



Semi-Supervised Learning in Network-Structured Data via Total Variation Minimization

Alexander Jung , Alfred O. Hero, III, *Fellow, IEEE*, Alexandru Cristian Mara , Saeed Jahromi, Ayelet Heimowitz , and Yonina C. Eldar , *Fellow, IEEE*

Abstract—We provide an analysis and interpretation of total variation (TV) minimization for semi-supervised learning from partially-labeled network-structured data. Our approach exploits an intrinsic duality between TV minimization and network flow problems. In particular, we use Fenchel duality to establish a precise equivalence of TV minimization and a minimum cost flow problem. This provides a link between modern convex optimization methods for non-smooth Lasso-type problems and maximum flow algorithms. We show how a primal-dual method for TV minimization can be interpreted as distributed network optimization. Moreover, we derive a condition on the network structure and available label information that ensures that TV minimization accurately learns (approximately) piece-wise constant graph signals. This condition depends on the existence of sufficiently large network flows between labeled data points. We verify our analysis in numerical experiments.

Index Terms—Machine learning, semisupervised learning, optimization, big data applications, network theory (graphs).

I. INTRODUCTION

WE CONSIDER machine learning using partially labeled network-structured datasets that arise in signal processing [2], image processing [3], social networks, internet and bioinformatics [4], [5]. Such data can be described by an “empirical graph,” whose nodes represent individual data points that

are connected by edges if they are “similar” in an application-specific sense. The notion of similarity can be based on physical proximity (in time or space), physical connection (communication networks), or statistical dependency (probabilistic graphical models) [6]–[8].

Besides graph structure, datasets carry additional information in the form of labels associated with individual data points. In a social network, we might define the personal preference for some product as the label associated with a data point (user profile). Acquiring labels is often costly and requires manual labor or experiment design. Therefore, we assume to have access to the labels of only a few data points of a small “training set.” This paper aims at learning or recovering the labels of all data points based on the knowledge of the labels of only a few data points.

Network models lend naturally to scalable algorithms via message passing over the empirical graph [9]. Moreover, semi-supervised learning (SSL) methods allow to overcome the absence of label information via borrowing statistical strength between connected data points [5]. Indeed, many SSL methods rely on a cluster assumption: labels of close-by data points are similar [5], [10]–[12]. This assumption is at the heart of many successful methods in graph signal processing [13], imaging [14], trend filtering [15], anomaly detection [16], information retrieval [17], and social networks [4]. We implement this cluster assumption by treating the labels of data points as graph signals with a small TV, which is the sum of the absolute values of signal differences along the edges in the empirical graph. This turns SSL into a TV minimization problem [1], [14], [15], [18], [19].

TV minimization problems in grid-structured image data have been studied in [14], [20]. For arbitrary networks, [15] studied the statistical properties of TV minimization when applied to noisy but fully observed labels. Considering partially labeled datasets (with arbitrary network structure), [1], [18], [21] offer sufficient conditions on the network structure and label information such that TV minimization accurately learns the labels of all data points. These conditions are somewhat difficult to verify, as they involve the (unknown) cluster structure of the empirical graph. We present a condition, which can be verified by network flow algorithms (see Sections V and VI-A), ensuring TV minimization to accurately learn labels that form a piece-wise constant graph signal.

The cluster assumption used in this paper is different from the smoothness condition widely used in graph signal processing [5], [10]. The smoothness condition requires connected

Manuscript received January 30, 2019; revised July 17, 2019 and October 25, 2019; accepted November 4, 2019. Date of publication November 15, 2019; date of current version December 3, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Mark A. Davenport. This work was supported in part by the Vienna Science Fund (WWTF) Grant ICT15-119 and in part by the U.S. Army Research Office Grant W911NF-15-1-0479. This paper was presented in part at the 12th International Conference on Sampling Theory and Applications, Tallinn, Estonia, July 2017. (*Corresponding author: Alexander Jung.*)

A. Jung is with the Department of Computer Science, Aalto University, Espoo 02150, Finland (e-mail: alexander.jung@aalto.fi).

A. O. Hero, III is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: hero@eecs.umich.edu).

A. C. Mara is with the Department of Electronics and Information Systems, Ghent University, Ghent 9052, Belgium (e-mail: alexandru.mara@ugent.be).

S. Jahromi is with the Faculty of Information Technology, Monash University, Melbourne, VIC 3800, Australia (e-mail: saeed.basirianjahromi@monash.edu).

A. Heimowitz is with The Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544 USA (e-mail: ayeltg@gmail.com).

Y. C. Eldar is with the Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 7610001, Israel (e-mail: yonina.eldar@weizmann.ac.il).

Digital Object Identifier 10.1109/TSP.2019.2953593

nodes to have similar labels by forcing them to live in a small subspace spanned by a few eigenvectors of the graph Laplacian. In contrast, the cluster assumption allows the labels to vary significantly over edges between two different clusters (see Section II-C for more details).

Minimizing TV as well as minimizing the Laplacian quadratic form are both special cases of p -Laplacian minimization [12], [22]. However, their statistical and computational properties are quite different. While the Laplacian quadratic form is a smooth convex function, the TV is a non-smooth convex function that requires more advanced optimization techniques such as proximal methods [14], [23]. Statistically, TV-based learning may be accurate in cases where the Laplacian quadratic form minimizer fails.

We analyze TV minimization using a variant of the nullspace property which provides necessary and sufficient conditions for the success of ℓ_1 based methods [24]–[27]. In a similar spirit [28] studies recovery of sparse signals defined on the edges of the empirical graph. In contrast, we study piece-wise constant signals defined on nodes.

This paper continues our studies [1], [18], [19] of statistical and computational aspects of SSL via TV regularization. The central theme of this paper is the duality between TV minimization and network flow problems. The relation between network flow problems and energy minimization has been studied mainly for discrete-valued graph signals [29]–[31]. However, it is not obvious how to generalize these methods to real-valued graph signals.

It turns out that the duality between TV minimization and network flow problems can be established in an elegant fashion using the concept of convex conjugate functions. This duality allows us to apply efficient convex optimization methods for TV minimization (see Algorithm 1) to solve network flow problems and, in the other direction, unleashes existing network-flow algorithms [32] for TV minimization.

Our detailed contributions are:

- Our main result is Proposition 1, which states that the dual of TV minimization is equivalent to a minimum-cost network flow problem (see Section III).
- An immediate consequence is Corollary 2, which characterizes the solutions of TV minimization. In contrast to our previous work, Corollary 2 does not involve any signal model, such as piece-wise constant signals.
- We provide an interpretation of a message passing algorithm [19, Algorithm 2] for TV minimization as distributed network flow optimization (see Section IV).
- Proposition 4 provides a new condition ensuring that TV minimization is accurate. In contrast to previous work [1], [18], this condition can be verified easily using existing network-flow algorithms (see Section VI-A).
- We verify our theoretical analysis of TV minimization by several numerical experiments (see Section VI).

Outline: In Section II, we formulate SSL for network-structured data as a convex TV minimization problem. We then discuss in Section III how a dual problem of TV minimization can be defined. Exploiting the relation between TV minimization and its dual, we discuss in Section IV how to

apply a particular instance of a proximal method [23] to obtain a solution to TV minimization (and its dual). As detailed in Section IV, the resulting algorithm can be implemented as message passing on the empirical graph. In Section V, we present a sufficient condition on the available label information and the empirical graph such that TV minimization delivers accurate label estimates. Numerical experiments are discussed in Section VI.

II. PROBLEM FORMULATION

We formalize SSL with network-structured data as an optimization problem. Section II-A introduces relevant concepts of graph theory. Section II-B introduces the cluster assumption using graph signals with a small TV. A particular class of such graph signals is constituted by piece-wise constant graph signals as defined in Section II-B. The cluster assumption leads naturally to a formulation of SSL as a TV minimization problem, which we define and discuss in Section II-C.

Let us fix some notation. Given a vector $\mathbf{x} = (x_1, \dots, x_n)^T$, we define the norms $\|\mathbf{x}\|_1 := \sum_{l=1}^n |x_l|$ and $\|\mathbf{x}\|_\infty := \max_{i=1, \dots, n} |x_i|$. The signum $\text{sign}\{\mathbf{x}\}$ of a vector $\mathbf{x} = (x_1, \dots, x_d)$ is the vector $(\text{sign}(x_1), \dots, \text{sign}(x_d)) \in \mathbb{R}^d$ with $\text{sign}(x_i) = 1$ for $x_i > 0$, $\text{sign}(x_i) = -1$ for $x_i < 0$.

The spectral norm of a matrix \mathbf{A} is denoted $\|\mathbf{A}\|_2 := \sup_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$. For a positive semidefinite (psd) matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$, with spectral decomposition $\mathbf{Q} = \mathbf{U}\mathbf{S}\mathbf{U}^T$ with the diagonal matrix $\mathbf{S} = \text{diag}\{s_i\}_{i=1}^n$. The square root of psd \mathbf{Q} is $\mathbf{Q}^{1/2} := \mathbf{U}\mathbf{S}^{1/2}\mathbf{U}^T$ with $\mathbf{S}^{1/2} := \text{diag}\{\sqrt{s_i}\}_{i=1}^n$. For a given psd \mathbf{Q} we define the norm $\|\mathbf{x}\|_{\mathbf{Q}} := \sqrt{\mathbf{x}^T \mathbf{Q} \mathbf{x}}$.

The subdifferential of a function $g(\mathbf{x})$ at $\mathbf{x}_0 \in \mathbb{R}^n$ is

$$\partial g(\mathbf{x}_0) := \{\mathbf{y} \in \mathbb{R}^n : g(\mathbf{x}) \geq g(\mathbf{x}_0) + \mathbf{y}^T (\mathbf{x} - \mathbf{x}_0) \text{ for any } \mathbf{x}\},$$

and its convex conjugate function is defined as [33]

$$g^*(\hat{\mathbf{y}}) := \sup_{\mathbf{y} \in \mathbb{R}^n} \mathbf{y}^T \hat{\mathbf{y}} - g(\mathbf{y}). \quad (1)$$

A. The Empirical Graph

Consider a dataset of N data points that can be represented as supported at the nodes of a simple undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, where \mathcal{V} are nodes, \mathcal{E} are edges and \mathbf{W} are edge weights. Following [5], we refer to the graph \mathcal{G} as the empirical graph associated with the dataset.

The nodes $i \in \mathcal{V} = \{1, \dots, N\}$ of the empirical graph \mathcal{G} represent the N individual data points. In many applications, the goal is to determine (or infer) some relevant property encoded as a numeric label x_i associated with the node $i \in \mathcal{V}$. The labels could represent instantaneous amplitudes of an audio signal, the greyscale values of image pixels, or the probabilities of social network members taking a particular action. The labels x_i define a graph signal $\mathbf{x} = (x_1, \dots, x_N)^T \in \mathbb{R}^N$ over the empirical graph with the signal value at node i given by the label x_i .

The undirected edges $\{i, j\} \in \mathcal{E}$ of the empirical graph \mathcal{G} connect data points which are considered similar (in some domain-specific sense). It will be convenient to represent the edges by the numbers $\{1, \dots, E = |\mathcal{E}|\}$.

For an edge $\{i, j\} \in \mathcal{E}$, the nonzero value $W_{i,j} > 0$ represents the strength of the connection $\{i, j\} \in \mathcal{E}$. The edge set \mathcal{E} is encoded in the non-zero pattern of the weight matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$,

$$\{i, j\} \in \mathcal{E} \text{ if and only if } W_{i,j} > 0. \quad (2)$$

The neighborhood $\mathcal{N}(i)$ and weighted degree (strength) d_i of node $i \in \mathcal{V}$ are defined, respectively, as

$$\mathcal{N}(i) := \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}\}, d_i := \sum_{j \in \mathcal{N}(i)} W_{i,j}. \quad (3)$$

The maximum (weighted) node degree is

$$d_{\max} := \max_{i \in \mathcal{V}} d_i \stackrel{(3)}{=} \max_{i \in \mathcal{V}} \sum_{j \in \mathcal{N}(i)} W_{i,j}. \quad (4)$$

Without loss of generality we consider only datasets whose empirical graph does not contain isolated nodes, i.e., we assume that $d_i > 0$ for every node $i \in \mathcal{V}$.

