves a regularized variant of the local loss $\mathcal{L}^{(i)}(\mathbf{w}^{(i)})$. The regularizer is the squared Euclidean distance to the argument $\mathbf{v}$ of the operator and will be used to enforce similar weight vectors at well-connected nodes (belonging to the same cluster). Note that the primal update operator (10) depends on the choice for the loss function. We discuss specific choices for loss functions and corresponding primal update operators in Appendix B.

Algorithm 1 involves dual updates, for each edge $e \in \mathcal{E}$, using the dual update operator

$$\mathcal{DU}^{(e)}\{\mathbf{v}\} := \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \lambda A_e \phi^*(\mathbf{z}/(\lambda A_e)) + (1/2\sigma_e)\|\mathbf{v} - \mathbf{z}\|^2. \quad (11)$$

Here, we used the convex conjugate $\phi^*(\mathbf{v}) := \sup_{\mathbf{z} \in \mathbb{R}^n} \mathbf{v}^T \mathbf{z} - \phi(\mathbf{z})$ of the GTV penalty function $\phi(\mathbf{v})$ (see (3)). The dual update operator (48) depends on the choice for the GTV penalty function $\phi(\mathbf{v})$ (3). For the “nLasso penalty” $\phi(\mathbf{v}) := \|\mathbf{v}\|_2$ [21], the dual update operator becomes a block-clipping $\mathcal{DU}^{(e)}\{\mathbf{v}\} := \mathcal{T}^{(\lambda A_e)}\{\mathbf{v}\}$. The “MOCHA penalty” $\phi(\mathbf{v}) := (1/2)\|\mathbf{v}\|_2^2$ [19], lends to the dual update operator $\mathcal{DU}^{(e)}\{\mathbf{v}\} := \mathbf{v}/(1 + (2\sigma_e/(\lambda A_e)))$. Yet another useful penalty is $\phi(\mathbf{v}) := \|\mathbf{v}\|_1$ with corresponding dual update operator $\mathcal{DU}^{(e)}\{\mathbf{v}\} := (\mathcal{T}^{(\lambda A_e)}(v_1), \dots, \mathcal{T}^{(\lambda A_e)}(v_n))^T$.

Algorithm 1 is to be understood as a template for an entire family of FL algorithms. This family is parametrized by choices for the loss function $\mathcal{L}^{(i)}(\cdot)$, defining the training error (2), or penalty function $\phi(\cdot)$ defining the GTV (3). Different choices for the local loss or penalty function result in different primal update and dual update operators (10) and (48), respectively.

Algorithm 1 Primal-Dual Method for Networked FL

Input: empirical graph $\mathcal{G}$; local loss $\mathcal{L}^{(i)}$ for $i \in \mathcal{M}$, GTV parameters λ, ϕ

Initialize: $k := 0; \hat{\mathbf{w}}_0 := \mathbf{0}; \hat{\mathbf{u}}_0 := \mathbf{0}; \sigma_e = 1/2$ and $\tau_i = 1/|\mathcal{N}_i|$

```

1: while stopping criterion is not satisfied do
2:   for all nodes  $i \in \mathcal{V}$  do
3:      $\hat{\mathbf{w}}_{k+1}^{(i)} := \hat{\mathbf{w}}_k^{(i)} - \tau_i \sum_{e \in \mathcal{E}} D_{e,i} \hat{\mathbf{u}}_k^{(e)}$ 
4:   end for
5:   for nodes in the training set  $i \in \mathcal{M}$  do
6:      $\hat{\mathbf{w}}_{k+1}^{(i)} := \mathcal{PU}^{(i)}\{\hat{\mathbf{w}}_{k+1}^{(i)}\}$ 
7:   end for
8:   for all edges  $e \in \mathcal{E}$  do
9:      $\hat{\mathbf{u}}_{k+1}^{(e)} := \hat{\mathbf{u}}_k^{(e)} + \sigma_e (2(\hat{\mathbf{w}}_{k+1}^{(e+)} - \hat{\mathbf{w}}_{k+1}^{(e-)}) - (\hat{\mathbf{w}}_k^{(e+)} - \hat{\mathbf{w}}_k^{(e-)}))$ 
10:     $\hat{\mathbf{u}}_{k+1}^{(e)} := \mathcal{DU}^{(e)}\{\hat{\mathbf{u}}_{k+1}^{(e)}\}$ 
11:   end for
12:    $k := k + 1$ 
13: end while

```

Algorithm 1 can be implemented as message passing over the edges of the empirical graph $\mathcal{G}$. During each iteration, local computations are carried out at each node and each edge. The primal update step (6) can be executed in parallel for each node $i \in \mathcal{M}$. It involves only the local weight vector $\hat{\mathbf{w}}_{k+1}^{(i)}$ and the local loss function $\mathcal{L}^{(i)}(\mathbf{w}^{(i)})$ (which encapsulates the local dataset $\mathcal{X}^{(i)}$). The dual update step (10) can be carried out in parallel for each edge $e \in \mathcal{E}$.

The results of the parallel primal and dual updates in step (6) and (10) are spread to neighbouring (incident) edges and nodes in the steps (3) and (9). For data whose network structure has a bounded maximum node degree, the computational complexity of Algorithm 1 scales linearly with the number of local datasets.

A computationally appealing property of Algorithm 1 is its robustness against errors occurring during the updates in step 6 and 10. . This is important for applications where the update operators (10) and (48) can be evaluated only approximately or where these updates have to be forwarded over imperfect wireless links. We can ensure convergence to an approximate solution of GTV minimization (5)

even if we implement the primal update operator (10) only numerically (with sufficient accuracy). Let us denote the perturbed update, e.g., obtained by some numerical method for solving (10), by $\tilde{\mathbf{w}}_{k+1}$ and the exact update by $\hat{\mathbf{w}}_{k+1}$. Then, Algorithm 1 converges to a solution of (5) as long as $\sum_{k=1}^{\infty} \|\tilde{\mathbf{w}}_{k+1} - \hat{\mathbf{w}}_{k+1}\| < \infty$ (see [34, Sec. 5]).

Algorithm 1 combines the information contained in the local datasets with their network structure to iteratively improve the weight vectors $\hat{\mathbf{w}}_k^{(i)}$ for each node $i \in \mathcal{V}$. Step 6 adapts the current weight vectors $\hat{\mathbf{w}}_k^{(i)}$ to better fit the labeled local datasets $\mathcal{X}^{(i)}$ for $i \in \mathcal{M}$. These updates are then propagated beyond the training set via steps 3 and 9 of Algorithm 1.

