

# Vector-Valued Graph Trend Filtering with Non-Convex Penalties

Rohan Varma, *Student Member*, Harlin Lee, *Student Member*, Jelena Kovačević, *Fellow*, and Yuejie Chi, *Senior Member*

**Abstract**—This work studies the denoising of piecewise smooth graph signals that exhibit inhomogeneous levels of smoothness over a graph, where the value at each node can be vector-valued. We extend the graph trend filtering framework to denoising vector-valued graph signals with a family of non-convex regularizers, which exhibit superior recovery performance over existing convex regularizers. Using an oracle inequality, we establish the statistical error rates of first-order stationary points of the proposed non-convex method for generic graphs. Furthermore, we present an ADMM-based algorithm to solve the proposed method and establish its convergence. Numerical experiments are conducted on both synthetic and real-world data for denoising, support recovery, event detection, and semi-supervised classification.

**Index Terms**—graph signal processing, graph trend filtering, semi-supervised classification, non-convex optimization

## I. INTRODUCTION

Signal estimation from noisy observations is a classic problem in signal processing and has applications in signal inpainting, collaborative filtering, recommendation systems and other large-scale data completion problems. Since noise can have deleterious, cascading effects in many downstream tasks, being able to efficiently and accurately filter and reconstruct a signal is of significant importance.

With the explosive growth of information and communication, signals are generated at an unprecedented rate from various sources, including social networks, citation networks, biological networks, and physical infrastructure [2]. Unlike time series or images, these signals lie on complex, irregular structures, and require novel processing techniques, leading to the emerging field of *graph signal processing* [3]-[5]. This framework models the structure by a graph and generalizes concepts and tools from classical discrete signal processing to graph signal processing. The associated graph-structured data are referred to as *graph signals*.

In graph signal processing, a common assumption is that the graph signal is smooth with respect to the graph, that is, the

signal coefficients do not vary much over local neighborhoods of the graph. However, this characterization is insufficient for many real-world signals that exhibit spatially inhomogeneous levels of smoothness over the graph. In social networks for example, within a given community or social circle, users' profiles tend to be homogeneous, while within a different social circle they will be of different, yet still have homogeneous values. Consequently, the signal is often characterized by large variations between regions and small variations within regions such that there are localized discontinuities and patterns in the signal. As a result, it is necessary to develop representations and algorithms to process and analyze such *piecewise smooth* graph signals.

In this work, we study the denoising of the class of piecewise smooth graph signals (including but not limited to piecewise constant graph signals), which is complementary to the class of smooth graph signals that exhibit homogeneous levels of smoothness over the graph. The reconstruction of smooth graph signals has been well studied in previous work both within graph signal processing [4]-[10] as well as in the context of Laplacian regularization [11], [12].

The Graph Trend Filtering (GTF) framework [13], which applies total variation denoising to graph signals [14], is a particularly flexible and attractive approach that regularizes discrete graph differences using the  $\ell_1$  norm. Although the  $\ell_1$  norm based regularization has many attractive properties [15], the resulting estimates are biased toward zero for large coefficients. To alleviate this bias effect, non-convex penalties such as the Smoothly Clipped Absolute Deviation (SCAD) penalty [16] and the Minimax Concave Penalty (MCP) [17] have been proposed as alternatives. These penalties behave similarly to the  $\ell_1$  norm when the signal coefficients are small, but tend to a constant when the signal coefficients are large. Notably, they possess the so-called *oracle property*: in the asymptotics of large dimension, they perform as well as the case where we know in advance the support of the sparse vectors [18]-[22].

In this work, we strengthen the GTF framework in [13] by considering a large family of possibly non-convex regularizers, including SCAD and MCP that exhibit superior reconstruction performance over  $\ell_1$  minimization for the denoising of piecewise smooth graph signals. Furthermore, we extend the GTF framework to allow vector-valued signals, e.g. time series [23], on each node of the graph, which greatly broadens the applicability of GTF to applications in social networks [24], gene networks, and semi-supervised classification [25]-[27]. Through theoretical analyses and empirical performance,

R. Varma, H. Lee and Y. Chi are with the Dept. of Electrical and Computer Engineering, Carnegie Mellon University. Emails: {rohanv, harlinl, yuejiec}@andrew.cmu.edu.

J. Kovačević is with the Tandon School of Engineering, New York University. Email: jelenak@nyu.edu.

This work is supported in part by NSF under grants CCF-1563918, CCF-1826519, CCF-1806154 and ECCS-1818571, by ONR under grant N00014-18-1-2142, by ARO under grant W911NF-18-1-0303, and by NIH under grant R01EB025018. A preliminary version of partial results in this paper was presented at the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing [1].

The first two authors contributed equally.

we demonstrate that the use of non-convex penalties improves the performance of GTF in terms of both reduced reconstruction error and improved support recovery, i.e. how accurately we can localize the discontinuities of the piecewise smooth signals. Our contributions can be summarized as follows:

- Theoretically, we derive the statistical error rates of the signal estimates, defined as first-order stationary points of the proposed GTF estimator. We derive the rates in terms of the noise level and the alignment of the ground truth signal with respect to the underlying graph, without making assumptions on the piecewise smoothness of the ground truth signal. The better the alignment, the more accurate the estimates. Importantly, the estimators do not need to be the global minima of the proposed non-convex problem, which are much milder requirements and important for the success of optimization. For denoising vector-valued signals, the GTF estimate is more accurate when each dimension of the signal shares similar patterns across the graph.
- Algorithmically, we propose an ADMM-based algorithm that is guaranteed to converge to a critical point of the proposed GTF estimator.
- Empirically, we demonstrate the performance improvements of the proposed GTF estimators with non-convex penalties on both synthetic and real data for signal estimation, support recovery, event detection, and semi-supervised classification.

The rest of this paper is organized as follows. Section II reviews related works and their relationships to the current paper. In Section III, we provide some background and definitions on graph signal processing and GTF. Section IV presents the proposed GTF framework with non-convex penalties and vector-valued graph signals. Section V develops its performance guarantees, and Section VI presents an efficient algorithm based on ADMM. Numerical performance of the proposed approach is examined on both synthetic and real-world data for denoising and semi-supervised classification in Section VII. Finally, we conclude in Section VIII and briefly discuss future work.

Throughout this paper, we use boldface letters  $\mathbf{a}$  and  $\mathbf{A}$  to represent vectors and matrices respectively. The transpose of  $\mathbf{A}$  is denoted as  $\mathbf{A}^\top$ . The  $\ell$ -th row of a matrix  $\mathbf{A}$  is denoted as  $\mathbf{A}_{\ell\cdot}$ , and the  $j$ -th column of a matrix  $\mathbf{A}$  is denoted as  $\mathbf{A}_{\cdot j}$ . The cardinality of a set  $T$  is denoted as  $|T|$ . For any set  $T \subseteq \{1, 2, \dots, r\}$  and  $\mathbf{x} \in \mathbb{R}^r$ , we denote  $(\mathbf{x})_T \in \mathbb{R}^{|T|}$  such that  $x_\ell \in (\mathbf{x})_T$  if and only if  $\ell \in T$  for  $\ell \in \{1, 2, \dots, r\}$ . Similarly, we define a submatrix  $\mathbf{A}_T \in \mathbb{R}^{|T| \times d}$  of  $\mathbf{A} \in \mathbb{R}^{r \times d}$  that corresponds to pulling out the rows of  $\mathbf{A}$  indexed by  $T$ . The  $\ell_2$  norm of a vector  $\mathbf{a}$  is defined as  $\|\mathbf{a}\|_2$ , and the spectral norm of a matrix  $\mathbf{A}$  is defined as  $\|\mathbf{A}\|$ . The pseudo-inverse of a matrix  $\mathbf{A}$  is defined as  $\mathbf{A}^\dagger$ . For a function  $h(\mathbf{x}) : \mathbb{R}^p \rightarrow \mathbb{R}$ , we write  $\nabla_{\mathbf{x}} h(\mathbf{x})|_{\mathbf{x}=\mathbf{x}^*}$  to denote the gradient or subdifferential of  $h(\mathbf{x})$ , if they exist, evaluated at  $\mathbf{x} = \mathbf{x}^*$ . When the intention is clear, this may be written concisely as  $\nabla h(\mathbf{x}^*)$ . We also follow the standard asymptotic notations. If for some constants  $C, N > 0$ ,  $|f(n)| \leq C|g(n)|$  for all  $n \geq N$ , then  $f(n) = O(g(n))$ ; if  $g(n) = O(f(n))$ , then  $f(n) = \Omega(g(n))$ . Finally,

Table I summarizes some key notations used in this paper for convenience.

| Symbol                 | Description                            | Dimension    |
|------------------------|--|--------------|
| $\Delta$               | oriented incidence matrix              | $m \times n$ |
| $\Delta^{(k+1)}$       | $k$ th order graph difference operator | $r \times n$ |
| $\beta$                | scalar-valued graph signal             | $n$          |
| $\mathbf{B}$           | vector-valued graph signal             | $n \times d$ |
| $\mathbf{y}$           | noisy observation of $\beta$           | $n$          |
| $\mathbf{Y}$           | noisy observation of $\mathbf{B}$      | $n \times d$ |
| $\Delta_{\ell\cdot}$   | $\ell$ -th row of $\Delta$             | $n$          |
| $\mathbf{B}_{\cdot j}$ | $j$ -th column of $\mathbf{B}$         | $n$          |
| $\ \Delta^{(k+1)}\ $   | spectral norm of $\Delta^{(k+1)}$      | 1            |

TABLE I: Key notations used in this paper.