For a given undirected empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, we orient the undirected edge $\{i, j\}$ by defining the head as $e^+ = \min\{i, j\}$ and the tail as $e^- = \max\{i, j\}$. The undirected edge $\{i, j\}$ with nodes $i < j$ becomes the directed edge (i, j) . We use \mathcal{G} and \mathcal{E} to also denote the oriented empirical graph and its directed edges, respectively. The *incidence matrix* $\mathbf{D} \in \mathbb{R}^{E \times N}$ of the empirical graph \mathcal{G} is

$$D_{e,i} = \begin{cases} W_e & \text{if } i = e^+ \\ -W_e & \text{if } i = e^- \\ 0 & \text{else.} \end{cases} \quad (5)$$

The rows of \mathbf{D} correspond to the edges $e \in \mathcal{E}$ while the columns represent nodes $i \in \mathcal{V}$ of the empirical graph \mathcal{G} . The row representing $e = \{i, j\}$ contains exactly two non-zero entries in the columns corresponding to the nodes $i, j \in \mathcal{V}$. It will be convenient to define the directed neighbourhoods (see (3)) of a node $i \in \mathcal{V}$ as

$$\begin{aligned} \mathcal{N}^+(i) &:= \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}, i < j\}, \text{ and} \\ \mathcal{N}^-(i) &:= \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}, i > j\}. \end{aligned} \quad (6)$$

B. Cluster Assumption

We assume that labels x_i are known at only a few nodes $i \in \mathcal{V}$ of a (small) training set $\mathcal{M} \subseteq \mathcal{V}$ (see Fig. 1). Our goal is then to learn the unknown labels x_i for all data points $i \in \mathcal{V} \setminus \mathcal{M}$ outside the training set. This learning problem, which is known as SSL, translates to a graph signal recovery problem within our setting.

Given the signal samples x_i for data points $i \in \mathcal{M}$ in the training set, we want to recover the entire graph signal $\mathbf{x} \in \mathbb{R}^N$. This learning (or recovery) problem is feasible if the underlying graph signal \mathbf{x} has a known structure. As mentioned above, a particular structure is obtained if the labels x_i conform with the cluster structure of the empirical graph \mathcal{G} . Consider the graph signal $\mathbf{x} \in \mathbb{R}^N$ constituted by the (mostly unknown) labels x_i of the data points $i \in \mathcal{V}$. The cluster assumption requires similar signal values $x_i \approx x_j$ at nodes $i, j \in \mathcal{V}$ in the same well-connected subset (cluster).




Fig. 1. Empirical graph \mathcal{G} whose nodes \mathcal{V} are grouped into two clusters \mathcal{C}_1 and \mathcal{C}_2 forming the partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2\}$. The boundary of the partition is $\partial\mathcal{F} = \{\{i, j\}\}$ having weight $W_{i,j} = 1/2$. The edges $e \in \mathcal{E}$ connecting nodes within the same cluster have weight $W_e = 1$. The nodes belonging to the training set \mathcal{M} are shaded.

We measure the “clusteredness” of a graph signal \mathbf{x} using the weighted TV [15], [34]

$$\|\mathbf{x}\|_{\text{TV}} := \sum_{\{i,j\} \in \mathcal{E}} W_{i,j} |x_j - x_i|. \quad (7)$$

As the notation indicates, $\|\mathbf{x}\|_{\text{TV}}$ defines a seminorm for graph signals \mathbf{x} . It is only a seminorm since it is zero also for non-zero (but constant) graph signals. The incidence matrix \mathbf{D} (5) of the (oriented) empirical graph \mathcal{G} allows us to represent the TV of a graph signal \mathbf{x} as

$$\|\mathbf{x}\|_{\text{TV}} = \|\mathbf{D}\mathbf{x}\|_1. \quad (8)$$

Using the TV (7) to guide learning (signal recovery) methods turns out to be useful statistically and computationally. Indeed, as we discuss below, minimizing TV results in labels (signals) which are constant over well-connected subsets (clusters) of data points. Moreover, TV minimization can be implemented as highly scalable message passing over the underlying empirical graph (see Algorithm 2).

The most simple model for graph signals conforming with the cluster assumption are piece-wise constant signals [2]

$$x_i = \sum_{l=1}^{|\mathcal{F}|} a_l \mathcal{I}_{\mathcal{C}_l}[i] \text{ with } a_l \in \mathbb{R}, \mathcal{I}_{\mathcal{C}_l}[i] := \begin{cases} 1 & \text{for } i \in \mathcal{C}_l \\ 0 & \text{else.} \end{cases} \quad (9)$$

The signal model (9) uses an arbitrary but fixed partition

$$\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$$

constituted by disjoint clusters $\mathcal{C}_l \subseteq \mathcal{V}$ (see Fig. 1). Our analysis will be applicable for an arbitrary choice of the partition underlying the signal model (9). However, our results are most useful for partitions which consist of well connected clusters (see Definition 3).

We emphasize that the learning algorithm we propose in Section IV does not require knowledge of the partition \mathcal{F} underlying the signal model (9). The partition is only required for the analysis of the learning accuracy of this algorithm (see Section V).

The signal model (9) is an idealization which crucially simplifies the analysis of the statistical properties of TV minimization (see Section II-C). The graph signals arising in many applications will typically not be perfectly constant over clusters.

However, Theorem 3 remains useful as long as the data (labels) can be well approximated by a piece-wise constant graph signal (9).

In Section V we characterize (see Definition 3) those partitions \mathcal{F} , used in the model (9), which allow for accurate recovery of a (approximately) piece-wise graph signal from its values x_i at the nodes $i \in \mathcal{M}$ of the training set. Our results indicate that piece-wise constant signals (9) can be learned accurately if the partition \mathcal{F} has a boundary with small weights. The boundary $\partial\mathcal{F}$ of \mathcal{F} consists of the edges connecting nodes from different clusters, i.e.,

$$\partial\mathcal{F} := \{\{i, j\} \in \mathcal{E} \text{ with } i \in \mathcal{C}_l \text{ and } j \in \mathcal{C}_{l'} \neq \mathcal{C}_l\}.$$

The boundary $\partial\mathcal{F}$ is the union of the cluster boundaries

$$\partial\mathcal{C}_l := \{\{i, j\} \in \mathcal{E} \text{ with } i \in \mathcal{C}_l \text{ and } j \in \mathcal{V} \setminus \mathcal{C}_l\}. \quad (10)$$

Recovering a piece-wise constant graph signal (9) may seem trivial given the availability of efficient clustering methods [35]–[37]. Indeed, it is natural to first obtain the partition \mathcal{F} underlying (9) using some clustering method and then perform cluster-wise averaging in order to obtain an estimate for the coefficients a_l in (9). Despite the conceptual simplicity of this approach, it has some challenges. Most existing clustering methods involve design parameters such as the number of clusters or distribution parameters of probabilistic (stochastic block) models. The proper choice (or learning) of these parameters can be non-trivial. Moreover, clustering methods do not exploit label information.

In what follows, we show how the recovery problem lends naturally to a TV minimization problem which, in turn, can be solved by efficient convex optimization methods. The resulting algorithm (Algorithm 1) does not involve any design parameters and can be implemented as scalable message passing (Algorithm 2) on the empirical graph.

C. TV Minimization

The TV of a piece-wise constant graph signal (9) is

$$\begin{aligned} \|\mathbf{x}\|_{\text{TV}} &\stackrel{(7)}{=} \sum_{\{i,j\} \in \mathcal{E}} W_{i,j} |x_j - x_i| \\ &\stackrel{(9)}{=} \sum_{\{i,j\} \in \partial\mathcal{F}} W_{i,j} |x_j - x_i| \\ &\leq \left(\sum_{\{i,j\} \in \partial\mathcal{F}} W_{i,j} \right) \max_{l,l' \in \{1, \dots, |\mathcal{F}|\}} |a_l - a_{l'}|. \quad (11) \end{aligned}$$

Thus, if the partition \mathcal{F} has a small weighted boundary $\sum_{\{i,j\} \in \partial\mathcal{F}} W_{i,j}$, the graph signals (9) have a small TV $\|\mathbf{x}\|_{\text{TV}}$ due to (11).

A sensible strategy for learning a piece-wise constant graph signal is therefore via minimizing the TV $\|\tilde{\mathbf{x}}\|_{\text{TV}}$ among all graph signals which are consistent with the known labels

$\{x_i\}_{i \in \mathcal{M}}$. This is formulated as the optimization problem

$$\begin{aligned} \hat{\mathbf{x}} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \underbrace{\sum_{\{i,j\} \in \mathcal{E}} W_{i,j} |\tilde{x}_j - \tilde{x}_i|}_{= \|\tilde{\mathbf{x}}\|_{\text{TV}}} \text{ s.t. } \tilde{x}_i = x_i \text{ for all } i \in \mathcal{M} \\ \stackrel{(8)}{=} \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \|\mathbf{D}\tilde{\mathbf{x}}\|_1 \text{ s.t. } \tilde{x}_i = x_i \text{ for all } i \in \mathcal{M}. \quad (12) \end{aligned}$$

Since the objective function and the constraints in (12) are convex, the optimization problem (12) is a convex optimization problem [33]. In fact, (12) can be reformulated as a linear program [33, Sec. 1.2.2].

The solution to (12) might not be unique.¹ Any such solution $\hat{\mathbf{x}}$ is characterized by two properties: (i) it is consistent with the initial labels, i.e., $\hat{x}_i = x_i$ for all nodes $i \in \mathcal{M}$ in the training set; and (ii) it has minimum TV among all such graph signals.

We solve (12) using a recently proposed primal-dual method [38]. This approach is appealing since it comes with a theoretical convergence guarantee and can be implemented efficiently as message passing over the underlying empirical graph (see Algorithm 2 below). The resulting algorithm bears some similarity to the class of label propagation (LP) algorithms for SSL on graphs [2], [39]. Indeed, LP algorithms can be interpreted as message passing methods for solving the optimization problem [5, Chap 11.3.4.]:

$$\begin{aligned} \hat{\mathbf{x}}^{(\text{LP})} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \sum_{\{i,j\} \in \mathcal{E}} W_{i,j}^2 (\tilde{x}_i - \tilde{x}_j)^2 \\ \text{ s.t. } \tilde{x}_i = x_i \text{ for all } i \in \mathcal{M}. \quad (13) \end{aligned}$$

The learning problem (13) amounts to minimizing the weighted sum of squared signal differences $(\tilde{x}_i - \tilde{x}_j)^2$ over edges $\{i, j\} \in \mathcal{E}$ in the empirical graph. In contrast, TV minimization (12) aims to minimize a weighted sum of absolute values of the signal differences $|\tilde{x}_i - \tilde{x}_j|$. It turns out that using the absolute values of the signal differences (the TV) instead of the sum of squared differences (as in LP) results in piece-wise constant graph signals (see (9)). In contrast, LP methods smooth out abrupt signal variations (see Section VI), making them unsuitable for data which can be (approximately) represented by piece-wise constant graph signals. LP methods have been shown to fail dramatically for random geometric graphs [11].

TV minimization (12) and LP (13) are special cases of p -Laplacian minimization [12]

$$\begin{aligned} \hat{\mathbf{x}}^{(p)} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \sum_{\{i,j\} \in \mathcal{E}} (W_{i,j} |\tilde{x}_i - \tilde{x}_j|)^p \\ \text{ s.t. } \tilde{x}_i = x_i \text{ for all } i \in \mathcal{M}. \quad (14) \end{aligned}$$

Indeed, TV minimization (12) is obtained from (14) when $p = 1$, while the LP problem (13) is obtained when $p = 2$. The limiting case of (14) for $p \rightarrow \infty$, known as the *minimal Lipschitz extension problem*, is studied in [22]. The work [22] presents efficient solvers and proves stability of the solutions

¹Assume that no initial labels are available such that the training set \mathcal{M} would be empty. Then, every constant graph signal solves (12).

for (14) in this limiting case. However, while the algorithms in [22] have high (combinatorial) complexity, we can solve TV minimization using efficient convex optimization methods (see Section IV).