We can use different stopping criteria for Algorithm 1 such as a fixed number of iterations or requiring a minimum decrease of the objective function (5) during each iteration. We can also mix criteria and stop iterating either after a maximum number of iterations has been reached or when the reduction of the objective function falls below a prescribed threshold.

4 Statistical Aspects

This section studies the statistical properties of Algorithm 1 by analyzing the exact solutions to (5). Our main analytical contribution is an upper bound on the deviation between solutions of (5) and weight vectors obtained by pooling local datasets using the true underlying cluster structure of $\mathcal{G}$. This upper bound applies only under certain conditions on the local loss functions and the cluster structure of $\mathcal{G}$. These conditions will be formalized as three explicit assumptions.

The first assumption restricts our analysis to loss functions $\mathcal{L}^{(i)}(\cdot)$ that are convex and differentiable. This is a rather mild restriction as it includes many popular choices for parametrized models and loss functions (see Appendix B).

Assumption 1 (Smooth and Convex Loss). *For each node $i \in \mathcal{V}$, the local loss $\mathcal{L}^{(i)}(\mathbf{v})$ is convex and smooth with a continuous Hessian $\nabla^2 \mathcal{L}^{(i)}(\mathbf{v}) \in \mathbb{R}^{n \times n}$ with entries $\frac{\partial^2 \mathcal{L}^{(i)}(\mathbf{v})}{\partial v_r \partial v_s}$. There is a positive constant $C > 0$ such that the eigenvalues of the Hessian $\nabla^2 \mathcal{L}^{(i)}$ are lower bounded as*

$$\nabla^2 \mathcal{L}^{(i)}(\mathbf{v}) \succeq 2C\mathbf{I}, \text{ for all } \mathbf{v} \in \mathbb{R}^n, i \in \mathcal{V}. \quad (12)$$

The factor 2 in (12) is only for notational convenience. For the special case of linear regression models and squared error loss (see Appendix B), the condition (12) is satisfied with high probability if the features of datapoints are not too strongly correlated.

The idea behind using GTV minimization (5) is that minimizers of loss functions $\mathcal{L}^{(i)}(\cdot)$ in the same cluster $i \in \mathcal{C}$ are all near to a common cluster-wise optimal weight

$$\bar{\mathbf{w}}^{(C)} = \operatorname{argmin}_{\mathbf{v} \in \mathbb{R}^n} \sum_{i \in C} \mathcal{L}^{(i)}(\mathbf{v}). \quad (13)$$

Thus, the collections of local loss functions are simultaneously minimized approximately by the piece-wise constant weight vectors

$$\bar{\mathbf{w}} : i \in \mathcal{V} \mapsto \bar{\mathbf{w}}^{(i)} \stackrel{(13)}{:=} \bar{\mathbf{w}}^{(C)} \text{ for all } i \in C. \quad (14)$$

Note that in practice we cannot use (13) to learn the model weights $\mathbf{w}^{(i)}$ since we do not know the true cluster structure of the empirical graph $\mathcal{G}$. Only if we had an oracle that provides us a true cluster $\mathcal{C}$, we could actually compute the optimal weight (13) for all nodes $i \in \mathcal{C}$. The role of (13) is to provide a (practically infeasible) benchmark for assessing the statistical properties of GTV minimization (5). We provide conditions on the (cluster structure of the) empirical graph $\mathcal{G}$ such that solutions of (5) are close to the oracle-based optimal weights (13).

Our main analytical result below is an upper bound on the deviation between $\bar{\mathbf{w}}$ and the solution $\hat{\mathbf{w}}$ of GTV minimization (5). This upper bound depends on the discrepancy between the node-wise (local) minimizers of $\mathcal{L}^{(i)}(\cdot)$, for some $i \in \mathcal{C}$, and the cluster-wise optimal weight vector $\bar{\mathbf{w}}^{(C)}$, we use the norm $\|\nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(C)})\|_2$. We now formalize our clustering assumption by requiring the norm $\|\nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(C)})\|_2$ to be uniformly bounded for all nodes in the empirical graph.

Assumption 2 (Clustering). *There is a constant U and a partition $\mathcal{P} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{P}|}\}$ of the empirical graph $\mathcal{G}$ into disjoint clusters such that for each cluster $\mathcal{C}_l \in \mathcal{P}$,*

$$\|\nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(\mathcal{C}_l)})\|_2 \leq U \text{ for all } i \in \mathcal{C}_l. \quad (15)$$

Here, we used the optimal weight vector (13) for cluster $\mathcal{C}_l$.

In principle, our analysis applies to an arbitrary choice for the partition $\mathcal{P}$. However, the analysis is most useful if the partition is such that the boundary

$$\partial \mathcal{P} := \{\{i, j\} \in \mathcal{E} : i \in \mathcal{C}_l, j \in \mathcal{C}_{l'}, l \neq l'\} \quad (16)$$

is small in a certain sense. In particular, we focus on partitions such that $\sum_{e \in \partial \mathcal{P}} A_e$ is small. We emphasize that the partition $\mathcal{P}$ used in Assumption 2 is only required for the analysis of Algorithm 1. Algorithm 1 itself does not require the specification of the partition $\mathcal{P}$.

Our third assumption requires that the GTV of approximately piece-wise constant (see Assumption 2) cannot significantly exceed the size of the weights on the training set $\mathcal{M}$. This assumption also restricts our analysis to GTV penalty functions $\phi(\cdot)$ being a norm [53].

Assumption 3 (Sampling Condition). *Consider the GTV minimization (5) with the penalty function $\phi(\cdot)$ (3) being a norm and the local loss functions satisfying (12) with $C > 1$. There are positive constants $L > 1$ and $K > 0$ such that*

$$L\|\mathbf{w}\|_{\partial \mathcal{P}} \leq K\|\mathbf{w}\|_{\mathcal{M}} + \|\mathbf{w}\|_{\mathcal{E} \setminus \partial \mathcal{P}}, \text{ for any } \mathbf{w} \in \mathbb{R}^{n \times \mathcal{V}}. \quad (17)$$

The norm $\|\mathbf{w}\|_{\mathcal{M}} := \sqrt{\sum_{i \in \mathcal{M}} \|\mathbf{w}^{(i)}\|_2^2}$ measures the size of the weights $\mathbf{w}^{(i)}$ on the training set $\mathcal{M}$.