## II. RELATED WORK AND CONNECTIONS

Estimators that adapt to spatial inhomogeneities have been well studied in the literature via regularized regression, total variation and splines [28]-[30]. Most of these methods involve locating change points or knots that denote a distinct change in the behavior of the function or the signal.

Our work is most related to the spatially adaptive GTF estimator introduced in [13] that *smoothens* or *filters* noisy signals to promote piecewise smooth behavior with respect to the underlying graph structure; see also [31]. In the same spirit as [28], the fused LASSO and univariate trend filtering framework developed in [14], [32], [33] use discrete difference operators to fit a time series signal using piecewise polynomials. The GTF framework generalizes univariate trend filtering by generalizing a path graph to arbitrarily complex graphs. Specifically, by appropriately defining the discrete difference operator, we can enforce piecewise constant, piecewise linear, and more generally piecewise polynomial behaviors over the graph structure. In comparison to previous work [13], in this paper, we have significantly expanded its scope by allowing vector-valued data over the graph nodes and a broader family of possibly non-convex penalties.

We note that while a significant portion of the relevant literature on GTF or the fused LASSO has focused on the sparsistency or support recovery conditions under which we can ensure the recovery of the location of the discontinuities or knots [34], [35], in this work, we study the asymptotic error rates of our estimator with respect to the mean squared error. Our analysis of error rates leverages techniques in [36], [37] that result in sharp error rates of total variation denoising via oracle inequalities, which we have carefully adapted to allow *non-convex* regularizers. The obtained error rates can be translated into bounds on support recovery or how well we can localize the boundary by leveraging techniques in [38].

Employing a graph-based regularizer that promotes similarities between the signal values at connected nodes has been investigated by many communities, such as graph signal processing, machine learning, applied mathematics, and network science. The Network LASSO proposed in [24], which is similar to the GTF framework with multi-dimensional or vector-valued data, focused on the development of efficient algorithms without any theoretical guarantees. The recent

works by Jung et al. [26], [27], [39] have analyzed the performance of Network LASSO for semi-supervised learning when the graph signal is assumed to be *clustered* according to the labels using the network null space property and the network compatibility condition inspired by related concepts in compressed sensing [40]. In contrast, our analysis does not make assumptions on the graph signal, and the error rate is adaptive to the alignment of the signal and the graph structure used in denoising.

A well-studied generalization of the sparse linear inverse problem is when there are multiple measurement vectors (MMV), and the solutions are assumed to have a *common sparsity pattern* [41]-[43]. Sharing information across measurements, and thereby exploiting the conformity of the sparsity pattern, has been shown to significantly improve the performance of sparse recovery in compressive sensing and sparse coding [44]-[48]. Motivated by these works, we consider vector-valued graph signals that are regarded as multiple measurements of scalar-valued graph signals sharing discontinuity patterns.

There are a few variants of non-convex penalties that promote sparsity such as SCAD, MCP, weakly convex penalties, and  $\ell_q$  ( $0 \leq q < 1$ ) minimization [18], [49]-[52]. In this paper, we consider and develop theory for a family of non-convex penalties parametrized similarly to that in [18], [50] with SCAD and MCP as our prime examples, although it is valid for other non-convex penalties.

### III. GRAPH SIGNAL PROCESSING, PIECEWISE SMOOTH SIGNALS, AND GRAPH TREND FILTERING

We consider an undirected graph  $G = (\mathcal{V}, \mathcal{E}, \mathbf{A})$ , where  $\mathcal{V} = \{v_1, \dots, v_n\}$  is the set of nodes,  $\mathcal{E} = \{e_1, \dots, e_m\}$  is the set of edges, and  $\mathbf{A} = [A_{jk}] \in \mathbb{R}^{n \times n}$  is the unweighted adjacency matrix – also known as the *graph shift operator* [3]. The edge set  $\mathcal{E}$  represents the connections of the undirected graph  $G$ , and the positive edge weight  $A_{jk}$  measures the underlying relation between the  $j$ th and the  $k$ th node, such as a similarity, a dependency, or a communication pattern. Let a scalar-valued *graph signal* be defined as

$$\boldsymbol{\beta} = [\beta_1, \beta_2, \dots, \beta_n]^\top \in \mathbb{R}^n,$$

where  $\beta_i$  denotes the signal coefficient at the  $i$ th node.

Let  $\boldsymbol{\Delta} \in \mathbb{R}^{m \times n}$  be the oriented incidence matrix of  $G$ , where each row corresponds to an edge. That is, if the edge  $e_i = (j, k) \in \mathcal{E}$  connects the  $j$ th node to the  $k$ th node ( $j < k$ ), the entries in the  $i$ th row of  $\boldsymbol{\Delta}$  is then given as

$$\Delta_{i\ell} = \begin{cases} -1, & \ell = j; \\ 1, & \ell = k; \\ 0, & \text{otherwise} \end{cases}.$$

The entries of the signal  $\boldsymbol{\Delta}\boldsymbol{\beta} = [(\beta_k - \beta_j)]_{(j,k) \in \mathcal{E}}$  specifies the unweighted pairwise differences of the graph signal over each edge. As a result,  $\boldsymbol{\Delta}$  can be interpreted as a *graph difference operator*. In graph signal processing, a signal is called smooth over a graph  $G$  if  $\|\boldsymbol{\Delta}\boldsymbol{\beta}\|_2^2 = \sum_{(j,k) \in \mathcal{E}} (\beta_k - \beta_j)^2$  is small.

#### A. Piecewise Smooth Graph Signals

In practice, the graph signal may not be necessarily smooth over the entire graph, but only locally within different pieces of the graph. To model inhomogeneous levels of smoothness over a graph, we say that a graph signal  $\boldsymbol{\beta}$  is piecewise constant over a graph  $G$  if many of the differences  $\beta_k - \beta_j$  are zero for  $(j, k) \in \mathcal{E}$ . Consequently, the difference signal  $\boldsymbol{\Delta}\boldsymbol{\beta}$  is sparse and  $\|\boldsymbol{\Delta}\boldsymbol{\beta}\|_0$  is small.

We can characterize *piecewise  $k$ th order polynomial* signals on a graph, where the piecewise constant case corresponds to  $k = 0$ , by generalizing the notion of graph difference operators. Specifically, we use the following recursive definition of the  $k$ th order graph difference operator  $\boldsymbol{\Delta}^{(k+1)}$  [13]. Let  $\boldsymbol{\Delta}^{(1)} = \boldsymbol{\Delta}$  for  $k = 0$ . For  $k \geq 1$ , let

$$\boldsymbol{\Delta}^{(k+1)} = \begin{cases} \boldsymbol{\Delta}^{(1)\top} \boldsymbol{\Delta}^{(k)} \in \mathbb{R}^{n \times n}, & \text{odd } k \\ \boldsymbol{\Delta}^{(1)} \boldsymbol{\Delta}^{(k)} \in \mathbb{R}^{m \times n}, & \text{even } k \end{cases}. \quad (1)$$

The signal  $\boldsymbol{\beta}$  is said to be a piecewise  $k$ th order polynomial graph signal if  $\|\boldsymbol{\Delta}^{(k+1)}\boldsymbol{\beta}\|_0$  is small. To further illustrate, let us consider the piecewise linear graph signal, corresponding to  $k = 1$ , as a signal whose value at a node can be linearly interpolated from the weighted average of the values at neighboring nodes. It is easy to see that this is the same as requiring the second-order differences  $\boldsymbol{\Delta}^\top \boldsymbol{\Delta}\boldsymbol{\beta}$  to be sparse. Similarly, we say that a signal has a piecewise quadratic structure over a graph if the differences between the second-order differences defined for piecewise linear signals are mostly zero, that is, if  $\boldsymbol{\Delta}\boldsymbol{\Delta}^\top \boldsymbol{\Delta}\boldsymbol{\beta}$  is sparse. Fig. 1 illustrates various orders of piecewise graph smooth signals over the Minnesota road network graph.