The TV minimization problem (12) is also closely related to graph trend filtering [15] and the more general network Lasso (nLasso) [21], [40]

$$\hat{\mathbf{x}}^{(\text{nL})} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \sum_{i \in \mathcal{M}} (\tilde{x}_i - x_i)^2 + \lambda \|\tilde{\mathbf{x}}\|_{\text{TV}}. \quad (15)$$

By Lagrangian duality [33], [41], there are values (which might depend on the initial labels x_i) for λ in (15) such that solutions of (15) coincide with those of (12). The tuning parameter $\lambda > 0$ in (15) allows us to trade a small empirical error $\sum_{i \in \mathcal{M}} (\hat{x}_i^{(\text{nL})} - x_i)^2$ against a small TV $\|\hat{\mathbf{x}}^{(\text{nL})}\|_{\text{TV}}$ of the learned graph signal $\hat{\mathbf{x}}^{(\text{nL})}$. Choosing a large value of λ enforces a small TV of the learned graph signal. Using a small value for λ puts more emphasis on the empirical error. In contrast to nLasso (15), TV minimization (12) does not require any parameter tuning.

III. THE DUAL OF TV MINIMIZATION

TV minimization (12) involves non-differentiable objective function, which rules out gradient (descent) methods. However, both the objective function and the constraint set of (12) have a simple structure individually. The compositional structure of (12) can be exploited by studying an equivalent dual problem. It turns out that this dual problem has an interpretation as network (flow) optimization [32]. Moreover, by jointly considered the primal TV minimization (12) and its dual we obtain an efficient method for simultaneously solving TV minimization (12) and its dual (see Section IV).

In order to formulate the dual problem we first reformulate TV minimization (12) as an equivalent unconstrained convex optimization problem

$$\hat{\mathbf{x}} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} f(\tilde{\mathbf{x}}) := g(\mathbf{D}\tilde{\mathbf{x}}) + h(\tilde{\mathbf{x}}), \quad (16)$$

with

$$g(\mathbf{y}) := \|\mathbf{y}\|_1, \text{ and } h(\tilde{\mathbf{x}}) := \begin{cases} \infty & \text{if } \tilde{\mathbf{x}} \notin \mathcal{Q} \\ 0 & \text{if } \tilde{\mathbf{x}} \in \mathcal{Q}. \end{cases}$$

The constraint set $\mathcal{Q} = \{\tilde{\mathbf{x}} \in \mathbb{R}^N : \tilde{x}_i = x_i \text{ for all } i \in \mathcal{M}\}$ collects all graph signals which match the labels x_i on the training set \mathcal{M} . The (extended-value) function $h(\mathbf{x})$ in (16) is the indicator function of the convex set \mathcal{Q} (see [33]).

We can view (16) as the primal problem (or formulation) of TV minimization (12). The dual problem associated with TV minimization is

$$\hat{\mathbf{y}} \in \arg \max_{\mathbf{y} \in \mathbb{R}^E} \tilde{f}(\mathbf{y}) := -h^*(-\mathbf{D}^T \mathbf{y}) - g^*(\mathbf{y}). \quad (17)$$

The objective function $\tilde{f}(\mathbf{y})$ of the dual problem (17) is composed of the convex conjugates (see (1)) of the components $h(\mathbf{x})$ and $g(\mathbf{y})$ of the primal problem (16). These convex conjugates

are given explicitly by

$$h^*(\tilde{\mathbf{x}}) = \sup_{\mathbf{z} \in \mathbb{R}^N} \mathbf{z}^T \tilde{\mathbf{x}} - h(\mathbf{z}) \stackrel{(16)}{=} \begin{cases} \infty & \text{if } \tilde{x}_i \neq 0 \text{ for some } i \in \mathcal{V} \setminus \mathcal{M} \\ \sum_{i \in \mathcal{M}} \tilde{x}_i x_i & \text{otherwise,} \end{cases} \quad (18)$$

and

$$g^*(\mathbf{y}) = \sup_{\mathbf{z} \in \mathbb{R}^E} \mathbf{z}^T \mathbf{y} - g(\mathbf{z}) \stackrel{(16)}{=} \sup_{\mathbf{z} \in \mathbb{R}^E} \mathbf{z}^T \mathbf{y} - \|\mathbf{z}\|_1 = \begin{cases} \infty & \text{if } \|\mathbf{y}\|_\infty > 1 \\ 0 & \text{otherwise.} \end{cases} \quad (19)$$

The relation between the primal problem (16) and the dual problem (17) is made precise in [42, Thm. 31.3]. First, the optimal values of (16) and (17) coincide:

$$\min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} g(\mathbf{D}\tilde{\mathbf{x}}) + h(\tilde{\mathbf{x}}) = \max_{\mathbf{y} \in \mathbb{R}^E} -h^*(-\mathbf{D}^T \mathbf{y}) - g^*(\mathbf{y}). \quad (20)$$

The identity (20) is useful for bounding the sub-optimality $\|\tilde{\mathbf{x}}\|_{\text{TV}} - \|\hat{\mathbf{x}}\|_{\text{TV}}$ of a given candidate $\tilde{\mathbf{x}}$ for the solution $\hat{\mathbf{x}}$ to the TV minimization (12). According to (20), given any (dual) vector $\mathbf{y} \in \mathbb{R}^E$, we can bound the sub-optimality as

$$\|\tilde{\mathbf{x}}\|_{\text{TV}} - \|\hat{\mathbf{x}}\|_{\text{TV}} \leq \|\tilde{\mathbf{x}}\|_{\text{TV}} + (h^*(-\mathbf{D}^T \mathbf{y}) + g^*(\mathbf{y})). \quad (21)$$

Another consequence of the duality result [42, Thm. 31.3] is a powerful characterization of the solutions of the primal (16) and dual problem (17). In particular, a pair of vectors $\hat{\mathbf{x}} \in \mathbb{R}^N$, $\hat{\mathbf{y}} \in \mathbb{R}^E$ are solutions to the primal (16) and dual problem (17), respectively, if and only if

$$-(\mathbf{D}^T \hat{\mathbf{y}}) \in \partial h(\hat{\mathbf{x}}), \mathbf{D}\hat{\mathbf{x}} \in \partial g^*(\hat{\mathbf{y}}). \quad (22)$$

Given any solution $\hat{\mathbf{y}} \in \mathbb{R}^E$ to the dual problem (17), any solution $\hat{\mathbf{x}}$ to the primal problem (16) and, in turn, to TV minimization (12) must be such that conditions (22) are satisfied. The optimality condition (22) is the launching point for a primal-dual method for solving (12) (see Section IV).

It turns out that the dual (17) of TV minimization (12) is an instance of network optimization for the empirical graph \mathcal{G} . To show this, we need the following definition.

Definition 1: A network flow $f : \mathcal{E} \rightarrow \mathbb{R}$ with supplies v_i , at the nodes $i \in \mathcal{V}$, assigns each directed edge $e = (i, j) \in \mathcal{E}$ some value $f_e \in \mathbb{R}$. The flow has to satisfy the conservation law:

$$\sum_{j \in \mathcal{N}^+(i)} f_{(i,j)} - \sum_{j \in \mathcal{N}^-(i)} f_{(j,i)} = v_i \text{ for each } i \in \mathcal{V}. \quad (23)$$

For a given empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, we will consider flows that satisfy the capacity constraints:

$$|f_e| \leq W_e \quad (24)$$

for some edges $e \in \mathcal{E}$. Thus, we interpret the weights W_e of the empirical graph as capacities of a flow network. At a later point, we will make explicit those edges for which the capacity constraints (24) have to be satisfied.

We can associate any dual vector $\mathbf{y} \in \mathbb{R}^E$ with a particular flow $f^{(\mathbf{y})}$ whose values are given by $f_e^{(\mathbf{y})} := W_e y_e$. It is then

easy to verify that the flow $f^{(y)}$ satisfies the capacity constraints (24) and the conservation law (23) with supplies v_i if and only if

$$\|\mathbf{y}\|_\infty \leq 1, \mathbf{D}^T \mathbf{y} = \mathbf{v} \text{ with } \mathbf{v} = (v_1, \dots, v_N)^T \in \mathbb{R}^N. \quad (25)$$

Thus, the magnitude $|y_e|$ of a dual vector entry represents the fraction of the edge capacity W_e flowing through edge $e \in \mathcal{E}$.

Proposition 1: The dual problem (17) of TV minimization (12) is equivalent to the network optimization problem

$$\max_{f \in \mathcal{R}} \sum_{i \in \mathcal{M}} x_i \sum_{j \in \mathcal{N}(i)} f_{(i,j)}, \quad (26)$$

with the constraint set \mathcal{R} consisting of all flows that conform with (24) and (23) with supplies v_i satisfying

$$v_i = 0 \text{ for all unlabeled nodes } i \in \mathcal{V} \setminus \mathcal{M}. \quad (27)$$

In particular, \mathbf{y} solves (17) if and only if the flow $f^{(y)}$, defined edge-wise by $f_e^{(y)} = W_e y_e$, solves (26).

Proof: The (extended-value) functions (18) and (19), which constitute the dual problem (17), implicitly constrain the dual vector \mathbf{y} to satisfy (25) with supplies of the form (27). Thus, any optimal dual vector $\hat{\mathbf{y}}$ induces a flow $f^{(\hat{\mathbf{y}})} \in \mathcal{R}$. For any $\mathbf{y} \in \mathbb{R}^E$ such that the flow $f^{(y)}$ belongs to \mathcal{R} , the objective functions in (26) and (17) coincide. ■

The problem (26) is an instance of a minimum-cost flow problem discussed in [32, Ch. 1]. Various methods for solving minimum-cost flow problems are presented in [32].

Combining Proposition 1 with the primal-dual optimality condition (22) provides a characterization of the solutions of TV minimization in terms of particular network flows.

Corollary 2: Given networked data with empirical graph \mathcal{G} and labels $\{x_i\}_{i \in \mathcal{M}}$, consider some flow \hat{f} which solves the minimum-cost flow problem (26). Let us denote the set of edges which are not saturated in \hat{f} by

$$\mathcal{U} := \{\{i, j\} \in \mathcal{E} : |\hat{f}_e| < W_e\}.$$

Then, any solution $\hat{\mathbf{x}}$ of (12) satisfies $\hat{x}_i = \hat{x}_j$ for each $e = \{i, j\} \in \mathcal{U}$. Thus, given some optimal flow \hat{f} (which solves (26)), any solution to TV minimization is constant along edges which are not saturated by \hat{f} .

Proof: For the optimal flow \hat{f} define the dual vector $\hat{\mathbf{y}} = \hat{f}_e / W_e$. According to Proposition 1, $\hat{\mathbf{y}}$ is a solution to the dual problem (17). For this particular (optimal) dual vector $\hat{\mathbf{y}}$, any solution $\hat{\mathbf{x}}$ to TV minimization has to satisfy the optimality condition (22). Using the right-hand condition in (22) and the properties of the sub-differential $\partial g^*(\mathbf{y})$ (see (19) and [42, Sec. 32]) yields the statement. ■

Note that, for a particular edge $e = \{i, j\} \in \mathcal{E}$ in the empirical graph, once we find at least one optimal flow \hat{f} such that $|\hat{f}_e| < W_e$ we are assured that every solution to TV minimization is constant along that edge e . However, to apply Corollary 2 we need an efficient means to construct or characterize flows which are optimal in the sense of (26). While there exist some well-known methods for solving minimum-cost flow problems (see [32]), we consider Corollary 2 mainly useful for (partially) characterizing the solutions of TV minimization. In order to

actually solve TV minimization we will apply a different method which starts directly from the optimality conditions (22).

IV. A PRIMAL-DUAL METHOD

The solutions $\hat{\mathbf{x}}$ of (16) are characterized by [42]

$$\mathbf{0} \in \partial f(\hat{\mathbf{x}}). \quad (28)$$

Proximal methods solve (16) via fixed-point iterations of an operator \mathcal{P} whose fixed-points are the solutions $\hat{\mathbf{x}}$ of (28),

$$\mathbf{0} \in \partial f(\hat{\mathbf{x}}) \text{ if and only if } \hat{\mathbf{x}} = \mathcal{P}\hat{\mathbf{x}}. \quad (29)$$

In general, the operator \mathcal{P} is not unique, i.e., there are different choices for \mathcal{P} such that (29) is valid. These choices result in different proximal algorithms [23]. One useful choice for \mathcal{P} in (29) is suggested by the characterization (22) of solutions to the primal (16) and dual (17) form of TV minimization (12). The resulting method has been presented in [19, Algorithm 1].