Assumption 3 ensures that piece-wise weight vectors, with sufficiently large GTV, cannot be arbitrarily small on $\mathcal{M}$. The condition (17) can be verified by the existence of sufficiently large network flows between nodes in the training set $\mathcal{M}$ and the boundary edges $\partial \mathcal{P}$ that connect nodes in different clusters [50, 54]. Using the flow condition [50, Lemma 6], it has been shown that (17) is satisfied with high probability for empirical graphs obtained as realizations of a stochastic block model (SBM) using a certain range of SBM parameters [55].

Theorem 1. *Consider local loss functions $\mathcal{L}^{(i)}(\cdot)$ that satisfy Assumption 1 and are approximately minimized by a piece-wise constant $\bar{\mathbf{w}}$ according to Assumption 2. We can access the local loss only for nodes in the training set $\mathcal{M} \subseteq \mathcal{V}$ which is such that Assumption 3 is valid. The estimation error incurred by solutions $\hat{\mathbf{w}}$ of GTV minimization (5) with $\lambda = UL\sqrt{M}/K$ satisfies*

$$\|\hat{\mathbf{w}} - \bar{\mathbf{w}}\|_{\text{GTV}} \leq \frac{8ULK\sqrt{M}}{C(L-1)^2} \quad (18)$$

Theorem 1 uses the specific choice $\lambda = UL\sqrt{M}/K$ for the GTV regularization parameter λ in (4). This provides a means for guiding the choice of λ based on (estimates for) the constants U, L, K and size M of the training set $\mathcal{M}$. If these constants cannot be determined reliably, it might be more convenient to tune λ via cross-validation techniques [49].

The bound (18) reveals an interesting interplay between the geometry of the loss functions $\mathcal{L}^{(i)}$, via constants C and U in (12) and 15, and the cluster geometry of the empirical graph $\mathcal{G}$, via constants L and K in Assumption 17. It can be shown that (17) holds with prescribed constant L if sufficiently large flows can be routed from cluster boundaries $\partial \mathcal{P}$ to the training set. According to (18), ensuring (17) with a larger L allows to tolerate a larger constant U in (15), which means that we can tolerate a larger discrepancy between true minimizers of $\mathcal{L}^{(i)}$ and the piece-wise constant weights (14).

5 Numerical Experiments

This section reports the results of some illustrative numerical experiments to verify the performance of Algorithm 1. We provide the code to reproduce these experiments at <https://github.com/sahelyiyi/FederatedLearning>. The experiments have been carried out using a standard desktop computer and revolve around synthetic datasets whose empirical graph is a realization of a SBM [56]. The nodes $\mathcal{V}$ are partitioned into two equal-sized clusters $\mathcal{C}_1, \mathcal{C}_2$. Two nodes in the same cluster are

connected by a (unit weight) edge with probability p_{in} . Two nodes in different clusters are connected by a (unit weight) edge with probability p_{out} (typically $p_{\text{out}} \ll p_{\text{in}}$). Each node $i \in \mathcal{V}$ of $\mathcal{G}$ holds a local dataset $\mathcal{X}^{(i)}$ consisting of $m_i = 5$ data points $(\mathbf{x}^{(i,1)}, y^{(i,1)}), \dots, (\mathbf{x}^{(i,m_i)}, y^{(i,m_i)})$ with feature vectors $\mathbf{x}^{(i,r)} \in \mathbb{R}^2$ and scalar labels $y^{(i,r)}$, for $r = 1, \dots, m_i$. The feature vectors are generated by i.i.d. realizations (draws) from a standard multivariate normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_{2 \times 2})$. The labels of the datapoints are generated by a linear model

$$y^{(i,r)} = (\bar{\mathbf{w}}^{(i)})^T \mathbf{x}^{(i,r)} + \sigma \varepsilon^{(i,r)}. \quad (19)$$

Here, we used i.i.d. realizations $\varepsilon^{(i,r)}$ of a standard Gaussian random variable $\mathcal{N}(0, 1)$. The true underlying weight vector $\bar{\mathbf{w}}^{(i)}$ is piece-wise constant on the partition $\mathcal{P} = \{\mathcal{C}_1, \mathcal{C}_2\}$.

We assume to have access only to the local datasets $\mathcal{X}^{(i)}$ in a training set $\mathcal{M}$ of relative size $\rho := |\mathcal{M}|/|\mathcal{V}|$. We choose the training set uniformly at random from all subsets of $\mathcal{V}$ with prescribed size $\rho|\mathcal{V}|$ (using rounding when necessary). Given the local datasets in the training set, we learn weights $\hat{\mathbf{w}}^{(i)}$ using Algorithm 1. We used a fixed number K of iterations ($K = 3000$ for Figure 2, $K = 2000$ for Figure 3) and GTV regularization parameter $\lambda = 10^{-2}$ in (5) empirically.

We compute $\text{MSE} := (1/|\mathcal{V}|) \sum_{i \in \mathcal{V}} \|\hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}\|_2^2$ for varying relative size ρ of training set, noise strength σ (see (19)) and intra-cluster edge probability p_{out} . In all experiments, the intra-cluster edge probability is fixed to $p_{\text{in}} = 1/2$. The curves and error bars represent the average and standard deviation of MSE values obtained using 10 simulation runs for a given choice of ρ, p_{out} and σ . Figure 2 depicts the results for the noise less case where $\sigma = 0$ in (19). Figure 3 depicts how the MSE varies with the noise level σ in (19) for a fixed relative training set size $\rho = 0.6$.

Figure 2: MSE of Algorithm 1 learning models (19) in the noiseless case $\sigma = 0$.

Figure 3: MSE of Algorithm 1 in learning (19) with relative training set size $\rho = 0.6$.

6 Conclusion

Using a multi-task learning perspective, we have proposed an studied and efficient method for FL with networks of local datasets. Each local dataset gives rise to a separate learning task which is

to learn the weights of a tailored (“personalized”) model for each local dataset. These individual tasks are coupled via a known network structure. We formalize this multi-task learning problem as RERM with the GTV of model weights as regularizer. We obtain a computationally and statistically appealing FL algorithm by solving this RERM with an established primal-dual method for large-scale optimization. Future research directions include the joint learning of network structure and networked model weights. We also aim at a more fine-grained convergence analysis of our FL algorithm that takes into account the cluster structure of the empirical graph.