#### B. Denoising Piecewise Smooth Graph Signals via GTF

Assume we observe a noisy signal  $\mathbf{y}$  over the graph under i.i.d Gaussian noise:

$$\mathbf{y} = \boldsymbol{\beta}^* + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (2)$$

and seek to reconstruct  $\boldsymbol{\beta}^*$  from  $\mathbf{y}$  by leveraging the graph structure. When  $\boldsymbol{\beta}$  is a smooth graph signal, Laplacian smoothing [11], [12], [53]-[55] can be used, which solves the following problem:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\Delta}\boldsymbol{\beta}\|_2^2, \quad (3)$$

where  $\lambda > 0$ . However, it cannot localize abrupt changes in the graph signal when the signal is piecewise smooth.

Graph trend filtering (GTF) [13] is a flexible framework for estimation on graphs that is adaptive to inhomogeneity in the level of smoothness of an observed signal across nodes. The  $k$ th order GTF estimate is defined as:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\Delta}^{(k+1)}\boldsymbol{\beta}\|_1, \quad (4)$$

which can be regarded as applying total variation or fused LASSO with the graph difference operator  $\boldsymbol{\Delta}^{(k+1)}$  [14], [56]. The sparsity-promoting properties of the  $\ell_1$  norm have been well-studied [57]. Consequently, applying the  $\ell_1$  penalty in GTF sets many of the (higher-order) graph differences to zero




Fig. 1: Illustration of piecewise smooth signals on the Minnesota road graph. From left to right: piecewise constant ( $k = 0$ ), piecewise linear ( $k = 1$ ), and piecewise quadratic ( $k = 2$ ) graph signals. Note that the highlighted change points, i.e. the support of  $\Delta^{(k+1)}\beta^*$ , are edges for even  $k$  and nodes for odd  $k$ .

while keeping a small fraction of non-zero values. GTF is then *adaptive* over the graph; its estimate at a node adapts to the smoothness in its localized neighborhood.

#### IV. VECTOR-VALUED GTF WITH NON-CONVEX PENALTIES

In this section, we first extend GTF to allow a broader family of non-convex penalties, and then extend it to handle vector-valued signals over the graph.

##### A. (Non-)convex Penalties

The  $\ell_1$  norm penalty considered in (4) is well-known to produce biased estimates [58], which motivates us to extend the GTF framework to a broader class of sparsity-promoting regularizers that are not necessarily convex. We wish to minimize the following generalized  $k$ th order GTF loss function:

$$f(\beta) = \frac{1}{2} \|\mathbf{y} - \beta\|_2^2 + g(\Delta^{(k+1)}\beta; \lambda, \gamma), \quad \beta \in \mathbb{R}^n, \quad (5)$$

where

$$g(\Delta^{(k+1)}\beta) \triangleq g(\Delta^{(k+1)}\beta; \lambda, \gamma) = \sum_{\ell=1}^r \rho((\Delta^{(k+1)}\beta)_\ell; \lambda, \gamma)$$

is a regularizer defined as the sum of the penalty function  $\rho(\cdot; \lambda, \gamma) : \mathbb{R} \rightarrow \mathbb{R}$  applied element-wise to  $\Delta^{(k+1)}\beta$ . Here,  $r = m$  for even  $k$  and  $r = n$  for odd  $k$  to account for different dimensions of  $\Delta^{(k+1)}$ ; see (1). We will refer to the GTF estimator that minimizes  $f(\beta)$  as *scalar-GTF*.

Similarly to [18], [20], [50], we consider a family of penalty functions  $\rho(\cdot; \lambda, \gamma)$  that satisfies the following assumptions.

**Assumption 1.** Assume  $\rho(\cdot; \lambda, \gamma)$  satisfies the following:

- (a)  $\rho(t; \lambda, \gamma)$  satisfies  $\rho(0; \lambda, \gamma) = 0$ , is symmetric around 0, and is non-decreasing on the real non-negative line.
- (b) For  $t \geq 0$ , the function  $t \mapsto \frac{\rho(t; \lambda, \gamma)}{t}$  is non-increasing in  $t$ . Also,  $\rho(t; \lambda, \gamma)$  is differentiable for all  $t \neq 0$  and sub-differentiable at  $t = 0$ , with  $\lim_{t \rightarrow 0^+} \rho'(t; \lambda, \gamma) = \lambda$ . This upper bounds  $\rho(t; \lambda, \gamma) \leq \lambda|t|$ .

(c) There exists  $\mu > 0$  such that  $\rho(t; \lambda, \gamma) + \frac{\mu}{2}t^2$  is convex.

Many penalty functions satisfy these assumptions. Besides the  $\ell_1$  penalty, the non-convex SCAD [16] penalty

$$\rho_{\text{SCAD}}(t; \lambda, \gamma) = \lambda \int_0^{|t|} \min\left(1, \frac{(\gamma - u/\lambda)_+}{\gamma - 1}\right) du, \quad \gamma \geq 2, \quad (6)$$

and the MCP [17]

$$\rho_{\text{MCP}}(t; \lambda, \gamma) = \lambda \int_0^{|t|} \left(1 - \frac{u}{\lambda\gamma}\right)_+ du, \quad \gamma \geq 1 \quad (7)$$

also satisfy them. We note that Assumption 1 (c) is satisfied for SCAD with  $\mu \geq \mu_{\text{SCAD}} = \frac{1}{\gamma-1}$  and for MCP with  $\mu \geq \mu_{\text{MCP}} = \frac{1}{\gamma}$ . Fig. 2 illustrates the  $\ell_1$ , SCAD and MCP penalties for comparison. While the non-convexity means that in general, we may not always find the global optimum of  $f(\beta)$ , it often affords us many other advantages. SCAD and MCP both taper off to a constant value, and hence apply less shrinkage for higher values. As a result, they mitigate the bias effect while promoting sparsity. Further, they are smooth and differentiable for  $t \geq 0$  and are both upper bounded by the  $\ell_1$  penalty for all  $t$ .

##### B. Vector-Valued GTF

In many applications, the signals on each node are in fact multi-dimensional or *vector-valued*, e.g. time series in social networks, multi-class labels in semi-supervised learning, feature vectors of different objects in feature selection. Therefore, it is natural to consider an extension to the graph signal denoising problem, where the graph signal on each node is a  $d$ -dimensional vector instead of a scalar. In this scenario, we define a vector-valued graph signal to be piecewise smooth if it is piecewise smooth in each of its  $d$  dimensions, and assume their discontinuities to coincide over the same small set of edges or nodes. Further, we denote the vector-valued signal of interest as  $\mathbf{B}^* \in \mathbb{R}^{n \times d}$ , such that the  $i$ th row of the




Fig. 2: Illustration of  $\rho(\cdot; \lambda, \gamma)$  for  $\ell_1$ , SCAD ( $\gamma = 3.7$ ), and MCP ( $\gamma = 1.4$ ), where  $\lambda = 2$ . Both SCAD and MCP move towards  $\ell_1$  as  $\gamma$  increases.

matrix  $\mathbf{B}$  corresponds to the  $i$ th node of the graph. The noise model for the observation matrix  $\mathbf{Y} \in \mathbb{R}^{n \times d}$  is defined as

$$\mathbf{Y} = \mathbf{B}^* + \mathbf{E}, \quad (8)$$

where each element of  $\mathbf{E} \in \mathbb{R}^{n \times d}$  is drawn i.i.d from  $\mathcal{N}(0, \sigma^2)$ . A naïve approach is to estimate each column  $\mathbf{B}_{\cdot j}$  of  $\mathbf{B}$  separately via scalar-GTF:

$$\min_{\mathbf{B} \in \mathbb{R}^{n \times d}} \sum_{j=1}^d f(\mathbf{B}_{\cdot j}). \quad (9)$$