Let us detail the derivation of [19, Algorithm 1] which is restated as Algorithm 1 below. Rewrite the two coupled conditions (22) as

$$\begin{aligned} \hat{\mathbf{x}} - \Gamma \mathbf{D}^T \hat{\mathbf{y}} &\in \hat{\mathbf{x}} + \Gamma \partial h(\hat{\mathbf{x}}) \\ 2\Lambda \mathbf{D} \hat{\mathbf{x}} + \hat{\mathbf{y}} &\in \Lambda \partial g^*(\hat{\mathbf{y}}) + \Lambda \mathbf{D} \hat{\mathbf{x}} + \hat{\mathbf{y}}, \end{aligned} \quad (30)$$

with the invertible diagonal matrices (cf. (2) and (3))

$$\begin{aligned} \Lambda &:= (1/2) \text{diag}\{\lambda_{\{i,j\}} = 1/W_{i,j}\}_{\{i,j\} \in \mathcal{E}} \in \mathbb{R}^{E \times E} \text{ and} \\ \Gamma &:= (1/2) \text{diag}\{\gamma_i = 1/d_i\}_{i=1}^N \in \mathbb{R}^{N \times N}. \end{aligned} \quad (31)$$

The particular choice (31) ensures that [38, Lemma 2]

$$\|\Gamma^{1/2} \mathbf{D}^T \Lambda^{1/2}\|_2 < 1,$$

which, in turn, guarantees convergence of the iterative algorithm we propose for solving (16).

Using the concept of resolvent operators [38, Sec. 1.1.], we further develop the characterization (30) of solutions $\hat{\mathbf{x}}$ to TV minimization (12). To this end we define the resolvent operators for the (set-valued) operators $\Lambda \partial g^*(\mathbf{y})$ and $\Gamma \partial h(\mathbf{x})$ (see (16)) as

$$\begin{aligned} (\mathbf{I} + \Lambda \partial g^*)^{-1}(\mathbf{y}) &:= \arg \min_{\mathbf{z} \in \mathbb{R}^E} g^*(\mathbf{z}) + (1/2) \|\mathbf{y} - \mathbf{z}\|_{\Lambda}^2 \\ (\mathbf{I} + \Gamma \partial h)^{-1}(\mathbf{x}) &:= \arg \min_{\mathbf{z} \in \mathbb{R}^N} h(\mathbf{z}) + (1/2) \|\mathbf{x} - \mathbf{z}\|_{\Gamma}^2. \end{aligned} \quad (32)$$

Applying [43, Prop. 23.2] and [43, Prop. 16.44] to the optimality condition (30) yields the equivalent condition (for $\hat{\mathbf{x}}, \hat{\mathbf{y}}$ to be primal and dual optimal)

$$\begin{aligned} \hat{\mathbf{x}} &= (\mathbf{I} + \Gamma \partial h)^{-1}(\hat{\mathbf{x}} - \Gamma \mathbf{D}^T \hat{\mathbf{y}}) \\ \hat{\mathbf{y}} - 2(\mathbf{I} + \Lambda \partial g^*)^{-1} \Lambda \mathbf{D} \hat{\mathbf{x}} &= (\mathbf{I} + \Lambda \partial g^*)^{-1}(\hat{\mathbf{y}} - \Lambda \mathbf{D} \hat{\mathbf{x}}). \end{aligned} \quad (33)$$

The characterization (33) of the solution $\hat{\mathbf{x}} \in \mathbb{R}^N$ for the TV minimization problem (12) leads naturally to the following coupled fixed-point iterations for finding a solution $\hat{\mathbf{x}}$ of (12):

$$\begin{aligned} \hat{\mathbf{y}}^{(k+1)} &:= (\mathbf{I} + \Lambda \partial g^*)^{-1}(\hat{\mathbf{y}}^{(k)} + \Lambda \mathbf{D}(2\hat{\mathbf{x}}^{(k)} - \hat{\mathbf{x}}^{(k-1)})) \\ \hat{\mathbf{x}}^{(k+1)} &:= (\mathbf{I} + \Gamma \partial h)^{-1}(\hat{\mathbf{x}}^{(k)} - \Gamma \mathbf{D}^T \hat{\mathbf{y}}^{(k+1)}). \end{aligned} \quad (34)$$

Algorithm 1: Primal-Dual Method for TV Minimization.

Input: empirical graph \mathcal{G} with incidence matrix $\mathbf{D} \in \mathbb{R}^{E \times N}$ (see (5)), training set \mathcal{M} with labels $\{x_i\}_{i \in \mathcal{M}}$.

Initialize: $k := 0$, $\bar{\mathbf{x}} = \hat{\mathbf{x}}^{(-1)} = \hat{\mathbf{x}}^{(0)} = \hat{\mathbf{y}}^{(0)} := \mathbf{0}$, $\gamma_i := 1/d_i$, $\lambda_{\{i,j\}} = 1/(2W_{i,j})$.

- 1: **repeat**
- 2: $\tilde{\mathbf{x}} := 2\hat{\mathbf{x}}^{(k)} - \hat{\mathbf{x}}^{(k-1)}$
- 3: $\hat{\mathbf{y}}^{(k+1)} := \hat{\mathbf{y}}^{(k)} + \mathbf{A}\mathbf{D}\tilde{\mathbf{x}}$ with $\mathbf{A} = \text{diag}\{\lambda_{\{i,j\}}\}_{\{i,j\} \in \mathcal{E}}$
- 4: $\hat{y}_e^{(k+1)} := \hat{y}_e^{(k+1)} / \max\{1, |\hat{y}_e^{(k+1)}|\}$ for every edge $e \in \mathcal{E}$
- 5: $\hat{\mathbf{x}}^{(k+1)} := \hat{\mathbf{x}}^{(k)} - \mathbf{\Gamma}\mathbf{D}^T\hat{\mathbf{y}}^{(k+1)}$ with $\mathbf{\Gamma} = \text{diag}\{\gamma_i\}_{i \in \mathcal{V}}$
- 6: $\hat{x}_i^{(k+1)} := x_i$ for every labeled node $i \in \mathcal{M}$
- 7: $k := k + 1$
- 8: $\bar{\mathbf{x}}^{(k)} := (1 - 1/k)\bar{\mathbf{x}}^{(k-1)} + (1/k)\hat{\mathbf{x}}^{(k)}$
- 9: **until** stopping criterion is satisfied

Output: labels $\hat{x}_i := \bar{x}_i^{(k)}$ for all nodes $i \in \mathcal{V}$

Here, we used the diagonal matrices defined in (31) as well as the incidence matrix \mathbf{D} (see (5)). The fixed-point iterations (34) are obtained as a special case of the iterations [38, Eq. (4)] when choosing $\theta = 1$ (using the notation in [38]).

We implement the updates in (34) by using simple closed-form expressions for the resolvent operators (32) (see [14, Sec. 6.2.] for more details):

$$(\mathbf{I} + \mathbf{A}\partial g^*)^{-1}(\mathbf{y}) = (\tilde{y}_1, \dots, \tilde{y}_N)^T, \tilde{y}_i = y_i / \max\{|y_i|, 1\}$$

$$(\mathbf{I} + \mathbf{\Gamma}\partial h)^{-1}(\tilde{\mathbf{x}}) = (t_1, \dots, t_N)^T, t_i = \begin{cases} x_i & \text{for } i \in \mathcal{M} \\ \tilde{x}_i & \text{otherwise.} \end{cases} \quad (35)$$

Inserting (35) into the updates (34) yields Algorithm 1 for solving TV minimization (12). Note that Algorithm 1 is a special case of [14, Algorithm 1] which uses a more general version of step 2 in Algorithm 1 of the form $\tilde{\mathbf{x}} := \hat{\mathbf{x}}^{(k)} + \theta(\hat{\mathbf{x}}^{(k)} - \hat{\mathbf{x}}^{(k-1)})$. Thus, step 2 in Algorithm 1 is obtained for the particular choice $\theta = 1$. This choice ensures convergence of Algorithm 1 with an optimal (worst-case) converge rate (see [19]). The tuning of θ is beyond the scope of this paper. Another difference between Algorithm 1 and [14, Algorithm 1] is the explicit computation of the running average in step 8 (which is required for the convergence analysis underlying Proposition 2).

We emphasize that Algorithm 1 does not require knowledge of the partition \mathcal{F} underlying signal model (9). It also does not involve any tuning parameters.

There are various possible stopping criteria in Algorithm 1, including using a fixed number of iterations or testing for sufficient decrease of the objective function (see [44] and Section VI). For testing if the objective function is decreased sufficiently, we can use the duality bound (21) on the sub-optimality of the current objective function value $\|\bar{\mathbf{x}}^{(k)}\|_{\text{TV}}$. When using a fixed number of iterations, the following characterization of the convergence rate of Algorithm 1 is helpful.

Proposition 2 ([19]): Consider the sequences $\hat{\mathbf{x}}^{(k)}$ and $\hat{\mathbf{y}}^{(k)}$ obtained from the update rule (34) and starting from some initializations $\hat{\mathbf{x}}^{(0)}$ and $\hat{\mathbf{y}}^{(0)}$. The averages

$$\bar{\mathbf{x}}^{(K)} = (1/K) \sum_{k=1}^K \hat{\mathbf{x}}^{(k)}, \text{ and } \bar{\mathbf{y}}^{(K)} = (1/K) \sum_{k=1}^K \hat{\mathbf{y}}^{(k)} \quad (36)$$

obtained after K iterations of (34), satisfy

$$\|\bar{\mathbf{x}}^{(K)}\|_{\text{TV}} - \|\hat{\mathbf{x}}\|_{\text{TV}} \leq (1/(2K))(\|\hat{\mathbf{x}}^{(0)} - \hat{\mathbf{x}}\|_{\mathbf{\Gamma}^{-1}}^2 + \|\hat{\mathbf{y}}^{(0)} - \bar{\mathbf{y}}^{(K)}\|_{\mathbf{A}^{-1}}^2) \quad (37)$$

with $\bar{\mathbf{y}}^{(K)} = \text{sign}\{\mathbf{D}\bar{\mathbf{x}}^{(K)}\}$. Moreover, the sequence $\|\hat{\mathbf{y}}^{(0)} - \bar{\mathbf{y}}^{(K)}\|_{\mathbf{A}^{-1}}$, for $K = 1, 2, \dots$ is bounded.

According to (37), the sub-optimality of Algorithm 1 after K iterations is bounded as

$$\|\bar{\mathbf{x}}^{(K)}\|_{\text{TV}} - \|\hat{\mathbf{x}}\|_{\text{TV}} \leq c/K, \quad (38)$$

where the constant c does not depend on K but might depend on the empirical graph \mathcal{G} , via its weighted incidence matrix \mathbf{D} (5), as well as on the initial labels $\{x_i\}_{i \in \mathcal{M}}$. The bound (38) suggests that in order to ensure reducing the sub-optimality by a factor of two, we need to run Algorithm 1 for twice as many iterations. The upper bound (38) is tight among all message passing (local) methods for solving (12). In particular, the rate $1/K$ cannot be improved for a chain-structured empirical graph (see [19]).

As indicated by [45, Thm. 3.2], Algorithm 1 is robust to numerical errors arising during the updates, which can be a crucial property for high-dimensional problems.

The computational cost of one iteration in Algorithm 1 is proportional to the number of edges in the empirical graph \mathcal{G} . This can be verified by noting that Algorithm 1 can be implemented as message passing on the empirical graph (see Algorithm 2). Thus, for a fixed number K of iterations, the computational cost of Algorithm 1 is proportional to the number of edges in the empirical graph. In contrast, the computational cost of state-of-the-art maximum flow algorithms can be considerably higher [46], [47]. Moreover, while Algorithm 1 allows for a rather straightforward implementation on modern big data computing frameworks (see Section VI-C), this is typically more challenging for maximum flow methods which are (partially) based on combinatorial search (see [38, Sec. 3.3.]).