References

- [1] S. Cui, A. Hero, Z.-Q. Luo, and J. Moura, Eds., *Big Data over Networks*. Cambridge Univ. Press, 2016.
- [2] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, ser. Adaptive computation and machine learning. MIT Press, 2009.
- [3] S. Banerjee, B. Carlin, and A. Gelfand, *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC, 2015.
- [4] M. Straathof, M. Sinke, T. Roelofs, E. Blezer, R. Sarabdjitsingh, A. van der Toorn, O. Schmitt, W. Otte, and R. Dijkhuizen, “Distinct structure-function relationships across cortical regions and connectivity scales in the rat brain,” *Scientific Reports*, vol. 10, no. 1, p. 56, 2020. [Online]. Available: <https://doi.org/10.1038/s41598-019-56834-9>
- [5] F. Zhao and L. J. Guibas, *Wireless Sensor Networks: An Information Processing Approach*. Amsterdam, The Netherlands: Morgan Kaufmann, 2004.
- [6] M. Wollschlaeger, T. Sauter, and J. Jasperneite, “The future of industrial communication: Automation networks in the era of the internet of things and industry 4.0,” *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, pp. 17–27, 2017.
- [7] M. Satyanarayanan, “The emergence of edge computing,” *Computer*, vol. 50, no. 1, pp. 30–39, Jan. 2017. [Online]. Available: <https://doi.org/10.1109/MC.2017.9>
- [8] H. Ates, A. Yetisen, F. Güder, and C. Dincer, “Wearable devices for the detection of covid-19,” *Nature Electronics*, vol. 4, no. 1, pp. 13–14, 2021. [Online]. Available: <https://doi.org/10.1038/s41928-020-00533-1>
- [9] H. Boyes, B. Hallaq, J. Cunningham, and T. Watson, “The industrial internet of things (iiot): An analysis framework,” *Computers in Industry*, vol. 101, pp. 1–12, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0166361517307285>
- [10] M. E. J. Newman, *Networks: An Introduction*. Oxford Univ. Press, 2010.
- [11] A. Barabási, N. Gulbahce, and J. Loscalzo, “Network medicine: a network-based approach to human disease,” *Nature Reviews Genetics*, vol. 12, no. 56, 2011.
- [12] K. Grantz, H. Meredith, D. Cummings, C. Metcalf, B. Grenfell, J. Giles, S. Mehta, S. Solomon, A. Labrique, N. Kishore, C. Buckee, and A. Wesolowski, “The use of mobile phone data to inform analysis of COVID-19 pandemic epidemiology,” *Nature Communications*, vol. 11, no. 1, p. 4961, 2020. [Online]. Available: <https://doi.org/10.1038/s41467-020-18190-5>
- [13] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. Fort Lauderdale, FL, USA: PMLR, 20–22 Apr 2017, pp. 1273–1282. [Online]. Available: <http://proceedings.mlr.press/v54/mcmahan17a.html>
- [14] P. Kairouz, B. McMahan, B. Avent, A. Bellet, and M. Bennis, “Advances and open problems in federated learning,” *HAL hal-02406503*, 2019.
- [15] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, May 2020.
- [16] Y. Cheng, Y. Liu, T. Chen, and Q. Yang, “Federated learning for privacy-preserving ai,” *Communications of the ACM*, vol. 63, no. 12, pp. 33–36, Dec. 2020.
- [17] N. Agarwal, A. Suresh, F. Yu, S. Kumar, and H. McMahan, “cpSGD: Communication-efficient and differentially-private distributed sgd,” in *Proc. Neural Inf. Proc. Syst. (NIPS)*, 2018.

- [18] F. Sattler, K. Müller, T. Wiegand, and W. Samek, “On the byzantine robustness of clustered federated learning,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, 2020, pp. 8861–8865.
- [19] V. Smith, C.-K. Chiang, M. Sanjabi, and A. Talwalkar, “Federated multi-task learning,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/6211080fa89981f66b1a0c9d55c61d0f-Paper.pdf>
- [20] Z. Xu and K. Kersting, “Multi-task learning with task relations,” in *2011 IEEE 11th International Conference on Data Mining*, 2011, pp. 884–893.
- [21] D. Hallac, J. Leskovec, and S. Boyd, “Network lasso: Clustering and optimization in large graphs,” in *Proc. SIGKDD*, 2015, pp. 387–396.
- [22] A. Jung and N. Tran, “Localized linear regression in networked data,” *IEEE Sig. Proc. Lett.*, vol. 26, no. 7, Jul. 2019.
- [23] N. Tran, H. Ambos, and A. Jung, “Classifying partially labeled networked data via logistic network lasso,” in *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, May 2020, pp. 3832–3836.
- [24] H. Ambos, N. Tran, and A. Jung, “Classifying big data over networks via the logistic network lasso,” in *Proc. 52nd Asilomar Conf. Signals, Systems, Computers*, Oct./Nov. 2018.
- [25] R. Nassif, S. Vlaski, C. Richard, J. Chen, and A. Sayed, “Multitask learning over graphs: An approach for distributed, streaming machine learning,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 14–25, 2020.
- [26] A. Jung, “Networked exponential families for big data over networks,” *IEEE Access*, vol. 8, pp. 202 897–202 909, 2020.
- [27] Y. Zhang and Q. Yang, “Learning sparse task relations in multi-task learning,” in *AAAI*, 2017, pp. 2914–2920. [Online]. Available: <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14430>
- [28] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity. The Lasso and its Generalizations*. CRC Press, 2015.
- [29] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [30] N. Komodakis and J.-C. Pesquet, “Playing with duality: An overview of recent primal?dual approaches for solving large-scale optimization problems,” *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 31–54, 2015.
- [31] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging,” *J. Math. Imag. Vis.*, vol. 40, no. 1, 2011.
- [32] —, “An introduction to continuous optimization for imaging,” *Acta Numer.*, vol. 25, pp. 161–319, 2016. [Online]. Available: <http://dx.doi.org/10.1017/S096249291600009X>
- [33] Z. Huo and H. Huang, “Straggler-agnostic and communication-efficient distributed primal-dual algorithm for high-dimensional data mining,” 2019.
- [34] L. Condat, “A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms,” *Journal of Opt. Th. and App.*, vol. 158, no. 2, pp. 460–479, Aug. 2013.
- [35] A. Chambolle, M. Ehrhardt, P. Richtárik, and C.-B. Schönlieb, “Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications,” *SIAM Journal on Optimization*, vol. 28, no. 4, pp. 2783–2808, 2018. [Online]. Available: <https://doi.org/10.1137/17M1134834>
- [36] P. Combettes and J.-C. Pesquet, “Stochastic quasi-fejér block-coordinate fixed point iterations with random sweeping,” *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 1221–1248, 2015. [Online]. Available: <https://doi.org/10.1137/140971233>
- [37] J. Rasch and A. Chambolle, “Inexact first-order primal–dual algorithms,” *Computational Optimization and Applications*, vol. 76, no. 2, pp. 381–430, 2020. [Online]. Available: <https://doi.org/10.1007/s10589-020-00186-y>