However, this formulation does not take full advantage of the multi-dimensionality of the graph signal. Instead, when the columns of  $\mathbf{B}$  are correlated, coupling them can be beneficial such that we encourage the sharing of information across dimensions or features. For example, if one column  $\mathbf{B}_{\cdot i}$  exhibits strong piecewise smoothness over the graph, and therefore has compelling evidence about the relationship between nodes, sharing that information to a related column  $\mathbf{B}_{\cdot j}$  can improve the overall denoising and filtering performance. As a result, we formulate a *vector-GTF* problem as follows:

$$\min_{\mathbf{B} \in \mathbb{R}^{n \times d}} \frac{1}{2} \|\mathbf{Y} - \mathbf{B}\|_F^2 + h(\Delta^{(k+1)} \mathbf{B}; \lambda, \gamma), \quad (10)$$

where the new penalty function  $h(\Delta^{(k+1)} \mathbf{B}) \triangleq h(\Delta^{(k+1)} \mathbf{B}; \lambda, \gamma) : \mathbb{R}^{r \times d} \rightarrow \mathbb{R}$  is the sum of  $\rho(\cdot; \lambda, \gamma)$  applied to the  $\ell_2$  norm of each row of  $\Delta^{(k+1)} \mathbf{B} \in \mathbb{R}^{r \times d}$ :

$$h(\Delta^{(k+1)} \mathbf{B}; \lambda, \gamma) = \sum_{\ell=1}^r \rho(\|(\Delta^{(k+1)} \mathbf{B})_{\ell}\|_2; \lambda, \gamma). \quad (11)$$

By enforcing sparsity on  $\{\|(\Delta^{(k+1)} \mathbf{B})_{\ell}\|_2\}_{1 \leq \ell \leq r}$ , we are coupling  $\Delta^{(k+1)} \mathbf{B}_{\cdot j}$  to be of similar sparsity patterns across  $j = 1, \dots, d$ . Note the difference from (9), where elements of  $(\Delta^{(k+1)} \mathbf{B})_{\ell}$  can be set to zero or non-zero independently.

## V. THEORETICAL GUARANTEES

In this section, we present the error rates and support recovery guarantees of the generalized GTF estimators, namely scalar-GTF (5) and vector-GTF (10), under the AWGN noise

model. Before continuing, we first define a few useful quantities. Let  $C_G$  be the number of connected components in the graph  $G$ , or equivalently, the dimension of the null space of  $\Delta^{(k+1)}$ . Further, let  $r$  be the number of rows of  $\Delta^{(k+1)}$ , and  $\zeta_k$  be the maximum  $\ell_2$  norm of the columns of  $\Delta^{(k+1)\dagger}$ .

### A. Error Rates of First-order Stationary Points

Due to non-convexity, global minima of the proposed GTF estimators may not be attainable. Therefore, it is more desirable to understand the statistical performance of any first-order stationary points of the GTF estimators. We call  $\hat{\beta} \in \mathbb{R}^n$  a stationary point of  $f(\beta)$ , if it satisfies

$$0 \in \nabla_{\beta} f(\beta)|_{\beta=\hat{\beta}}.$$

We further introduce the compatibility factor, which generalizes the notion used in [36] to allow vector-valued signals.

**Definition 1** (Compatibility factor). *Let  $\Delta^{(k+1)}$  be fixed. The compatibility factor  $\kappa_{T,d}$  of a set  $T \subseteq \{1, 2, \dots, r\}$  is defined as  $\kappa_{\emptyset,d} = 1$ , and for nonempty set  $T$ ,*

$$\kappa_{T,d}(\Delta^{(k+1)}) = \inf_{\mathbf{B} \in \mathbb{R}^{n \times d}} \left\{ \frac{\sqrt{|T|} \cdot \|\mathbf{B}\|_F}{\sum_{\ell \in T} \|(\Delta^{(k+1)} \mathbf{B})_{\ell}\|_2} \right\}.$$

To further build intuition, consider  $\sqrt{|T|} \kappa_{T,1}(\Delta)^{-1} = \sup_{\beta \in \mathbb{R}^n} \{ \|(\Delta)_T \beta\|_1 / \|\beta\|_2 \}$ . This is precisely the definition of  $\|(\Delta)_T\|_{1,2}$ , an induced norm of the  $|T| \times n$  submatrix of  $\Delta$ . If we consider signals with fixed power  $\|\beta\|_2^2 = 1$ ,  $\|(\Delta)_T \beta\|_1$  will depend on how much the  $T$  edges are connected to each other. Together with  $\|(\Delta)_T \beta\|_1 \leq \sqrt{|T|} \|(\Delta)_T \beta\|_2$ ,  $\kappa_{T,1}(\Delta)$  can be related to the restricted eigenvalue condition, which is often used to bound the performance of LASSO [40]. With slight abuse of notation, we write  $\kappa_T := \kappa_{T,d}$ .

We have the following oracle inequality that is applicable to the stationary points of the GTF estimators. The proof in Appendix A follows a construction that is similar to Theorem 2 in [36]. The oracle inequality holds for any  $\hat{\beta}$  that satisfies the first-order optimality condition, allowing the use of non-convex penalties. This mild condition on  $\hat{\beta}$  is a key difference from [13, Theorem 3] and [1, Theorem 1] that are applicable to the global minimizer, which is difficult to guarantee when using non-convex penalties. We also stress that although GTF was motivated by piecewise smooth graph signals, Theorem 1 holds for any graph  $G$  and graph signal  $\beta^*$ .

**Theorem 1** (Oracle inequality of GTF stationary points). *Assume  $\mu < 1 / \|\Delta^{(k+1)}\|^2$ . Fix  $\delta \in (0, 1)$ . For scalar-GTF (5), let  $\hat{\beta}$  be a stationary point. Set  $\lambda = \sigma \zeta_k \sqrt{2 \log(\frac{er}{\delta})}$ , then*

$$\frac{\|\hat{\beta} - \beta^*\|_2^2}{n} \leq \inf_{\beta \in \mathbb{R}^n} \left\{ \frac{\|\beta - \beta^*\|_2^2 + 4g((\Delta^{(k+1)} \beta)_{T^c})}{n} \right\} + \frac{2\sigma^2 \left[ C_G + 2\sqrt{2C_G \log(\frac{1}{\delta})} + \frac{8\zeta_k^2 |T|}{\kappa_T^2} \log(\frac{er}{\delta}) \right]}{n(1 - \mu \|\Delta^{(k+1)}\|^2)} \quad (12)$$

with probability at least  $1 - 2\delta$  for any  $T \subseteq \{1, 2, \dots, r\}$ . Similarly, for vector-GTF (10), let  $\hat{\mathbf{B}}$  be a stationary point.

Set  $\lambda = \sigma\zeta_k\sqrt{2d\log(\frac{edr}{\delta})}$ , then

$$\frac{\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2}{dn} \leq \inf_{\mathbf{B} \in \mathbb{R}^{n \times d}} \left\{ \frac{\|\mathbf{B} - \mathbf{B}^*\|_{\text{F}}^2 + 4h((\mathbf{\Delta}^{(k+1)}\mathbf{B})_{T^c})}{dn} \right\} + \frac{2\sigma^2 \left[ C_G + 2\sqrt{2C_G \log(\frac{d}{\delta})} + \frac{8\zeta_k^2|T|}{\kappa_T^2} \log(\frac{edr}{\delta}) \right]}{n(1 - \mu\|\mathbf{\Delta}^{(k+1)}\|^2)} \quad (13)$$

with probability at least  $1 - 2\delta$  for any  $T \subseteq \{1, 2, \dots, r\}$ .

**Remark 1.** Recall that  $\mu$  is defined in Assumption 1 (c), which characterizes how “non-convex” the regularizer is, and dictates the inflection point in Fig. 2. The assumption  $\mu < 1/\|\mathbf{\Delta}^{(k+1)}\|^2$  in Theorem 1 therefore implicitly constrains the level of non-convexity of the regularizer. Take MCP in (7) for example: since  $\mu \geq 1/\gamma$ , we can guarantee the existence of a valid  $\mu$  such that  $\mu < 1/\|\mathbf{\Delta}^{(k+1)}\|^2$  as long as we set  $\gamma > \|\mathbf{\Delta}^{(k+1)}\|^2$ .

Theorem 1 allows one to select  $\beta$  and  $T$  to optimize the error bounds on the right hand side of (12) and (13). For example, pick  $\beta = \beta^*$  in (12) (hence an “oracle”) to have

$$\frac{\|\widehat{\beta} - \beta^*\|_2^2}{n} \leq \frac{4g((\mathbf{\Delta}^{(k+1)}\beta^*)_{T^c})}{n} + \frac{2\sigma^2 [C_G^\delta + 8\zeta_k^2\kappa_T^{-2}|T|\log(\frac{er}{\delta})]}{n(1 - \mu\|\mathbf{\Delta}^{(k+1)}\|^2)}, \quad (14)$$

where  $C_G^\delta = C_G + 2\sqrt{2C_G \log(\frac{1}{\delta})}$ .