We now show how to obtain a scalable implementation of Algorithm 1 using message passing over the underlying empirical graph \mathcal{G} . This message passing formulation, summarized in Algorithm 2 (being a slight reformulation of [19, Algorithm 2]), is obtained by implementing the application of the graph incidence matrix \mathbf{D} and its transpose \mathbf{D}^T (cf. steps 2 and 5 of Algorithm 1) by local updates of the labels \hat{x}_i , i.e., updates which involve only the neighbourhoods $\mathcal{N}(i)$, $\mathcal{N}(j)$ of all edges $\{i, j\} \in \mathcal{E}$ in the empirical graph \mathcal{G} .

Note that executing Algorithm 2 does not require global knowledge (such as the maximum node degree d_{\max} (4)) about the entire empirical graph. Indeed, if we associate each node in the data graph with a computational unit, execution of Algorithm 2 requires each node $i \in \mathcal{V}$ only to store the neighboring values $\{\hat{y}_{\{i,j\}}, W_{i,j}\}_{j \in \mathcal{N}(i)}$ and $\hat{x}_i^{(k)}$. Moreover, the number

Algorithm 2: Distributed Implementation of Algorithm 1.

Input: empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, training set \mathcal{M} with labels $\{x_i\}_{i \in \mathcal{M}}$.

Initialize: $k := 0$, $\bar{\mathbf{x}} = \hat{\mathbf{y}}^{(0)} = \hat{\mathbf{x}}^{(-1)} = \hat{\mathbf{x}}^{(0)} := \mathbf{0}$,

$\gamma_i := 1/d_i$.

1: **repeat**

2: for all nodes $i \in \mathcal{V}$: $\tilde{x}_i := 2\hat{x}_i^{(k)} - \hat{x}_i^{(k-1)}$

3: for all edges $e = (i, j) \in \mathcal{E}$:

$$\hat{y}_e^{(k+1)} := \hat{y}_e^{(k)} + (1/2)(\tilde{x}_{e^+} - \tilde{x}_{e^-})$$

4: for all edges $e \in \mathcal{E}$:

$$\hat{y}_e^{(k+1)} := \hat{y}_e^{(k+1)} / \max\{1, |\hat{y}_e^{(k+1)}|\}$$

5: for all nodes $i \in \mathcal{V}$:

$$\hat{x}_i^{(k+1)} := \hat{x}_i^{(k)} - \gamma_i \left[\sum_{j \in \mathcal{N}^+(i)} W_{i,j} \hat{y}_{(i,j)}^{(k+1)} - \sum_{j \in \mathcal{N}^-(i)} W_{i,j} \hat{y}_{(j,i)}^{(k+1)} \right]$$

6: for all labeled nodes $i \in \mathcal{M}$: $\hat{x}_i^{(k+1)} := x_i$

7: $k := k + 1$

8: for all nodes $i \in \mathcal{V}$: $\bar{x}_i := (1 - 1/k)\bar{x}_i + (1/k)\hat{x}_i^{(k)}$

9: **until** stopping criterion is satisfied

Output: labels $\hat{x}_i := \hat{x}_i^{(k)}$ for all $i \in \mathcal{V}$

of arithmetic operations required at each node $i \in \mathcal{V}$ during each time step is proportional to the number $|\mathcal{N}(i)|$ of its neighbours $\mathcal{N}(i)$. Thus, Algorithm 2 can be scaled to large datasets which can be represented as sparse networks having small maximum degree d_{\max} (4). The datasets generated in many important applications are accurately represented by such sparse networks [48].

Algorithm 1 also implicitly solves the dual problem (17) of TV minimization (12). We might therefore interpret Algorithm 2 as a message passing method for network optimization. In particular, associate the current approximation $\hat{\mathbf{y}}^{(k)}$ for the optimal dual vector $\hat{\mathbf{y}}$ (see (17)) with the flow $f^{(k)} : \mathcal{E} \rightarrow \mathbb{R}$ having values $f_e^{(k)} := W_e y_e^{(k)}$. Then, step 4 of Algorithm 2 aims at enforcing the capacity constraint (24) for the flow $f^{(k)}$. Moreover, step 5 amounts to updating the current signal estimate $\hat{x}_i^{(k)}$, for each unlabeled node $i \in \mathcal{V} \setminus \mathcal{M}$, by the (scaled) demand induced by the current flow $f^{(k)}$ (23). Thus, for each unlabeled node $i \in \mathcal{V} \setminus \mathcal{M}$, we might interpret the signal estimates $\hat{x}_i^{(k)}$ as the (scaled) cumulative demand induced by the flows $f^{(k')}$ for $k' = 1, \dots, k$. The labeled nodes $i \in \mathcal{M}$ have a constant supply $\hat{x}_i^{(k)} = x_i$ whose amount is the label x_i . Step 3 of Algorithm 2 balances discrepancies between accumulated demands $\hat{x}_i^{(k)}$ at the different nodes by adapting the flow $f_{(i,j)}^{(k)}$ through an edge $e = (i, j) \in \mathcal{E}$ according to the difference $(\tilde{x}_i - \tilde{x}_j)$.

V. WHEN IS TV MINIMIZATION ACCURATE?

We now provide conditions which ensure that any solution $\hat{\mathbf{x}}$ of TV minimization (12) is close to the true underlying graph signal $\mathbf{x} = (x_1, \dots, x_N)^T \in \mathbb{R}^N$ which can be well approximated by a piece-wise constant graph signal (9).

Since TV minimization (12) is a particular case of ℓ_1 minimization [27], successful recovery is ensured by the stable analysis nullspace property (see [18, Lemma 5]).

As we show in Proposition 4, the stable analysis nullspace property is ensured if the nodes in the training set are sufficiently well connected to the cluster boundaries $\partial\mathcal{F}$. To this end, we define the notion of resolving training sets.

Definition 3: Consider a partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{|\mathcal{F}|}\}$ of the empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$ into disjoint subsets of nodes (clusters) $\mathcal{C}_l \subseteq \mathcal{V}$. A training set $\mathcal{M} \subseteq \mathcal{V}$ resolves the partition \mathcal{F} if, for any collection of signs $\{b_e \in \{-1, 1\}\}_{e \in \partial\mathcal{F}}$, there exists a flow $f : \mathcal{E} \rightarrow \mathbb{R}$ such that

$$\begin{aligned} f_{(i,j)} &= b_{(i,j)} 2W_{i,j} \text{ for each } (i,j) \in \partial\mathcal{F} \\ |f_{(i,j)}| &\leq W_{i,j} \text{ for each } (i,j) \in \mathcal{E} \setminus \partial\mathcal{F} \\ \sum_{(i,j) \in \mathcal{E}} f_{(i,j)} - \sum_{(j,i) \in \mathcal{E}} f_{(j,i)} &= 0 \text{ for each } i \in \mathcal{V} \setminus \mathcal{M}. \end{aligned} \quad (39)$$

We highlight that Definition 3 is only required for the analysis of the solutions of TV minimization (12). In order to use Algorithm 1 for solving (12), we do not need to place any requirements on the training set \mathcal{M} . We can perfectly use Algorithm 1 also when the training set \mathcal{M} does not resolve the partition \mathcal{F} underlying the signal model (9). However, in this case we cannot guarantee that the estimate delivered by Algorithm 1 is close to the true underlying graph signal.

It is important to note that Definition 3 involves both the labeled training set \mathcal{M} and the partition \mathcal{F} . For a given training set \mathcal{M} , we can increase the chance of satisfying (39) by optimizing the partition \mathcal{F} underlying (9). Enlarging the training set \mathcal{M} (by acquiring more labels), will increase the chance of satisfying (39) as there are fewer unlabeled nodes for which the last condition in (39) has to be ensured.

Definition 3 requires a sufficiently large network flow (across cluster boundaries) between the labeled nodes \mathcal{M} . These network flows have to be such that the boundary edges $e \in \partial\mathcal{F}$ are flooded (or saturated) with an amount of flow at least $2W_e$. The training set $\mathcal{M} \subseteq \mathcal{V}$ depicted in Fig. 1 resolves the partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2\}$.

Proposition 3 (Thm. 4 in [18]): Consider data with empirical graph \mathcal{G} and true labels x_i forming a graph signal $\mathbf{x} \in \mathbb{R}^N$. We are provided with observed labels x_i at nodes in the training set \mathcal{M} . If \mathcal{M} resolves the partition $\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$, any solution $\hat{\mathbf{x}}$ of (12) satisfies

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_{\text{TV}} \leq 6 \min_{\{a_l\}_{l=1}^{|\mathcal{F}|}} \left\| \mathbf{x} - \sum_{l=1}^{|\mathcal{F}|} a_l \mathcal{I}_{\mathcal{C}_l}[\cdot] \right\|_{\text{TV}}, \quad (40)$$

For convenience, we spell out a bound on the error $\hat{x}_i - x_i$ itself which is a direct consequence of (40).

Corollary 4: Under the same assumptions as in Proposition 3, any solution of (12) satisfies

$$\max_{i \in \mathcal{V}} |\hat{x}_i - x_i| \leq 6d_{\max} \min_{\{a_l\}_{l=1}^{|\mathcal{F}|}} \left\| \mathbf{x} - \sum_{l=1}^{|\mathcal{F}|} a_l \mathcal{I}_{\mathcal{C}_l}[\cdot] \right\|_1. \quad (41)$$

Proof: The bound (41) is obtained from (40) using the inequality $\|\mathbf{z}\|_{\text{TV}} \leq d_{\max} \|\mathbf{z}\|_1$ (see (7)) with the maximum weighted degree d_{\max} (4). ■

Thus, if the training set \mathcal{M} resolves the partition underlying (9), any solution $\hat{\mathbf{x}}$ to TV minimization (12) is close (in TV seminorm) to the true labels if they can be well approximated by a piece-wise constant graph signal (9). For labels forming exactly a piece-wise constant signal, we can specialize Proposition 3 as follows.

Corollary 5 (Thm. 3 in [18]): Consider data with empirical graph \mathcal{G} and true labels x_i forming a piece-wise constant graph signal $\mathbf{x} \in \mathbb{R}^N$ (see (9)) over the partition $\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$. If the training set \mathcal{M} resolves \mathcal{F} , the solution $\hat{\mathbf{x}}$ of (12) is unique and coincides with \mathbf{x} .

We emphasize that Algorithm 1 does not require knowledge of the partition $\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$. Indeed, we could use Algorithm 1 to determine the clusters \mathcal{C}_l if the underlying labels x_i form a piece-wise constant signal $x_i = \sum_{l=1}^{|\mathcal{F}|} a_l \mathcal{I}_{\mathcal{C}_l}[i]$ with $a_l \neq a_{l'}$ for different clusters $l \neq l'$.

Proposition 3 and Corollary 5 require the partition \mathcal{F} in (9) to be resolved by the training set \mathcal{M} . The direct verification that a given partition is resolved by a particular training set is computationally challenging as it involves an exponential number of constraints (39) to be evaluated. However, if the empirical graph is modeled using a probabilistic model, such as the stochastic block model (SBM) [37], we can make use of large deviation results to determine network parameter regimes such that (39) is satisfied with high probability [49].

We now show how to verify the validity of (39) using maximum flow algorithms [32], [50]. To this end, we define a particular subgraph \mathcal{G}_l associated with the clusters \mathcal{C}_l of a partition $\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$ which is resolved by \mathcal{M} .