- [38] L. Jacob, J.-P. Vert, and F. Bach, “Clustered multi-task learning: A convex formulation,” in *Advances in Neural Information Processing Systems*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds., vol. 21. Curran Associates, Inc., 2009. [Online]. Available: <https://proceedings.neurips.cc/paper/2008/file/fccb3cdc9acc14a6e70a12f74560c026-Paper.pdf>
- [39] T. Evgeniou, C. Micchelli, and M. Pontil, “Learning multiple tasks with kernel methods,” *Journal of Machine Learning Research*, vol. 6, no. 21, pp. 615–637, 2005. [Online]. Available: <http://jmlr.org/papers/v6/evgeniou05a.html>
- [40] B. Romera-Paredes, A. Argyriou, N. Berthouze, and M. Pontil, “Exploiting unrelated tasks in multi-task learning,” in *Proc. of the Fifteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, N. D. Lawrence and M. Girolami, Eds., vol. 22. La Palma, Canary Islands: PMLR, 21–23 Apr 2012, pp. 951–959. [Online]. Available: <http://proceedings.mlr.press/v22/romera12.html>
- [41] A. Maurer, M. Pontil, and B. Romera-Paredes, “The benefit of multitask representation learning,” *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2853–2884, Jan. 2016.
- [42] F. R. Bach, “Consistency of the group lasso and multiple kernel learning,” *J. Mach. Learn. Research*, vol. 9, pp. 1179–1225, 2008.
- [43] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *J. Roy. Stat. Soc. Ser. B*, vol. 68, no. 1, pp. 49–67, 2006.
- [44] P. Bühlmann and S. van de Geer, *Statistics for High-Dimensional Data*. New York: Springer, 2011.
- [45] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, “An efficient framework for clustered federated learning,” in *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada, 2020.
- [46] F. Sattler, K. Müller, and W. Samek, “Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints,” *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [47] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge Univ. Press, 2004.
- [48] Y. SarcheshmehPour, M. Leinonen, and A. Jung, “Federated learning from big data over networks,” in *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, preprint: <https://arxiv.org/pdf/2010.14159.pdf>, 2021.
- [49] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, ser. Springer Series in Statistics. New York, NY, USA: Springer, 2001.
- [50] A. Jung, N. Tran, and A. Mara, “When is Network Lasso Accurate?” *Front. Appl. Math. Stat.*, vol. 3, Jan. 2018.
- [51] B. Nadler, N. Srebro, and X. Zhou, “Statistical analysis of semi-supervised learning: The limit of infinite unlabelled data,” in *Advances in Neural Information Processing Systems 22*, 2009, pp. 1330–1338.
- [52] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton Univ. Press, 1970.
- [53] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD: Johns Hopkins University Press, 1996.
- [54] A. Jung, “On the duality between network flows and network lasso,” *IEEE Sig. Proc. Lett.*, vol. 27, pp. 940 – 944, 2020.
- [55] —, “Clustering in partially labeled stochastic block models via total variation minimization,” in *Proc. 54th Asilomar Conf. Signals, Systems, Computers*, Pacific Grove, CA, Nov. 2020.
- [56] E. Abbe, “Community detection and stochastic block models: Recent developments,” *Journal of Machine Learning Research*, vol. 18, no. 177, pp. 1–86, 2018.
- [57] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. New York: McGraw-Hill, 1976.
- [58] D. Liu and J. Nocedal, “On the limited memory method for large scale optimization,” *Mathematical Programming B*, vol. 5, no. 3, pp. 503–528, 1989.
- [59] T. Pock and A. Chambolle, “Diagonal preconditioning for first order primal-dual algorithms in convex optimization,” in *IEEE ICCV*, Barcelona, Spain, Nov. 2011.

- [60] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. New York: Springer, 2011.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#) See Section 2 with our conceptual contribution, i.e. formulating federated multi-task learning as GTV minimization. See Section 3 with Algorithm 1 as our main methodological contribution and Section 4 contains our technical contribution which is Theorem 1.
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See Section 1, Section 4 and Section 6.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) See Assumptions 1, 2 and 3 in Section 4.
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) See Appendix A
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) The code is available at <https://github.com/sahelyiyi/FederatedLearning>.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See Section 5.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) See Section 5.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) See Section 5.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#) See Section 5.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)

A Proof of Theorem 1

Consider networked weights $\hat{\mathbf{w}} \in \mathbb{R}^{\mathcal{V} \times n}$ that solve GTV minimization (4). According to Assumption 2, the empirical graph has an intrinsic cluster structure $\mathcal{P}$ such that the (oracle-based) of piece-wise constant weights $\bar{\mathbf{w}}$ (13) approximately minimize each local loss functions $\mathcal{L}^{(i)}(\mathbf{w}^{(i)})$.

The high-level outline of the proof consists of two steps. As the first step, we verify that the estimation error

$$\tilde{\mathbf{w}}^{(i)} := \hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}, \text{ for all nodes } i \in \mathcal{V}, \quad (20)$$

is approximately clustered in the sense that the GTV of the error over inter-cluster edges $\partial\mathcal{P}$ can be bounded via its GTV over intra-cluster edges $\mathcal{E} \setminus \partial\mathcal{P}$. Using this fact, as the second step, we invoke Assumption 3 to upper bound the size of the estimation error.