- By setting  $T$  as an empty set, we have

$$\frac{\|\widehat{\beta} - \beta^*\|_2^2}{n} \leq \frac{4g(\mathbf{\Delta}^{(k+1)}\beta^*)}{n} + \frac{2\sigma^2 C_G^\delta}{n(1 - \mu\|\mathbf{\Delta}^{(k+1)}\|^2)}, \quad (15)$$

which suggest that the reconstruction accuracy improves when the ground truth  $\beta^*$  is better aligned with the graph structure, and consequently the value of  $g(\mathbf{\Delta}^{(k+1)}\beta^*)$  is small.

- On the other hand, by setting  $T$  as the support of  $\mathbf{\Delta}^{(k+1)}\beta^*$ , we achieve

$$\frac{\|\widehat{\beta} - \beta^*\|_2^2}{n} \leq \frac{2\sigma^2 [C_G^\delta + 8\zeta_k^2\kappa_T^{-2}\|\mathbf{\Delta}^{(k+1)}\beta^*\|_0 \log(\frac{er}{\delta})]}{n(1 - \mu\|\mathbf{\Delta}^{(k+1)}\|^2)},$$

which grows linearly as we increase the sparsity level  $\|\mathbf{\Delta}^{(k+1)}\beta^*\|_0$ .

Similar discussions can be conducted for vector-GTF by choosing  $\mathbf{B} = \mathbf{B}^*$  in (13). More importantly, we can directly compare the performance of vector-GTF with scalar-GTF, which was formulated for vector-valued graph signals in (9). The error bound of vector-GTF pays a small price in the order of  $\log d$ , but is tighter than scalar-GTF if  $h((\mathbf{\Delta}^{(k+1)}\mathbf{B}^*)_{T^c}) \ll \sum_{j=1}^d g((\mathbf{\Delta}^{(k+1)}\mathbf{B}^*)_{T^c})$ . This suggests that vector-GTF is much more advantageous when the support sets of  $\mathbf{\Delta}^{(k+1)}\mathbf{B}^*_j$  for  $j = 1, \dots, d$  overlap, i.e. when the local discontinuities and patterns in  $\mathbf{B}^*_j$  are shared.

## B. Comparison with Scalar-GTF using $\ell_1$ Regularization

We compare our error bound for scalar-GTF that is on  $\|\widehat{\beta} - \beta^*\|_2^2/n$  with [13, Theorem 3], which is obtained for GTF with the  $\ell_1$  penalty, reproduced below for convenience.

**Theorem 2** (Basic error bound of  $\ell_1$  GTF minimizer). *If  $\lambda = \Theta(\sigma\zeta_k\sqrt{\log r})$ , then  $\widehat{\beta}$ , the minimizer of (4), satisfies*

$$\frac{\|\widehat{\beta} - \beta^*\|_2^2}{n} = O\left(\frac{\lambda\|\mathbf{\Delta}^{(k+1)}\beta^*\|_1}{n} + \frac{\sigma^2 C_G}{n}\right).$$

The above bound is comparable to our bound in the special case of setting  $T$  to an empty set, i.e. (15). The first term of the bound in (15) is upper bounded by that of Theorem 2. The non-convex regularization yields especially tighter bounds when  $\mathbf{\Delta}^{(k+1)}\beta^*$  contains large coefficients, so that  $g(\mathbf{\Delta}^{(k+1)}\beta^*) \ll \lambda\|\mathbf{\Delta}^{(k+1)}\beta^*\|_1$ . On the other hand, the second term of (15) contains  $1 - \mu\|\mathbf{\Delta}^{(k+1)}\|^2$  in the denominator, which makes it an upper bound of the second term in Theorem 2. This gap can be brought down by choosing a larger  $\gamma$ , which allows one to pick a smaller  $\mu$ , as mentioned in Remark 1. However, as  $\gamma \rightarrow \infty$ , non-convex SCAD and MCP also tends to  $\ell_1$ , which erases the improvement from using non-convex regularizers in the first term of the bound. This indicates a trade-off in the overall error bound based on  $\gamma$ , or the “non-convexity” of the regularizers chosen for scalar-GTF.

To sum up, despite being non-convex, we can guarantee that any stationary point of the proposed GTF estimator possesses strong statistical guarantees.

## C. Error Rates for Erdős-Rényi Graphs

We next specialize Theorem 1 to the Erdős-Rényi random graphs using spectral graph theory [59]. Let  $d_{\max}$  and  $d_0$  respectively be the maximum and expected degree of the graph. It is known that for any graph it holds [13]

$$\zeta_k \leq \lambda_{\min}(\mathbf{\Delta}^{(2)})^{-\frac{k+1}{2}}, \quad (16)$$

where  $\lambda_{\min}(\mathbf{\Delta}^{(2)})$  is the smallest non-zero eigenvalue of the graph Laplacian matrix  $\mathbf{\Delta}^{(2)}$ . Moreover, we have  $\|\mathbf{\Delta}^{(k+1)}\|^2 = (\lambda_{\max}(\mathbf{\Delta}^{(2)}))^{k+1}$ , and  $d_{\max} + 1 \leq \lambda_{\max}(\mathbf{\Delta}^{(2)}) \leq 2d_{\max}$  [60]. Next, we present a simple lower bound on  $\kappa_T$ , which is proved in Appendix B.

**Proposition 1** (Bound on  $\kappa_T$ ).  $\kappa_T$  is bounded for any  $T$  and  $d$  as

$$\kappa_T(\mathbf{\Delta}^{(k+1)}) \geq (2d_{\max})^{-\frac{k+1}{2}}.$$

For an Erdős-Rényi random graph, if  $d_0 = \Omega(\log(n))$ , we have  $d_{\max} = O(d_0)$  almost surely [61, Corollary 8.2] and  $C_G = 1$ . Furthermore,  $\lambda_{\min}(\mathbf{\Delta}^{(2)}) = \Omega(d_0 - \sqrt{d_0})$  [13], [59], [62], and  $r = n$  for odd  $k$  and  $r = O(nd_0)$  for even  $k$ . Therefore, with probability at least  $1 - n^{-10}$ , we have

$$\frac{\|\widehat{\beta} - \beta^*\|_2^2}{n} \lesssim \frac{\sigma^2 \sqrt{\log n}}{n} + \min\left\{\frac{g(\mathbf{\Delta}^{(k+1)}\beta^*)}{n}, \frac{\sigma^2 \|\mathbf{\Delta}^{(k+1)}\beta^*\|_0 \log n}{n}\right\},$$

where  $g(\Delta^{(k+1)}\beta^*) \lesssim \frac{\sigma\|\Delta^{(k+1)}\beta^*\|_1\sqrt{\log n}}{d_0^{(k+1)/2}}$  by plugging in  $g(\Delta^{(k+1)}\beta^*) \leq \lambda\|\Delta^{(k+1)}\beta^*\|_1$ .

These results are also applicable to  $d_0$ -regular Ramanujan graphs [62].

#### D. Support Recovery

An alternative yet important metric for gauging the success of the proposed GTF estimators is support recovery, which aims to localize the discontinuities in the piecewise smooth graph signals, i.e. the support set of  $\Delta^{(k+1)}\beta^*$ , that is

$$S_k(\beta^*) = \left\{ t \in \{1, \dots, r\} : (\Delta^{(k+1)}\beta^*)_t \neq 0 \right\}. \quad (17)$$

In particular, for odd  $k$ , the discontinuities correspond to graph nodes; and for even  $k$ , they correspond to the edges. Let  $\hat{\beta}$  be the GTF estimate of the graph signal. The quality of the support recovery can be measured using the *graph screening distance* [38]. For any  $t_1 \in S_k(\beta^*)$  and  $t_2 \in S_k(\hat{\beta})$ , let  $d_G(t_1, t_2)$  denote the length of the shortest path between them. The distance of  $S_k(\hat{\beta})$  from  $S_k(\beta^*)$  is then defined as

$$d_G(S_k(\hat{\beta})|S_k(\beta^*)) = \begin{cases} \max_{t_1 \in S_k(\beta^*)} \min_{t_2 \in S_k(\hat{\beta})} d_G(t_1, t_2), & \text{if } S_k(\beta^*) \neq \emptyset \\ \infty & \text{otherwise} \end{cases}. \quad (18)$$

Interestingly, Lin et al. [38] showed recently that under mild assumptions, one can translate the error bound into a support recovery guarantee. Specifically, letting  $R_n$  be the RHS of (12) that bounds the error  $\|\hat{\beta} - \beta^*\|_2^2/n$  in Theorem 1, we have

$$d_G(S_k(\hat{\beta})|S_k(\beta^*)) = \begin{cases} O\left(\frac{R_n}{H_r^2}\right), & k = 0 \\ O\left(\frac{R_n^{1/3}}{H_r^{2/3}}\right), & k = 1 \end{cases}, \quad (19)$$

where  $H_r$  quantifies the minimum level of discontinuity, defined as the minimum absolute value of the non-zero values of  $\Delta^{(k+1)}\beta^*$ , i.e.

$$H_r = \min_{t \in S_k(\beta^*)} |(\Delta^{(k+1)}\beta^*)_t|. \quad (20)$$

Consequently, this leads to support recovery guarantees of the proposed GTF estimators. Numerical experiment in Section VII-A verifies the superior performance of the non-convex regularizers over the  $\ell_1$  regularizer for support recovery.