Definition 6: For a given cluster $\mathcal{C}_l \subseteq \mathcal{V}$ within the empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, we define the augmented cluster subgraph $\mathcal{G}_l = (\mathcal{C}_l \cup \{0\}, \mathcal{E}_l, \mathbf{C}^{(l)})$ whose nodes are constituted by the cluster \mathcal{C}_l and the additional node 0. The edge set \mathcal{E}_l of \mathcal{G}_l is defined as

$$\mathcal{E}_l = \{\{i, j\} \in \mathcal{E} : i, j \in \mathcal{C}_l\} \cup \{\{0, i\} : i \in \partial\mathcal{C}_l \cap \mathcal{C}_l\}. \quad (42)$$

Thus, the edges \mathcal{E}_l of the augmented cluster subgraph \mathcal{G}_l are constituted by (i) the intra-cluster edges $\{\{i, j\} \in \mathcal{E} : i, j \in \mathcal{C}_l\}$ connecting nodes within cluster \mathcal{C}_l of the empirical graph \mathcal{G} and (ii) one additional edge $\{0, i\}$ for each node $i \in \partial\mathcal{C}_l \cap \mathcal{C}_l$ on the boundary of cluster \mathcal{C}_l . The weights $C_e^{(l)}$ of the edges $e \in \mathcal{E}_l$ in the graph \mathcal{G}_l are defined as

$$C_e^{(l)} := W_{i,j} \text{ for every edge } e = \{i, j\} \in \mathcal{E} \text{ with } i, j \in \mathcal{C}_l \quad (43)$$

and

$$C_{\{0,i\}}^{(l)} := 2 \sum_{j \in \mathcal{N}^{(i)} \setminus \mathcal{C}_l} W_{i,j} \text{ for each node } i \in \partial\mathcal{C}_l \cap \mathcal{C}_l. \quad (44)$$




Fig. 2. Augmented subgraphs $\mathcal{G}_1, \mathcal{G}_2$ obtained from the partitioned empirical graph in Fig. 1. Each subgraph is obtained from a cluster \mathcal{C}_l by adding edges from each boundary node $i^{(l)} \in \partial\mathcal{C}_l$ to the augmented node 0. The numbers indicate the capacity constraints (24) along the edges.

To illustrate Definition 6, Fig. 2 depicts the augmented subgraphs of the clusters in the empirical graph in Fig. 1.

Proposition 4: Consider an empirical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$ which is partitioned into the clusters $\mathcal{F} = \{\mathcal{C}_1, \dots, \mathcal{C}_{|\mathcal{F}|}\}$. Assume that each cluster \mathcal{C}_l contains at least one labeled node $i^{(l)} \in \mathcal{C}_l \cap \mathcal{M}$ from the training set $\mathcal{M} \subseteq \mathcal{V}$. If, for each cluster \mathcal{C}_l , the corresponding subgraph \mathcal{G}_l (see Definition 6) supports a network flow (using the capacities (43) and (44) for the capacity constraints (24)) of value $2 \sum_{e \in \partial\mathcal{C}_l} W_e$ between the source node $i^{(l)}$ and the sink node 0, then the training set \mathcal{M} resolves the partition \mathcal{F} .

Proof: Consider a particular cluster \mathcal{C}_l containing the labeled node $i^{(l)} \in \mathcal{C}_l \cap \mathcal{M}$. By assumption, the associated subgraph \mathcal{G}_l supports a network flow between $i^{(l)}$ and the extra node 0 of value $2 \sum_{e \in \partial\mathcal{C}_l} W_e$. The “max-flow min-cut” theorem (see [51, Thm. 6.1.6]) implies that this flow value can only be achieved if, for each subset $\mathcal{A} \subseteq \mathcal{C}_l \setminus \{i^{(l)}\}$, the total capacity of the edges $\{\{i, j\} \in \mathcal{E} : i \in \mathcal{A}, j \in \mathcal{C}_l \setminus \mathcal{A}\}$ is at least as large as twice the total capacity of the edges $\{\{i, j\} \in \mathcal{E} : i \in \mathcal{A}, j \in \mathcal{V} \setminus \mathcal{C}_l\}$,

$$\sum_{\{i,j\} \in \mathcal{E} : i \in \mathcal{A}, j \in \mathcal{C}_l \setminus \mathcal{A}} W_{i,j} \geq 2 \sum_{\{i,j\} \in \mathcal{E} : i \in \mathcal{A}, j \in \mathcal{V} \setminus \mathcal{C}_l} W_{i,j}. \quad (45)$$

The validity of (45), for each cluster \mathcal{C}_l of the partition \mathcal{F} , implies via Hoffman’s circulation theorem² [51, Thm. 10.2.7] the existence of a network flow satisfying the requirements (39) for the training set \mathcal{M} to resolve the partition \mathcal{F} . ■

In Section VI-A, we will demonstrate the usefulness of Proposition 4 for certifying the accuracy of Algorithm 1. Moreover, we can combine Proposition 4 with existing results from graph sampling to characterize TV minimization for empirical graphs that can be well approximated by an SBM. In particular, [49, Theorem 2.1] allows us to verify if the conditions of Proposition 4 are satisfied (with high probability) based on the expected values of cuts in the graph \mathcal{G} . According to Proposition 4, TV minimization is accurate if there exists a flow from the labeled nodes $\mathcal{M} \cap \mathcal{C}_l$ in each cluster to its boundary $\partial\mathcal{C}_l$ of value $2 \sum_{e \in \partial\mathcal{C}_l} W_e$. A simple argument based on [49, Theorem 2.1]

²See also [52] and [32, Ch. 3] for how Hoffman’s circulation theorem can be derived from the max-flow min-cut theorem.

shows that this condition is satisfied with high probability for an SBM (with cluster sizes not too small), whenever

$$|\mathcal{M} \cap \mathcal{C}_l| p_{\text{in}} \gg 2p_{\text{out}}(|\mathcal{V}| - |\mathcal{C}_l|). \quad (46)$$

Here, p_{in} (p_{out}) denotes the probability that two nodes from the same cluster (from different clusters) are connected by an edge. Condition (46) allows to characterize parameter regimes for the SBM such that TV minimization can recover piece-wise constant signals from a given number of labeled nodes. We will verify condition (46) empirically in Section VI-B.

Proposition 3 and Corollary 5 requires each cluster \mathcal{C}_l in (9) to contain at least one labeled node $i \in \mathcal{M}$ (see Definition 3). However, even if this condition is not met we still can say something about the solutions of TV minimization (12). In particular, the optimality condition (22) requires any solution $\hat{\mathbf{x}}$ of TV minimization (12) to be constant around labeled nodes $i \in \mathcal{M}$. The graph signal $\hat{\mathbf{x}}$ can only change along edges $e = \{i, j\} \in \mathcal{E}$ which are saturated, i.e., $|\hat{y}_e| = 1$ holds for every dual solution $\hat{\mathbf{y}}$ of (17) (see Corollary 2).

VI. NUMERICAL EXPERIMENTS

We assess the statistical and computational performance of Algorithm 1 using numerical experiments involving synthetic and “real-world” data. The first experiment discussed in Section VI-A revolves around an ensemble of synthetic datasets whose empirical graphs consist of two clusters with varying levels of connectivity. We verify the recovery condition provided by Proposition 3 by computing the recovery error of Algorithm 1 as the cluster connectivity varies. Section VI-B discusses the application of TV minimization to a synthetic empirical graph generated using an SBM. In Section VI-C, we verify the scalability of Algorithm 1 by implementing its message passing formulation Algorithm 2 in a big data framework. Finally, in Section VI-D, we discuss the application of Algorithm 1 to data obtained from a Danish road network.

To allow for reproducible research, we have made the source code for the numerical experiments discussed in Section VI-A and Section VI-B available at <https://github.com/alexjungaalto/ResearchPublic/tree/master/TVMIn>. The source code for the numerical experiments discussed in Section VI-C and Section VI-D can be found at <https://github.com/DruMara/GraphSignalRecovery>.

A. Two-Cluster Graph

In this experiment, we generate an empirical graph \mathcal{G} by first generating two clusters \mathcal{C}_1 and \mathcal{C}_2 of size $N/2 = 100$ drawn from an Erdős-Rényi ensemble with varying edge occurrence probability. We then connected those two clusters by randomly placing edges between them. The resulting empirical graph \mathcal{G} is then assigned a piece-wise constant graph signal \mathbf{x} of the form (9) using the partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2\}$. We apply Algorithm 1 to recover the graph signal \mathbf{x} based only on its values at the nodes in the training set \mathcal{M} which contains exactly one node from each of the two clusters, i.e., $|\mathcal{M}| = 2$.

Using Proposition 4, we can verify that the partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2\}$ is resolved by the training set \mathcal{M} by computing, for each

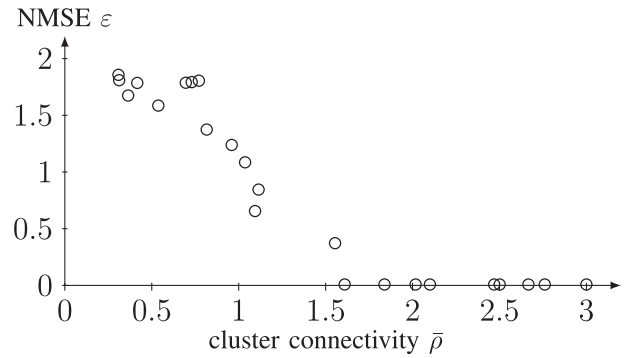


Fig. 3. NMSE achieved by Algorithm 1 for a two-cluster graph (see Fig. 1) with varying connectivity $\bar{\rho}$.




Fig. 4. Graph signal estimates delivered by Algorithm 1 for a two-cluster graph (see Fig. 1) with connectivity $\bar{\rho} \approx 0.26$. True graph signal (“x”) is piece-wise constant over clusters $\mathcal{C}_1 = \{1, 2, \dots, 100\}$, $\mathcal{C}_2 = \{101, 102, \dots, 200\}$. The signal is estimated from its values on $\mathcal{M} = \{1, 200\}$ using TV minimization (12) (“o”) and nLasso (15) for $\lambda = 10^{-4}$ (“x”), $\lambda = 10^{-3}$ (“⊗”) and $\lambda = 10^{-2}$ (“Δ”).

cluster \mathcal{C}_l the network flow between the labeled node $i \in \mathcal{C}_l \cap \mathcal{M}$ and the boundary $\partial\mathcal{C}_l$. Let $\rho^{(l)}$ denote the resulting flow value, normalized by the total weight of the boundary $2 \sum_{e \in \partial\mathcal{C}_l} W_e$. According to Proposition 4, the partition \mathcal{F} is resolved by \mathcal{M} if $\rho^{(l)} \geq 2$ for all $l = 1, 2$.

In Fig. 3, we depict the normalized mean squared error (NMSE) $\varepsilon := \|\mathbf{x} - \tilde{\mathbf{x}}^{(k)}\|_2^2 / \|\mathbf{x}^{(k)}\|_2^2$ incurred by Algorithm 1 (averaged over 10 i.i.d. simulation runs) for varying connectivity, as measured by the empirical average $\bar{\rho}$ of $\rho^{(1)}$ and $\rho^{(2)}$ (which have the same distribution due to the symmetric graph construction). The results in Fig. 3 agrees with our analysis (see Proposition 4 and Proposition 3) which predicts that TV minimization Algorithm 1 is accurate (incurring small NMSE) if the clusters \mathcal{C}_1 and \mathcal{C}_2 are well connected such that $\rho^{(1)}, \rho^{(2)} \geq 2$.

As indicated in Fig. 3, TV minimization fails to recover piece-wise constant graph signals \mathbf{x} of the form (9) if the cluster connectivity $\bar{\rho}$ is too small. We have depicted, for one particular realization of the two-cluster graph \mathcal{G} with connectivity $\bar{\rho} \approx 0.26$, the graph signal estimate obtained from TV minimization via Algorithm 1 in Fig. 4. Clearly, in this case TV minimization fails to correctly identify the underlying cluster structure and assigns many nodes to the signal values of the wrong cluster.




Fig. 5. NMSE achieved by Algorithm 1 for an empirical graph obtained from an SBM with varying edge probabilities p_{out} and p_{in} .

We have also included the estimates obtained from nLasso (15) for different choices of the parameter λ . According to Fig. 4, the nLasso estimates tend to be forced towards zero while TV minimization results in signal values closer to the initial labels $x_1 = 1/10$ and $x_{200} = -1/10$.

B. Stochastic Block Model

In this experiment, we generate an empirical graph \mathcal{G} using the SBM [37]. The graph \mathcal{G} consists of three clusters $\mathcal{C}_1, \mathcal{C}_2$ and \mathcal{C}_3 , each consisting of 10 nodes. An edge is placed between nodes i, j with probability p_{in} if they are in the same cluster and with probability p_{out} if they are from different clusters.

The empirical graph \mathcal{G} is then assigned a piece-wise constant graph signal \mathbf{x} (see (9)) using the partition $\mathcal{F} = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3\}$. We apply Algorithm 1 to recover the graph signal \mathbf{x} from its values at the nodes in the training set \mathcal{M} which contains exactly five nodes from each cluster such that $|\mathcal{M}| = 15$.