Step 1. We start with the (rather trivial) observation that any solution $\hat{\mathbf{w}}$ of GTV minimization (4) must satisfy

$$\sum_{i \in \mathcal{M}} \mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) + \lambda \|\hat{\mathbf{w}}\|_{\text{GTV}} \leq \sum_{i \in \mathcal{M}} \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) + \lambda \|\bar{\mathbf{w}}\|_{\text{GTV}}. \quad (21)$$

Here, $\bar{\mathbf{w}}^{(i)}$ denotes the cluster-wise optimal weight vectors (14). We use $\bar{\mathbf{w}}^{(i)}$ as a benchmark for the statistical properties of GTV minimization. Note that the weights $\bar{\mathbf{w}}^{(i)}$ are based on an oracle that provides us the true underlying partition $\mathcal{P}$, which in practice is unknown. Our proof uses that, by their very definition, the weights $\bar{\mathbf{w}}^{(i)}$ do not change along any edge in $\mathcal{E} \setminus \partial\mathcal{P}$ and therefore

$$\|\bar{\mathbf{w}}\|_{\text{GTV}} = \|\bar{\mathbf{w}}\|_{\partial\mathcal{P}} \text{ as well as } \|\tilde{\mathbf{w}}\|_{\mathcal{E} \setminus \partial\mathcal{P}} = \|\hat{\mathbf{w}}\|_{\mathcal{E} \setminus \partial\mathcal{P}}. \quad (22)$$

To further develop (21) we use [47, Sec. 3.1.3]

$$\mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) - \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) \geq \underbrace{(\hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)})}_{\tilde{\mathbf{w}}^{(i)}}^T \nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}). \quad (23)$$

Inserting (23) into (21),

$$\begin{aligned} \|\hat{\mathbf{w}}\|_{\text{GTV}} &\leq \|\bar{\mathbf{w}}\|_{\text{GTV}} + (1/\lambda) \sum_{i \in \mathcal{M}} (\tilde{\mathbf{w}}^{(i)})^T \nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) \\ &\stackrel{(a)}{\leq} \|\bar{\mathbf{w}}\|_{\text{GTV}} + (U/\lambda) \sum_{i \in \mathcal{M}} \|\tilde{\mathbf{w}}^{(i)}\|_2 \\ &\stackrel{(b)}{\leq} \|\bar{\mathbf{w}}\|_{\text{GTV}} + \underbrace{(U\sqrt{M}/\lambda)}_{:=\gamma} \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \end{aligned} \quad (24)$$

Here, step (a) uses the Cauchy–Schwarz inequality and Assumption 2 and step (b) uses the inequality $|a_1| + \dots + |a_M| \leq \sqrt{M} \sqrt{|a_1|^2 + \dots + |a_M|^2}$ (which is a special case of Cauchy–Schwarz inequality).

We further develop (24) using the decomposition property $\|\mathbf{w}\|_{\text{GTV}} = \|\mathbf{w}\|_{\partial\mathcal{P}} + \|\mathbf{w}\|_{\mathcal{E} \setminus \partial\mathcal{P}}$ (see (3)) along with the triangle inequality for the semi-norm $\|\cdot\|_{\text{GTV}}$ (see Assumption 3),

$$\begin{aligned} \|\hat{\mathbf{w}}\|_{\mathcal{E} \setminus \partial\mathcal{P}} &\leq \|\bar{\mathbf{w}}\|_{\text{GTV}} - \|\hat{\mathbf{w}}\|_{\partial\mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \\ &\stackrel{(a)}{=} \|\bar{\mathbf{w}}\|_{\partial\mathcal{P}} - \|\hat{\mathbf{w}}\|_{\partial\mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \\ &\leq \|\tilde{\mathbf{w}}\|_{\partial\mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}}. \end{aligned} \quad (25)$$

According to (22), the LHS of (25) is equal to $\|\tilde{\mathbf{w}}\|_{\mathcal{E} \setminus \partial\mathcal{P}}$ and therefore

$$\begin{aligned} \|\tilde{\mathbf{w}}\|_{\mathcal{E} \setminus \partial\mathcal{P}} &\leq \|\bar{\mathbf{w}}\|_{\partial\mathcal{P}} - \|\hat{\mathbf{w}}\|_{\partial\mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \\ &\leq \|\tilde{\mathbf{w}}\|_{\partial\mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}}. \end{aligned} \quad (26)$$

Here, we have used the shorthand $\gamma := U\sqrt{M}/\lambda$. Choosing a sufficiently large λ (making γ small), (26) implies that the error weights $\tilde{\mathbf{w}}^{(i)} = \hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}$ change mainly over boundary edges $\partial\mathcal{P}$ (16). In other words, the estimation error is approximately piece-wise constant over the clusters $\mathcal{C}_l$ in the partition $\mathcal{P} = \{\mathcal{C}_1, \dots, \mathcal{C}_P\}$.

Step 2. The second final step in our proof of Theorem 1 is to verify that the estimation error $\tilde{\mathbf{w}} = \hat{\mathbf{w}} - \bar{\mathbf{w}}$ cannot be too large. Another application of the decomposition property and the triangle inequality for the GTV semi-norm to (21) yields

$$\underbrace{\sum_{i \in \mathcal{M}} \mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) - \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)})}_{f(\hat{\mathbf{w}}) - f(\bar{\mathbf{w}})} \leq \lambda \|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}}. \quad (27)$$

We next lower bound the LHS of (27) in terms of the norm $\|\tilde{\mathbf{w}}\|_{\mathcal{M}}$. To this end we apply Taylor's theorem [57] separately to each local loss function,

$$\begin{aligned} \mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) &= \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) + (\nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}))^T (\hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}) \\ &\quad + (1/2) (\bar{\mathbf{w}}^{(i)} - \hat{\mathbf{w}}^{(i)})^T \nabla^2 \mathcal{L}^{(i)}(\mathbf{z}) (\bar{\mathbf{w}}^{(i)} - \hat{\mathbf{w}}^{(i)}) \\ &\text{for } \mathbf{z} = \alpha \bar{\mathbf{w}}^{(i)} - (1 - \alpha) \hat{\mathbf{w}}^{(i)}, \text{ with some } \alpha \in [0, 1]. \end{aligned} \quad (28)$$

Inserting our assumptions $\|\nabla \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)})\|_2 \leq U$ (see (15)) and $\nabla^2 \mathcal{L}^{(i)}(\mathbf{z}) \succeq 2C\mathbf{I}$ (see (12)) into (28),

$$\mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) - \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) \geq -U \|\hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}\|_2 + C \|\hat{\mathbf{w}}^{(i)} - \bar{\mathbf{w}}^{(i)}\|_2^2. \quad (29)$$