## VI. ADMM ALGORITHM AND ITS CONVERGENCE

There are many algorithmic approaches to optimize the vector-GTF formulation in (10), since scalar-GTF (5) can be regarded as a special case with  $d = 1$ . In this section, we illustrate the approach adopted in this paper, which is the Alternating Direction Method of Multipliers (ADMM) framework for solving separable optimization problems [63].

Via a change of variable as  $Z = \Delta^{(k+1)}B$ , we can transform (10) to

$$\min_{B \in \mathbb{R}^{n \times d}} \frac{1}{2} \|\mathbf{Y} - B\|_F^2 + h(Z; \lambda, \gamma) \quad \text{s.t. } Z = \Delta^{(k+1)}B.$$

---

### Algorithm 1 ADMM for solving (10)

---

```

1: Inputs:  $\mathbf{Y}$ ,  $\Delta^{(k+1)}$ , and parameters  $\lambda$ ,  $\gamma$ ,  $\tau$ 
2: Initialize:
    $B \leftarrow \mathbf{Y}$  or  $B_{\text{init}}$  if given.
    $D \leftarrow \Delta^{(k+1)}$ ,  $Z \leftarrow DB$ ,  $U \leftarrow DB - Z$ 
    $X \leftarrow (I + \tau D^\top D)^{-1}$ 
3: repeat
4:   for  $j \leftarrow 1$  to  $\text{num\_cols}(B)$  do
5:      $B_{\cdot j} \leftarrow X(\tau D^\top(Z_{\cdot j} - U_{\cdot j}) + Y_{\cdot j})$ 
6:   end for
7:   for  $\ell \leftarrow 1$  to  $\text{num\_rows}(DB)$  do
8:      $Z_{\ell} \leftarrow \text{Prox}_\rho(\|D_{\ell} B + U_{\ell}\|_2; \lambda/\tau)$ 
9:   end for
10:   $U \leftarrow U + DB - Z$ 
11: until termination

```

---

Its corresponding Lagrangian can be written as:

$$\mathcal{L}(B, Z, U) = \frac{1}{2} \|\mathbf{Y} - B\|_F^2 + h(Z; \lambda, \gamma) + \frac{\tau}{2} \|\Delta^{(k+1)}B - Z + U\|_F^2 - \frac{\tau}{2} \|U\|_F^2, \quad (21)$$

where  $U \in \mathbb{R}^{r \times d}$  is the Lagrangian multiplier, and  $\tau$  is the parameter. Alg. 1 shows the ADMM updates based on the Lagrangian in (21). Recall the proximal operator is defined as  $\text{Prox}_f(\mathbf{v}; \alpha) = \arg\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + \alpha f(\mathbf{x})$  for a function  $f(\cdot)$ .  $\ell_1$ , SCAD and MCP all admit closed-form solutions of  $\text{Prox}$ , which are simple thresholding operations [64]. Furthermore, we have the following convergence guarantee for Alg. 1, whose proof is provided in Appendix C. Theorem 3 implies that the output of Alg. 1 satisfies Theorem 1.

**Theorem 3.** *Let  $\tau \geq \mu$ , then Alg. 1 converges to a stationary point of (10).*

In addition, we provide a detailed time complexity analysis of Alg. 1 in Table II. Note that since  $\Delta$  is a sparse matrix with exactly  $2m$  non-zero entries, Alg. 1 can run much faster when  $k = 0$ . As a preprocessing step for each  $D$ , we compute  $V \in \mathbb{R}^{n \times n}$  and  $S \in \mathbb{R}^{n \times n}$ , the eigenvectors and eigenvalues of  $D^\top D$ , exactly once.  $X = V(I + \tau S)^{-1}V^\top$  can then be initialized very efficiently for all experiments that use  $D$ .

|                                | $k \geq 1$        | $k = 0$          |
|--------------------------------|-------------------|------------------|
| $D^\top D$ eigen decomposition | $O(rn^2 + n^3)$   | $O(m^2 + n^3)$   |
| $Z$ initialization             | $O(rnd)$          | $O(md)$          |
| $X$ initialization             | $O(n^2)$          | $O(n^2)$         |
| $B$ update                     | $O(d(nr + n^2))$  | $O(d(m + n^2))$  |
| $DB$ calculation               | $O(rnd)$          | $O(md)$          |
| $Z, U$ update                  | $O(rd)$           | $O(rd)$          |
| Total after $t$ iterations     | $O(tdrn + tdn^2)$ | $O(tdm + tdn^2)$ |

TABLE II: Time complexity analysis of Alg. 1.

## VII. NUMERICAL EXPERIMENTS

For the following experiments, we fixed  $\gamma = 3.7$  for SCAD,  $\gamma = 1.4$  for MCP. Further, the graphs we use in the following

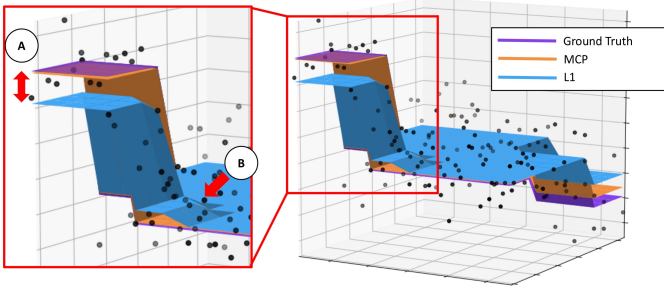


Fig. 3: Scalar-GTF with MCP (orange) has much lower bias than scalar-GTF with  $\ell_1$  (blue) when estimating a piecewise constant signal over a  $12 \times 12$  grid graph. See highlighted regions pointed by red arrows in A and B. The scatter points correspond to a noisy signal with 5dB SNR.




Fig. 4: The ROC curve for classifying whether an edge lies on a boundary for the Minnesota road graph signal shown in Fig. 5. The input SNR of the noisy piecewise constant signal is 7.8dB.

experiments satisfy Assumption 2 for this choice of  $\gamma$ . Unless explicitly mentioned, we tuned  $\lambda$  and  $\frac{\tau}{\lambda}$  for each experiment using the Hyperopt toolbox [65]. To meet the convergence criteria in Theorem 3, we enforced  $\tau \geq 1/\gamma$ . SCAD/MCP were warm-started with the GTF estimate with  $\ell_1$  penalty. Python packages PyGSP [66] and NetworkX [67] were used to construct and plot graphs. The input signal SNR was calculated as  $10 \log_{10}(\|\mathbf{B}^*\|_F / \sigma^2 nd)$ , while the reconstructed signal SNR was calculated as  $10 \log_{10}(\|\mathbf{B}^*\|_F / \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F)$ , where  $\widehat{\mathbf{B}}$  was the reconstruction. Computation time was measured with MacBook Pro 2017 with an 2.9 GHz Intel Core i7 and 16GB RAM. Our code is available at <https://github.com/HarlinLee/nonconvex-GTF-public>.

#### A. Denoising via GTF with Non-convex Regularizers

We first highlight via synthetic examples two important advantages that non-convex regularizers provide over the  $\ell_1$  penalty.

- **Bias Reduction:** We demonstrate the reduction in signal bias in Fig. 3 for the graph signal defined over a  $12 \times 12$

2D-grid graph, using both the  $\ell_1$  penalty and the MCP penalty. Clearly, the MCP estimate (orange) has less bias than the  $\ell_1$  estimate (blue), and can recover the ground truth surface (purple) more closely.