The (non-rigorous) condition (46) suggests that Algorithm 1 delivers an accurate estimate of \mathbf{x} whenever $p_{in}/p_{out} \gg (2/|\mathcal{M} \cap \mathcal{C}_l|)(|\mathcal{V}| - |\mathcal{C}_l|)$ for all $l = 1, 2, 3$. Inserting the particular SBM parameters used in this experiment yields the condition $p_{in}/p_{out} \gg 8$.

In Fig. 5, we depict the normalized mean squared error (NMSE) $\varepsilon := \|\mathbf{x} - \tilde{\mathbf{x}}^{(k)}\|_2^2 / \|\mathbf{x}^{(k)}\|_2^2$ incurred by Algorithm 1 (averaged over 100 i.i.d. simulation runs) for varying ratio p_{in}/p_{out} of SBM edge probabilities p_{in}, p_{out} . The results in Fig. 5 agree with the (non-rigorous) condition $p_{in}/p_{out} \gg 8$ such that TV minimization correctly recovers a piece-wise constant graph signal from few labeled nodes.

C. Big Data Framework Implementation

We have implemented Algorithm 2 using the higher-level programming interface GRAPHX [53] for the large-scale distributed computation framework SPARK [54]. The central concept of this framework is the distributed data structure (RDD) which is used to represent graph nodes, edges and associated signal values. Computations on graph data amount to transformations applied to RDDs. These RDD transformations are executed using efficient low-level distributed computing primitives [54].

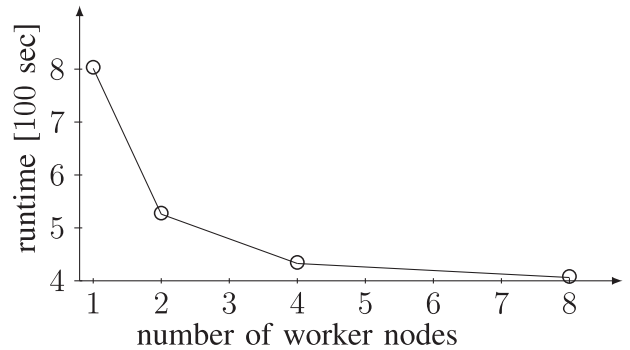
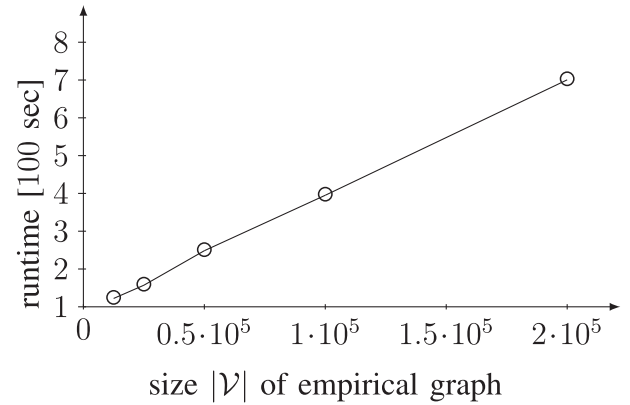


Fig. 6. (Top) Runtime of Algorithm 2 as size $|\mathcal{V}|$ of empirical graph increases. (Bottom) Runtime for varying number of “worker nodes”.

Using this implementation, we applied 2 to synthetic data obtained from the Lancichinetti-Fortunato-Radicchi (LFR) network model [55]. The probabilistic LFR model is widely used for benchmarking network algorithms [55] and aims at imitating some key characteristics of “real-world” networks such as the internet [4].

In order to study the scalability of Algorithm 2, we generated empirical graphs (using the LFR model) of varying size. We then measured the execution time of Algorithm 2 for a fixed number of 100 iterations.

As indicated by Fig. 6, the execution time scales linearly with the size (number of nodes) of the empirical graph. Fig. 6 also illustrates the effect of adding worker nodes to the cluster. In particular, for an empirical graph with size $|\mathcal{V}| = 10^5$, we determined the execution time of Algorithm 2 when the number of worker nodes is increased from 1 up to 8. As expected, the execution time decreases with increasing number of worker nodes. This decrease in execution time is, however, not exactly proportional to the increase of worker nodes due to communication overhead and data fragmentation associated with parallel computation frameworks [56].

D. Road Network

In this experiment we consider a dataset with empirical graph $\mathcal{G}_3 = (\mathcal{V}, \mathcal{E}, \mathbf{W})$ representing a road network in North Jutland (Denmark) [57], [58]. The edges \mathcal{E} of the graph \mathcal{G}_3 represent segments of road, and the nodes \mathcal{V} are intersections or terminations

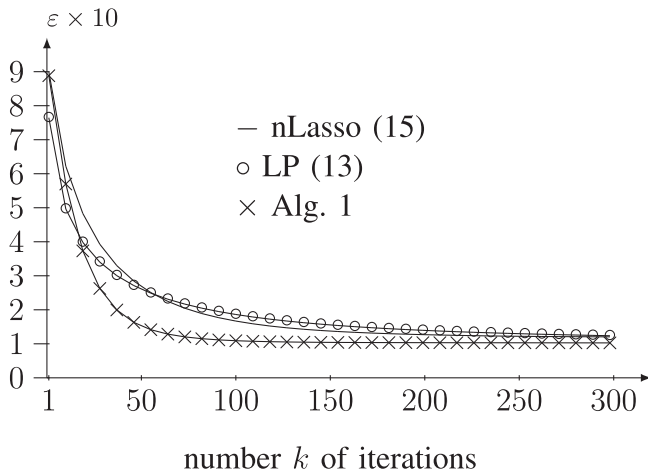


Fig. 7. NMSE ϵ incurred by learning methods applied to empirical graph \mathcal{G}_3 for increasing number of iterations.

of roads. The empirical graph \mathcal{G}_3 contains $N \approx 4 \cdot 10^5$ nodes and $E \approx 3.7 \cdot 10^5$ edges. The edge weights $W_{i,j}$ are obtained from the great-circle distances between intersections, measured in kilometres.

Each node $i \in \mathcal{V}$ of \mathcal{G}_3 is labeled with the elevation $x_i \in \mathbb{R}$ (relative to sea level) of the corresponding location in the road network. We construct a training set \mathcal{M} by selecting $|\mathcal{V}|/10$ nodes of \mathcal{G}_3 uniformly at random. Based on the labels of the nodes in the training set, we recover (predict) the labels on the remaining nodes using Algorithm 1, nLasso (15) and LP (13). The results are presented in Fig. 7, which depicts the NMSE achieved by the different algorithms after a certain number k of iterations (the iterations of the three methods having similar computational complexity).

As indicated by Fig. 7, TV minimization Algorithm 1 converges rapidly to a solution with smaller NMSE than nLasso (with manually tuned λ in (15)) and LP (13).

VII. CONCLUSION

We provided an analysis of the computational and statistical properties of TV minimization from a network flow perspective. This approach allowed us to derive conditions on network structure and available label information such that TV minimization accurately learns piece-wise constant graph signals. We also suggested an interpretation of primal-dual methods for TV minimization as distributed methods for network flow optimization.

Several topics for follow-up research can be identified. First, we plan to extend our network-flow based analysis of TV minimization to the closely related nLasso problem. This seems to be quite straightforward and might require merely a minor modification of the network flow constraints used to measure cluster connectivity. Regarding computational properties of TV based methods, we consider extending work on partial linear convergence of non-smooth Lasso type problems to TV minimization and nLasso. It turns out that such problems can be solved by iterative methods that converge linearly (at a geometric rate) to a solution up to a sub-optimality on the order of the

intrinsic estimation error, which cannot be overcome by any algorithm. Using the duality of TV minimization and network flows, such results would have immediate consequences for the complexity of network flow (and clustering) problems. As to the statistical properties of TV minimization, it would be interesting to extend our analysis from piece-wise constant to piece-wise smooth graph signals.

We expect our work to initiate cross-fertilization between network science and compressive graph signal processing. Convex methods for TV minimization are computationally attractive methods for handling massive networks and it would be interesting to investigate if they might outperform state-of-the-art network algorithms in some settings.

On the other hand, the duality of TV minimization and network flow optimization suggests new routes for combining primal-dual methods for TV minimization with existing methods for clustering and computing (approximating) maximum network flows. In particular, we might use maximum flow methods to (approximately) solve the dual of TV minimization in order to obtain an initial solution for TV minimization via the primal-dual optimality condition presented in Section III. The initial estimates for the solutions of the primal and dual problem might then, in turn, be used to warm-start the primal-dual iterations underlying Algorithm 1.

REFERENCES


- [1] A. Jung, A. Heimowitz, and Y. C. Eldar, "The network nullspace property for compressed sensing over networks," in *Proc. Int. Conf. Sampling Theory Appl.*, Tallinn, Estonia, Jul. 2017, pp. 644–648.
- [2] S. Chen, A. Sandryhaila, J. M. F. Moura, and J. Kovačević, "Signal recovery on graphs: Variation minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 17, pp. 4609–4624, Sep. 2015.
- [3] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, Aug. 2000.
- [4] M. E. J. Newman, *Networks: An Introduction*. London, U.K.: Oxford Univ. Press, 2010.
- [5] O. Chapelle, B. Schölkopf and A. Zien, Eds., *Semi-Supervised Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [6] S. L. Lauritzen, *Graphical Models*. Oxford, U.K.: Clarendon, 1996.
- [7] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [8] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (Adaptive Computation and Machine Learning). Cambridge, MA, USA: MIT Press, 2009.
- [9] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*, vol. 3. Hanover, MA, USA: Now Publishers, 2010.
- [10] M. Belkin, I. Matveeva, and P. Niyogi, "Regularization and semi-supervised learning on large graphs," in *Proc. Int. Conf. Comput. Learn. Theory*, 2004, vol. 3120, pp. 624–638.
- [11] B. Nadler, N. Srebro, and X. Zhou, "Statistical analysis of semi-supervised learning: The limit of infinite unlabelled data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, vol. 22, pp. 1330–1338.
- [12] A. E. Alaoui, X. Cheng, A. Ramdas, M. J. Wainwright, and M. I. Jordan, "Asymptotic behavior of ℓ_p -based Laplacian regularization in semi-supervised learning," in *Proc. Conf. Learn. Theory*, Jun. 2016, pp. 879–906.
- [13] S. Chen, R. Varma, A. Singh, and J. Kovačević, "Representations of piecewise smooth signals on graphs," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, China, Mar. 2016, pp. 6370–6374.
- [14] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imag. Vision*, vol. 40, no. 1, pp. 120–145, 2011.
- [15] Y.-X. Wang, J. Sharpnack, A. Smola, and R. Tibshirani, "Trend filtering on graphs," *J. Mach. Learn. Res.*, vol. 17, no. 105, pp. 1–41, 2016.