Summing (29) over all nodes $i \in \mathcal{M}$,

$$\sum_{i \in \mathcal{M}} \mathcal{L}^{(i)}(\hat{\mathbf{w}}^{(i)}) - \mathcal{L}^{(i)}(\bar{\mathbf{w}}^{(i)}) \geq -U\sqrt{M} \|\tilde{\mathbf{w}}\|_{\mathcal{M}} + C \|\tilde{\mathbf{w}}\|_{\mathcal{M}}^2. \quad (30)$$

Combining (30) with (27),

$$C \|\tilde{\mathbf{w}}\|_{\mathcal{M}}^2 \leq \lambda \|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}} + U\sqrt{M} \|\tilde{\mathbf{w}}\|_{\mathcal{M}}. \quad (31)$$

To obtain another relation between $\|\tilde{\mathbf{w}}\|_{\mathcal{M}}$ and $\|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}}$, we insert (26) into (17) yielding

$$(L-1) \|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}} \leq (K+\gamma) \|\tilde{\mathbf{w}}\|_{\mathcal{M}}. \quad (32)$$

Inserting (32) into (31),

$$\|\tilde{\mathbf{w}}\|_{\mathcal{M}} \leq (\lambda/C) \left((K+\gamma)/(L-1) + \underbrace{U\sqrt{M}/\lambda}_{=\gamma} \right). \quad (33)$$

Putting together the pieces,

$$\begin{aligned} \|\tilde{\mathbf{w}}\|_{\text{GTV}} &= \|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}} + \|\tilde{\mathbf{w}}\|_{\mathcal{E} \setminus \partial \mathcal{P}} \\ &\stackrel{(26)}{\leq} 2\|\tilde{\mathbf{w}}\|_{\partial \mathcal{P}} + \gamma \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \\ &\stackrel{(32)}{\leq} (2(K+\gamma)/(L-1) + \gamma) \|\tilde{\mathbf{w}}\|_{\mathcal{M}} \\ &\stackrel{(33)}{\leq} (\lambda/C) (2(K+\gamma)/(L-1) + \gamma) \left((K+\gamma)/(L-1) + \gamma \right). \end{aligned} \quad (34)$$

We obtain (18) by minimizing the bound (34) with respect to the GTV regularization parameter λ .

B Some Special Cases

Networked Linear Regression. We can use Algorithm 1 to learn a separate linear predictor $\hat{y} = h(\mathbf{x}; \mathbf{w}^{(i)}) = (\mathbf{w}^{(i)})^T \mathbf{x}$, for each $i \in \mathcal{V}$. For a labeled dataset $\mathcal{X}^{(i)}$, with some $i \in \mathcal{M}$, we measure the quality of the linear predictor $h(\mathbf{x}; \mathbf{w}^{(i)})$ via the squared error loss

$$\mathcal{L}(\mathcal{X}^{(i)}; \mathbf{v}) := (1/m_i) \sum_{r=1}^{m_i} (y_r^{(i)} - \mathbf{v}^T \mathbf{x}^{(i,r)})^2 \quad (35)$$

Inserting (35) into (10) yields the primal update operator

$$\mathcal{PU}^{(i)}(\mathbf{v}) = [\mathbf{I} + 2\tau_i \tilde{\mathbf{Q}}^{(i)}]^{-1} [\mathbf{v} + 2\tau_i \tilde{\mathbf{y}}^{(i)}]. \quad (36)$$

Here, we used $\tilde{\mathbf{Q}}^{(i)} := (1/m_i)(\mathbf{X}^{(i)})^T \mathbf{X}^{(i)}$ with $\mathbf{X}^{(i)} := (\mathbf{x}^{(i,1)}, \dots, \mathbf{x}^{(i,m_i)})^T$ and the normalized label vector $\tilde{\mathbf{y}}^{(i)} := (1/m_i)(\mathbf{X}^{(i)})^T \mathbf{y}^{(i)}$ with $\mathbf{y}^{(i)} := (y_1^{(i)}, \dots, y_{m_i}^{(i)})^T$.

Networked Logistic Regression. We can use Algorithm 1 to learn linear classifiers $h^{(i)}(\mathbf{x}; \mathbf{w}^{(i)}) = \mathbf{x}^T \mathbf{w}^{(i)}$ for $i \in \mathcal{V}$. To measure the quality of $h(\mathbf{x}; \mathbf{v})$ at node i , we use the logistic loss given by (for label values $y \in \{-1, 1\}$) $\mathcal{L}(\mathcal{X}^{(i)}; \mathbf{v}) = (1/m_i) \sum_{r=1}^{m_i} \log(1 + \exp(-y_r^{(i)} \mathbf{v}^T \mathbf{x}^{(i,r)}))$. The corresponding primal update operator is

$$\mathcal{PU}^{(i)}(\mathbf{v}) = \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} (1/m_i) \sum_{r=1}^{m_i} \log(1 + \exp(-y_r^{(i)} \mathbf{z}^T \mathbf{x}^{(i,r)})) + (1/2\tau_i) \|\mathbf{z} - \mathbf{v}\|^2. \quad (37)$$

While lacking a closed-form solution, (37) can be solved numerically with efficient optimization methods such as L-BFGS [58].

Networked Regularized Linear Regression. Many application domains generate high-dimensional local datasets $\mathcal{X}^{(i)}$, with a number of datapoints much smaller than number of features, $m_i \ll n$ [44, 28]. Regularized linear regression methods allow to handle this high-dimensional regime by augmenting the squared error loss (35) with a regularization term $\eta \mathcal{R}(\mathbf{w})$. Two well-known special cases are ridge regression, using $\mathcal{R}(\mathbf{w}) = \|\mathbf{w}\|_2^2$, and the Lasso, which is obtained for $\mathcal{R}(\mathbf{w}) = \|\mathbf{w}\|_1$ [28]. Plugging the regularized loss into (10) yields the primal update operator

$$\mathcal{PU}^{(i)}(\mathbf{v}) = \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} (1/m_i) \sum_{r=1}^{m_i} (\mathbf{z}^T \mathbf{x}^{(i,r)} - y_r^{(i)})^2 + \eta \mathcal{R}(\mathbf{z}) + (1/2\tau_i) \|\mathbf{z} - \mathbf{v}\|^2. \quad (38)$$