- **Support Recovery:** We illustrate the improved support recovery performance of non-convex penalties on localizing the boundaries for a piecewise constant signal on the Minnesota road graph, shown in Fig. 5. Particularly, we look at how well our estimator localizes the support of  $\Delta^{(k+1)}\beta^*$ , that is, the discontinuity of the piecewise constant graph signal by looking at how well we can classify an edge as connecting two nodes in the same piece or being a cut edge across two pieces. By sweeping the regularization parameter  $\lambda$ , we obtain the ROC curve in Fig. 4, i.e. the true positive rate versus the false positive rate of classifying a cut edge correctly, and see that scalar-GTF with MCP and SCAD consistently outperforms the scalar-GTF with  $\ell_1$  penalty.

Then, we compare the performance of GTF using non-convex regularizers such as SCAD and MCP with that using the  $\ell_1$  norm more rigorously. For the ground truth signal  $\beta^*$ , we construct a piecewise constant signal on a  $20 \times 20$  2D-grid graph and the Minnesota road graph [66] as shown in the left panel of Fig. 5, and add different levels of noise following (2). We recover the signal by scalar-GTF with Alg. 1, and plot the SNR of the reconstructed signal versus the SNR of the input signal *in solid lines* in the middle panel of Fig. 5, averaged over 10 and 20 realizations, respectively. SCAD/MCP consistently outperforms  $\ell_1$  in denoising graph signals defined over both regular and irregular structures.

#### B. Denoising Vector-valued Signals via GTF

We compare the performance of vector-GTF in (10) with (9), which applies scalar-GTF to each column of the vector-valued graph signal. The convex  $\ell_1$  norm, and the non-convex SCAD and MCP are employed. We reuse the same ground truth graph signals over the 2D-grid graph and the Minnesota road graph constructed in Section VII-A in Fig. 5.  $d$  independent noisy realizations of the graph signal are concatenated to construct a noisy vector-valued graph signal with dimension  $d = 10$  on the 2D-grid graph and with  $d = 20$  on the Minnesota road graph. We recover the vector-valued signal by minimizing vector-GTF (10) with Alg. 1.

The middle panel of Fig. 5 plots the average SNR of the reconstructed signal versus the average SNR of the input signal *in dotted lines*. We emphasize that the performance of (9) is the same as applying scalar-GTF to each realization, which is shown in the middle panel of Fig. 5 in solid lines. As before, SCAD/MCP consistently outperforms  $\ell_1$  in denoising signals over both regular and irregular graphs. Furthermore, as expected, due to the sharing of information across realizations, vector-GTF consistently outperforms scalar-GTF, especially in the low SNR regime.

The right panel of Fig. 5 plots the computation time versus the gain in SNR from denoising via vector-GTF. 10 trials are performed for each regularizer with the input signal SNR fixed at 20dB. Parameter tuning and eigen decomposition of  $\Delta^{(2)}$



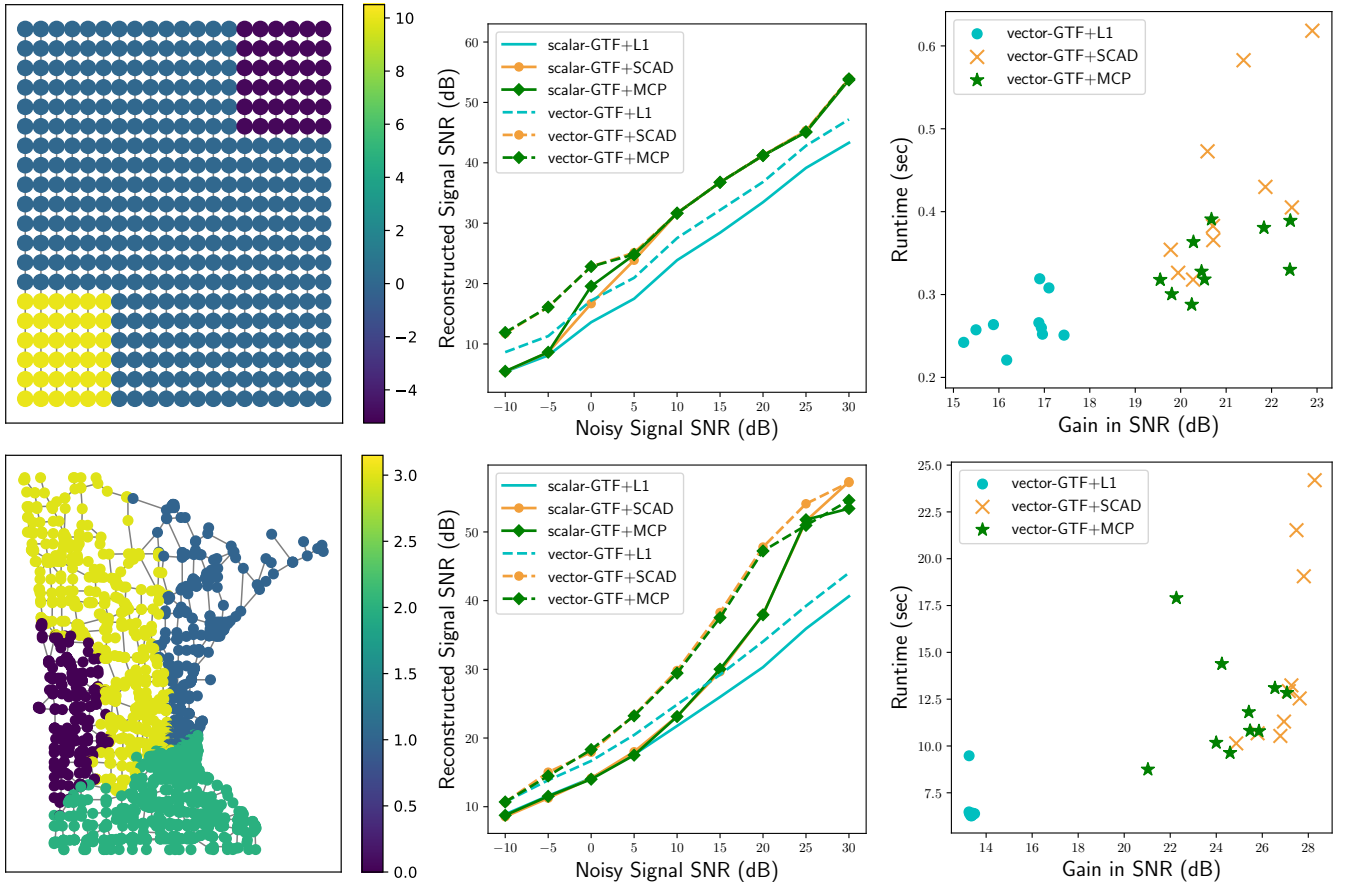


Fig. 5: The left panel shows the ground truth piecewise constant signals on  $20 \times 20$  2D-grid graph (top), and Minnesota road graph (bottom). The middle panel shows their corresponding plots of input signal SNR versus reconstructed signal SNR, averaged over 10 and 20 realizations, respectively. Finally, the right panel plots the computation time against gain in SNR from denoising via vector-GTF. 10 trials were performed for each regularizer, where the input signal SNR was fixed at 20dB.

are preprocessing steps, and hence they are not included in the time measurement; but for reference, the eigen decomposition took 0.025 and 2.5 seconds for 2D-grid and Minnesota graphs, respectively. Since GTF with non-convex regularizers are warm-started by the  $\ell_1$  estimate, the runtime for  $\ell_1$  GTF is added to the SCAD/MCP runtime. Overall, running vector-GTF with SCAD/MCP after once with  $\ell_1$  takes more time, but with large benefits in the denoising performance. Even with the additional computation time, Vector-GTF runs reasonably fast; with the Minnesota road network, where  $n = 2642$  and  $m = 3304$ , computation takes less than 25 seconds.

We further investigate the benefit of sharing information across measurements or realizations in the following experiment, using the same ground truth signal on the 2D-grid graph. We stack eight noisy realizations of this same piecewise constant signal to build a vector-valued signal. We construct these noisy measurements by scaling each one of them differently and randomly such that each will have  $\text{SNR} \sim \log_{10} \text{Uniform}[-10, 30]\text{dB}$  under (2). This has the effect of rendering some measurements more *informative* than others, and potentially allowing vector-GTF to reap the benefits of sharing information across measurements. We recover the 8-dimensional graph signal via Alg. 1 using  $\ell_1$ , SCAD, and

|                       | Average   | #1        | #2        | #3        | #4        | #5        | #6        | #7        | #8        |
|-----------------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| Input SNR (dB)        | 8.7       | -14       | 0         | 0         | 3.5       | 5.8       | 12        | 29        | 34        |
| Vector-GTF + $\ell_1$ | 29        | <b>10</b> | <b>20</b> | <b>23</b> | <b>26</b> | <b>36</b> | <b>37</b> | 39        | 38        |
| Scalar-GTF + $\ell_1$ | 21        | 0         | 11        | 13        | 16        | 18        | 26        | 41        | 45        |
| Vector-GTF + SCAD     | <b>32</b> | <b>10</b> | <b>20</b> | 22        | 25        | <b>36</b> | 35        | <b>49</b> | <b>61</b> |
| Scalar-GTF + SCAD     | 29        | 0         | 15        | 17        | 25        | 35        | 34        | 47        | 60        |
| Vector-GTF + MCP      | <b>32</b> | <b>10</b> | <b>20</b> | 22        | 25        | <b>36</b> | 35        | <b>49</b> | <b>61</b> |
| Scalar-GTF + MCP      | 29        | 0         | 15        | 22        | 24        | 30        | 33        | <b>49</b> | 60        |

TABLE III: Noisy input and reconstructed signal SNRs for eight measurements of varying input SNRs, rounded to two significant figures. Highest reconstructed signal SNR for each measurement is in **bold**.