- [16] Z. Fan and L. Guan, "Approximate ℓ_0 -penalized estimation of piecewise-constant signals on graphs," *Ann. Statist.*, vol. 46, no. 6B, pp. 3217–3245, Dec. 2018.
- [17] O. Kurland, "The cluster hypothesis in information retrieval," in *Proc. Eur. Conf. Inf. Retrieval*, 2014, pp. 823–826.
- [18] A. Jung and M. Hulsebos, "The network nullspace property for compressed sensing of big data over networks," *Frontiers Appl. Math. Statist.*, vol. 4, p. 9, Apr. 2018.
- [19] A. Jung, "On the complexity of sparse label propagation," *Frontiers Appl. Math. Statist.*, vol. 4, p. 22, Jul. 2018.
- [20] J.-C. Hütter and P. Rigollet, "Optimal rates for total variation denoising," in *Proc. Annu. Conf. Learn. Theory*, Jun. 2016, vol. 49, pp. 1115–1146.
- [21] A. Jung, N. Quang, and A. Mara, "When is network Lasso accurate?" *Frontiers Appl. Math. Statist.*, vol. 3, p. 28, Jan. 2018.
- [22] R. Kyng, A. Rao, S. Sachdeva, and D. Spielman, "Algorithms for Lipschitz learning on graphs," in *Proc. Conf. Learn. Theory*, 2015, pp. 1190–1223.
- [23] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 123–231, 2013.
- [24] Y. C. Eldar, *Sampling Theory: Beyond Bandlimited Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2015.
- [25] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [26] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. New York, NY, USA: Springer, 2012.
- [27] M. Kabanava and H. Rauhut, "Cosparsity in compressed sensing," in *Compressed Sensing and Its Applications*, H. Boche, R. Calderbank, G. Kutyniok, and J. Vybiral, Eds. New York, NY, USA: Springer, 2015, pp. 315–339.
- [28] M. Zhao, M. Kaba, R. Vidal, D. P. Robinson, and E. Mallada, "Sparse recovery over graph incidence matrices," in *Proc. IEEE Conf. Decis. Control*, Dec. 2018, pp. 364–371.
- [29] D. Goldfarb and W. Yin, "Parametric maximum flow algorithms for fast total variation minimization," *SIAM J. Sci. Comput.*, vol. 31, pp. 3712–3743, 2009.
- [30] A. Chambolle, "Total variation minimization and a class of binary MRF models," in *Proc. Int. Workshop Energy Minimization Methods Comput. Vision Pattern Recognit.*, 2005, pp. 136–152.
- [31] V. Kolmogorov and R. Zabih, "What energy functions can be minimized via graph cuts?" *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 2, pp. 147–159, Feb. 2004.
- [32] D. P. Bertsekas, *Network Optimization: Continuous and Discrete Models*. Belmont, MA, USA: Athena Scientific, 1998.
- [33] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [34] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [35] U. von Luxburg, "A tutorial on spectral clustering," *Statist. Comput.*, vol. 17, no. 4, pp. 395–416, Dec. 2007.
- [36] D. A. Spielman and S. Hua Teng, "A local clustering algorithm for massive graphs and its application to nearly-linear time graph partitioning," *SIAM J. Comput.*, vol. 42, no. 1, pp. 1–26, Jan. 2013.
- [37] E. Abbe, "Community detection and stochastic block models: Recent developments," *J. Mach. Learn. Res.*, vol. 18, no. 177, pp. 1–86, 2018.
- [38] T. Pock and A. Chambolle, "Diagonal preconditioning for first order primal-dual algorithms in convex optimization," in *Proc. IEEE Int. Conf. Comput. Vision*, Barcelona, Spain, Nov. 2011, pp. 1762–1769.
- [39] A. Anis, A. Gadde, and A. Ortega, "Efficient sampling set selection for bandlimited graph signals using graph spectral proxies," *IEEE Trans. Signal Process.*, vol. 64, no. 14, pp. 3775–3789, Jul. 2016.
- [40] D. Hallac, J. Leskovec, and S. Boyd, "Network Lasso: Clustering and optimization in large graphs," in *Proc. Special Interest Group Knowl. Discovery Data Mining*, 2015, pp. 387–396.
- [41] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA, USA: Athena Scientific, Jun. 1999.
- [42] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 1970.
- [43] H. Bauschke and P. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, 2nd ed. New York, NY, USA: Springer, 2017.
- [44] S. Becker, J. Bobin, and E. J. Candès, "NESTA: A fast and accurate first-order method for sparse recovery," *SIAM J. Imag. Sci.*, vol. 4, no. 1, pp. 1–39, 2011.
- [45] L. Condat, "A primal–dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms," *J. Optim. Theory Appl.*, vol. 158, no. 2, pp. 460–479, Aug. 2013.
- [46] A. V. Goldberg and R. E. Tarjan, "Efficient maximum flow algorithms," *Commun. ACM*, vol. 57, no. 8, pp. 82–89, 2014.
- [47] J. B. Orlin, "Max flows in $O(nm)$ time, or better," in *Proc. 55th Annu. ACM Symp. Theory Comput.*, 2013, pp. 765–774.
- [48] A.-L. Barabási and M. Pósfai, *Network Science*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
- [49] D. Karger, "Random sampling in cut, flow, and network design problems," *Math. Oper. Res.*, vol. 24, no. 2, pp. 383–413, 1999.
- [50] J. Kleinberg and E. Tardos, *Algorithm Design*. Reading, MA, USA: Addison-Wesley, 2006.
- [51] D. Jungnickel, *Graphs, Networks and Algorithms*, 4th ed. Berlin, Germany: Springer, 2013.
- [52] A. J. Hoffman, "Some recent applications of the theory of linear inequalities to extremal combinatorial analysis," in *Proc. Symp. Appl. Math., Am. Math. Soc.*, 1960, pp. 113–127.
- [53] R. S. Xin, J. E. Gonzalez, M. J. Franklin, and I. Stoica, "GraphX: A resilient distributed graph system on spark," in *Proc. 1st Int. Workshop Graph Data Manage. Experience Syst.*, Jun. 2013.
- [54] M. Zaharia, M. Chowdhury, M. Franklin, S. Shenker, and I. Stoica, "Spark: Cluster computing with working sets," in *Proc. 2nd USENIX Conf. Hot Topics Cloud Comput.*, Berkeley, CA, USA, 2010.
- [55] A. Lancichinetti, S. Fortunato, and F. Radicchi, "Benchmark graphs for testing community detection algorithms," *Phys. Rev. E*, vol. 78, Oct. 2008, Art. no. 046110. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevE.78.046110>
- [56] G. M. Amdahl, "Validity of the single processor approach to achieving large-scale computing capabilities," in *Proc. Spring Joint Comput. Conf.*, 1967, pp. 483–485.
- [57] M. Lichman, "UCI machine learning repository," 2013. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [58] M. Kaul, B. Yang, and C. S. Jensen, "Building accurate 3D spatial networks to enable next generation intelligent transportation systems," in *Proc. IEEE Int. Conf. Mobile Data Manage.*, Milan, Italy, 2013, pp. 137–146.



Alexander Jung received the Ph.D. (Sub auspiciis Praesidentis) degree in statistical signal processing from the Vienna University of Technology (TU Vienna), Vienna, Austria, in 2012. He was a Postdoc with ETH Zurich and TU Vienna before he joined Aalto University, Espoo, Finland, as an Assistant Professor for machine learning in 2015. The focus of his research is to understand fundamental limits and efficient methods for machine learning problems arising in various application domains. The quality of his research has been recognized by a Best Student

Paper Award at the conference IEEE International Conference on Acoustics, Speech, and Signal Processing 2011 and an Amazon Web Services Machine Learning Award in 2018. He coauthored a paper that was finalist for the Best Student Paper Award at Asilomar 2017. While at Aalto University, he has redesigned the main course on machine learning and developed a new online course "Machine Learning with Python." He was elected as Teacher of the Year 2018 by the Department of Computer Science, Aalto University. He is the Chair of the Signal Processing and Circuits and Systems Chapter within the IEEE Finland Section.



Alfred O. Hero, III (F'97) received the B.S. (*summa cum laude*) degree from Boston University, Boston, MA, USA, in 1980 and the Ph.D. degree from Princeton University, Princeton, NJ, USA, in 1984, both in electrical engineering. Since 1984, he has been with the University of Michigan, Ann Arbor, MI, USA, where he is the John H. Holland Distinguished University Professor of electrical engineering and computer science and the R. Jamison and Betty Williams Professor of Engineering. His primary appointment is in the Department of Electrical Engineering and

Computer Science and he also has appointments, by courtesy, in the Department of Biomedical Engineering and the Department of Statistics. He is the Chair of the Committee on Applied and Theoretical Statistics of the U.S. National Academies of Science. His recent research interests are in high-dimensional spatio-temporal data, multi-modal data integration, statistical signal processing, and machine learning with particular interests in applications to social networks, network security and forensics, computer vision, and personalized health. He is a Section Editor of the *SIAM Journal on Mathematics of Data Science* and a Senior Editor of the *IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING*. He is on the editorial board of the *Harvard Data Science Review*. He is a moderator for the Electrical Engineering and Systems Science category of the arXiv. He is Co-General Chair of the IEEE International Symposium on Information Theory.

He was the founding Co-Director of the Michigan Institute for Data Science (2015–2018). From 2008 to 2013, he held the Digiteo Chaire d'Excellence, sponsored by Digiteo Research Park, Paris, located at the Ecole Supérieure d'Electricité, Gif-sur-Yvette, France. He has held other visiting positions at LIDS Massachusetts Institute of Technology (2006), Boston University (2006), I3S University of Nice, Sophia-Antipolis, France (2001), Ecole Normale Supérieure de Lyon (1999), Ecole Nationale Supérieure des Télécommunications, Paris, France (1999), Lucent Bell Laboratories (1999), Scientific Research Labs of the Ford Motor Company, Dearborn, MI, USA (1993), Ecole Nationale Supérieure des Techniques Avancées, Ecole Supérieure d'Electricité, Paris (1990), and M.I.T. Lincoln Laboratory (1987–1989).

He was the recipient of the University of Michigan Distinguished Faculty Achievement Award (2011) and the Stephen S. Attwood Excellence in Engineering Award (2017). He has been plenary and keynote speaker at several workshops and conferences. He is the recipient of several best paper awards including an IEEE Signal Processing Society Best Paper Award (1998), a Best Original Paper Award from the *Journal of Flow Cytometry* (2008), a Best Magazine Paper Award from the IEEE Signal Processing Society (2010), an SPIE Best Student Paper Award (2011), an IEEE ICASSP Best Student Paper Award (2011), an AISTATS Notable Paper Award (2013), and an IEEE ICIP Best Paper Award (2013). He is also the recipient of an IEEE Signal Processing Society Meritorious Service Award (1998), an IEEE Third Millennium Medal (2000), an IEEE Signal Processing Society Distinguished Lecturership (2002), an IEEE Signal Processing Society Technical Achievement Award (2014), the 2015 Society Award from the IEEE Signal Processing Society, and the 2020 Fourier Award from the IEEE.

He was the President of the IEEE Signal Processing Society (2006–2007). He was a member of the IEEE TAB Society Review Committee (2009), the IEEE Awards Committee (2010–2011), and served on the Board of Directors of the IEEE (2009–2011) as Director of Division IX (Signals and Applications). He served on the IEEE TAB Nominations and Appointments Committee (2012–2014). He is currently a member of the Big Data Special Interest Group of the IEEE Signal Processing Society. Since 2011, he has been a member of the Committee on Applied and Theoretical Statistics of the U.S. National Academies of Science and currently serves as a Chair.



Alexandru Cristian Mara was born in Aiud, Romania, in 1991. He received the B.Sc. degree (with first-class Hons.) in computer science from the Technical University of Madrid, Madrid, Spain, in 2014, and the M.Sc. (with Hons.) degree in computer communication and information sciences from Aalto University, Espoo, Finland, in 2017. He is currently working toward the Ph.D. degree in computer science engineering with Ghent University, Ghent, Belgium. His research interests include machine learning, graph mining, and big data analytics. Currently,

his research focuses on representation learning and dimensionality reduction.



Saeed Jahromi received the M.Sc. degree in computer science from Aalto University, Espoo, Finland, in 2018. He is currently a graduate student with the Faculty of Information Technology, Monash University, Melbourne, VIC, Australia.



Ayelet Heimowitz received the B.Sc. degree in computer engineering and the M.Sc. and Ph.D. degrees in electrical engineering from Bar Ilan University, Ramat Gan, Israel, in 2009, 2011, and 2016, respectively. She is currently a Postdoctoral Research Associate at the Program for Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA.



Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in physics and the other B.Sc. degree in electrical engineering both from Tel Aviv University, Tel Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002.

She is currently a Professor with the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel. She was previously

a Professor with the Department of Electrical Engineering, Technion, where she held the Edwards Chair in engineering. She is also a Visiting Professor with MIT, a Visiting Scientist with the Broad Institute, and an Adjunct Professor with Duke University and was a Visiting Professor with Stanford University. She is a member of the Israel Academy of Sciences and Humanities (elected 2017), and a EURASIP Fellow. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, learning and optimization methods, and their applications to biology and optics.

She is the recipient of many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyu Tomiyasu Award (2016). She was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. She is also the recipient of the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel and David Jacknow Award for Excellence in Teaching, and the Technion's Award for Excellence in Teaching (two times). She is the recipient of several best paper awards and best demo awards together with her research students and colleagues including the SIAM Outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel.

She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She is the Editor-in-Chief of *Foundations and Trends in Signal Processing*, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio-Inspired Signal Processing technical committees, and was an Associate Editor for the *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She was the Co-Chair and Technical Co-Chair of several international conferences and workshops.