C Derivation of Algorithm 1

A necessary and sufficient condition for $\hat{\mathbf{w}}$ to solve (5) and $\hat{\mathbf{u}}$ to solve (6) is

$$-\mathbf{D}^T \hat{\mathbf{u}} \in \partial f(\hat{\mathbf{w}}), \text{ and } \mathbf{D} \hat{\mathbf{w}} \in \partial g^*(\hat{\mathbf{u}}). \quad (39)$$

We can rewrite the optimality condition (39) as [32]

$$\mathbf{M}^{-1} \begin{pmatrix} \partial f & \mathbf{D}^T \\ -\mathbf{D} & \partial g^* \end{pmatrix} \begin{pmatrix} \hat{\mathbf{w}} \\ \hat{\mathbf{u}} \end{pmatrix} \ni \mathbf{0} \text{ with } \mathbf{M} := \begin{pmatrix} \mathbf{T}^{-1} & -\mathbf{D}^T \\ -\mathbf{D} & \Sigma^{-1} \end{pmatrix}. \quad (40)$$

The condition (40) involves the invertible matrix $\mathbf{M}$ which is tuned to the structure of the empirical graph via its incidence matrix $\mathbf{D}$. It is constructed using two positive definite block-diagonal matrices

$$(\Sigma)_{e,e} := \sigma_e \mathbf{I}_n, \text{ for } e \in \mathcal{E}, (\mathbf{T})_{i,i} := \tau_i \mathbf{I} \text{ for } i \in \mathcal{V},$$

with $\sigma_e := 1/2$ for $e \in \mathcal{E}$ and $\tau_i := 1/|\mathcal{N}_i|$ for $i \in \mathcal{V}$.

In principle, the condition (40) can use an arbitrary invertible matrix $\mathbf{M}$ to characterize solution of (5). However the specific choice for $\mathbf{M}$ in (40) serves two main purposes. The off-diagonal blocks partially decouples solving (5) into optimizing (regularized variants of) $f(\mathbf{w})$ and $g(\mathbf{u})$ separately. These separate optimization steps are coupled via simple applications of the block-incidence matrix $\mathbf{D}$ (1) and its transpose $\mathbf{D}^T$. The main diagonal blocks of $\mathbf{M}$ serve as a preconditioning of (5) for irregular network structures with highly varying node degrees [59].

We solve the optimality conditions (40) using a proximal point algorithm [60, Thm. 23.41]. This algorithm consists of the iterations

$$\begin{pmatrix} \hat{\mathbf{w}}_{k+1} \\ \hat{\mathbf{u}}_{k+1} \end{pmatrix} := \left(\mathbf{I} + \mathbf{M}^{-1} \begin{pmatrix} \partial f & \mathbf{D}^T \\ -\mathbf{D} & \partial g^* \end{pmatrix} \right)^{-1} \begin{pmatrix} \hat{\mathbf{w}}_k \\ \hat{\mathbf{u}}_k \end{pmatrix}. \quad (41)$$

The iteration (41) can be rewritten as [32]

$$\hat{\mathbf{w}}_{k+1} := (\mathbf{I} + \mathbf{T} \partial f)^{-1} (\hat{\mathbf{w}}_k - \mathbf{T} \mathbf{D}^T \hat{\mathbf{u}}_k) \quad (42)$$

$$\hat{\mathbf{u}}_{k+1} := (\mathbf{I} + \Sigma \partial g^*)^{-1} (\hat{\mathbf{u}}_k + \Sigma \mathbf{D} (2\hat{\mathbf{w}}_{k+1} - \hat{\mathbf{w}}_k)). \quad (43)$$

The sequences $\widehat{\mathbf{w}}_k$ and $\widehat{\mathbf{u}}_k$ obtained by iterating (42), (43) converge to solutions of (5) and (6). This convergence is guaranteed for any initialization $\widehat{\mathbf{w}}_0$ and $\widehat{\mathbf{u}}_0$ [59].

To compute the updates (42), we use the identity [59]

$$(\mathbf{I} + \mathbf{S}\partial h)^{-1}(\mathbf{a}) := \underset{\mathbf{b}}{\operatorname{argmin}} h(\mathbf{b}) + (1/2)\|\mathbf{a} - \mathbf{b}\|_{\mathbf{S}^{-1}}^2, \quad (44)$$

which is valid for any convex function h and positive definite matrix $\mathbf{S}$. By inserting (2) into (42), and using (44), we can decompose the primal update into node-wise updates

$$\mathbf{w}_{k+1}^{(i)} := \mathcal{PU}^{(i)}\left\{\mathbf{w}_k^{(i)} - \tau_i \sum_{e \in \mathcal{E}} D_{e,i} \mathbf{u}_k^{(e)}\right\}, \text{ for each } i \in \mathcal{M}, \quad (45)$$

and $\mathbf{w}_{k+1}^{(i)} := \mathbf{w}_k^{(i)} - \tau_i \sum_{e \in \mathcal{E}} D_{e,i} \mathbf{u}_k^{(e)}$ for $i \in \mathcal{V} \setminus \mathcal{M}$. Here, we used the primal update operator

$$\mathcal{PU}^{(i)}\{\mathbf{v}\} := \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \mathcal{L}^{(i)}(\mathbf{z}) + (1/2\tau_i)\|\mathbf{v} - \mathbf{z}\|^2. \quad (46)$$

The dual update (43) decomposes into independent updates for each edge $e \in \mathcal{E}$. Inserting (7) into (43),

$$\mathbf{u}_{k+1}^{(e)} := \mathcal{DU}^{(e)}\left\{\mathbf{u}_k^{(e)} + \sigma_e \left(2(\mathbf{w}_{k+1}^{(e+)} - \mathbf{w}_{k+1}^{(e-)}) - (\mathbf{w}_k^{(e+)} - \mathbf{w}_k^{(e-)})\right)\right\}, \text{ for each } e \in \mathcal{E}, \quad (47)$$

with the dual update operator

$$\mathcal{DU}^{(e)}\{\mathbf{v}\} := \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \lambda A_e \phi^*(\mathbf{z}/(\lambda A_e)) + (1/2\sigma_e)\|\mathbf{v} - \mathbf{z}\|^2. \quad (48)$$

Here, we used the convex conjugate $\phi^*(\mathbf{v}) := \sup_{\mathbf{z} \in \mathbb{R}^n} \mathbf{v}^T \mathbf{z} - \phi(\mathbf{z})$ of the GTV penalty function $\phi(\mathbf{v})$ (see (3)).