MCP regularizers, and in Table III, report the input signal and reconstructed signal SNRs for each measurement in addition to the average SNRs.  $\lambda$  is fixed at  $0.5\sigma^2$ .

First of all, notice that as before, using SCAD/MCP generally achieves results with higher SNR than using  $\ell_1$ , and that on average, minimizing (10) outperforms minimizing (9). The effect of sharing information across measurements is




Fig. 6: Top left: the noisy signal on the Manhattan road network is the change in the taxi pickup and dropoff count during the 2011 NYC Gay Pride. Top right: areas of Pride events, where the traffic was blocked off. Bottom: the GTF estimates using  $\ell_1$  and MCP. The GTF estimate with MCP better detects and localizes the event, compared to the one using  $\ell_1$  penalty.

most apparent in low SNR settings, when information about the boundaries of the graph signal can be borrowed from higher SNR signals to improve the estimation. On the other hand, sharing information with noisier signals does not help denoising signals with high input SNR. However, it is worth noting that, unlike  $\ell_1$ , SCAD/MCP does not see decrease in its performance in the high SNR settings.

### C. Event Detection with NYC Taxi Data

To further illustrate graph trend filtering on a real-world dataset, we consider the road network of Manhattan where the nodes correspond to junctions [68]. We map the pickups and dropoffs of the NYC taxi trip dataset [69], [70] to the nearest road junctions, and define the total count at that junction to be the signal value on the corresponding graph node. The signal of interest, plotted on the top left panel of Fig. 6, is the difference between the *event* graph signal on the day of NYC Gay Pride parade, 12-2pm on June 26, 2011, and the *seasonal average* graph signal at the same time during the 8 nearest Sundays. During the event, no pickups and dropoffs could occur in the areas shown in the top right panel of Fig. 6 [71]. We denoise the signal via GTF using both  $\ell_1$  and MCP,

where we chose  $\lambda$  such that  $\|\Delta\hat{\beta}\|_0 \approx 200$ . Once again, we observe the GTF estimate with MCP produces sharper traces around the parade route, indicating better capabilities of event detection and localization.

### D. Semi-supervised Classification

|                      | Heart        | Wine quality | Wine  | Iris  | Breast       | Car          |
|----------------------|--------------|--------------|-------|-------|--------------|--------------|
| # of samples ( $n$ ) | 303          | 1599         | 178   | 150   | 569          | 1728         |
| # of classes ( $K$ ) | 2            | 6            | 3     | 3     | 2            | 4            |
| $k = 0$              |              |              |       |       |              |              |
| $\ell_1$             | 0.148        | 0.346        | 0.038 | 0.036 | 0.042        | 0.172        |
| SCAD                 | 0.148        | <i>0.353</i> | 0.038 | 0.033 | 0.042        | <b>0.149</b> |
| p-value              | 1.           | <i>0.06</i>  | 1.    | 0.27  | 1.           | <b>0.06</b>  |
| MCP                  | 0.144        | 0.351        | 0.037 | 0.035 | 0.040        | <b>0.148</b> |
| p-value              | 0.23         | 0.18         | 0.34  | 0.34  | 0.35         | <b>0.05</b>  |
| $k = 1$              |              |              |       |       |              |              |
| $\ell_1$             | 0.143        | 0.351        | 0.034 | 0.039 | 0.035        | 0.104        |
| SCAD                 | 0.144        | 0.350        | 0.034 | 0.039 | 0.035        | 0.104        |
| p-value              | 0.30         | 0.43         | 0.34  | 1.    | 0.71         | 0.66         |
| MCP                  | <i>0.146</i> | 0.350        | 0.034 | 0.039 | <b>0.034</b> | 0.103        |
| p-value              | <i>0.05</i>  | 0.44         | 0.34  | 1.    | <b>0.02</b>  | 0.23         |

TABLE IV: Misclassification rates averaged over 10 trials, with p-values from running sampled t-tests between SCAD/MCP misclassification rates and the corresponding rates using  $\ell_1$ . Cases where non-convex penalties perform better than  $\ell_1$  with p-value below 0.1 are highlighted in **bold**, and where they perform worse are in *italic*.

Graph-based learning provides a flexible and attractive way to model data in semi-supervised classification problems when vast amounts of unlabeled data are available compared to labeled data, and labels are expensive to acquire [11], [12], [55]. One can construct a nearest-neighbor graph based on the similarities between each pair of samples, and hope to propagate the label information from labeled samples to unlabeled ones. We move beyond our original problem in (5) to a  $K$ -class classification problem in a semi-supervised learning setting, where for a given dataset with  $n$  samples, we observe a subset of the one-hot encoded class labels,  $\mathbf{Y} \in \mathbb{R}^{n \times K}$ , such that  $Y_{ij} = 1$  if  $i$ th sample has been observed to be in  $j$ th class, and  $Y_{ij} = 0$  otherwise. A diagonal indicator matrix  $\mathbf{M} \in \mathbb{R}^{n \times n}$  denotes samples whose class labels have been observed. Then, we can define the modified absorption problem [12], [13], [55] using a variation of GTF to estimate the unknown class probabilities  $\mathbf{B} \in \mathbb{R}^{n \times K}$ :

$$\begin{aligned} \tilde{\mathbf{B}} = \operatorname{argmin}_{\mathbf{B} \in \mathbb{R}^{n \times K}} & \frac{1}{2} \|\mathbf{M}(\mathbf{Y} - \mathbf{B})\|_{\text{F}}^2 \\ & + \sum_{j=1}^K g(\Delta^{(k+1)} \mathbf{B}_{:j}; \lambda, \gamma) + \epsilon \|\mathbf{R} - \mathbf{B}\|_{\text{F}}^2, \end{aligned} \quad (22)$$

where  $\mathbf{R} \in \mathbb{R}^{n \times K}$  (set to be uniform in the experiment) is a fixed prior belief, and  $\epsilon > 0$  determines how much emphasis to be given to the prior belief. The labels  $\tilde{\mathbf{Y}}$  can be estimated using  $\tilde{\mathbf{B}}$  such that  $\tilde{Y}_{ij} = 1$  if and only if  $j = \operatorname{argmax}_{1 \leq \ell \leq K} \tilde{B}_{i\ell}$ , and otherwise  $\tilde{Y}_{ij} = 0$ . Note that this

can be completely separated into  $K$  scalar-GTF problems, one corresponding to each class.

We applied the algorithm in (22) to 6 popular UCI classification datasets [72] with  $\epsilon = 0.01$ . For each dataset, we normalized each feature to have zero mean and unit variance, and constructed a 5-nearest-neighbor graph of the samples based on the Euclidean distance between their features, with edge weights from the Gaussian radial basis kernel. We observed the labels of 20% of samples in each class randomly. Table IV shows the misclassification rates averaged over 10 repetitions, which demonstrates that the performance using non-convex penalties such as SCAD/MCP are at least competitive with, and often better than, those with the  $\ell_1$  penalty.

## VIII. CONCLUSIONS

We presented a framework for denoising piecewise smooth signals on graphs that generalizes the graph trend filtering framework to handle vector-valued signals using a family of non-convex regularizers. We provided theoretical guarantees on the error rates of our framework, and derived a general ADMM-based algorithm to solve this generalized graph trend filtering problem. Furthermore, we demonstrated the superior performance of these non-convex regularizers in terms of reconstruction error, bias reduction, and support recovery on both synthetic and real-world data. In particular, its performance on semi-supervised classification is investigated. In future work, we plan to further study this approach when the graph signals are observed by indirect and incomplete measurements.

## REFERENCES

- [1] R. Varma, H. Lee, Y. Chi, and J. Kovačević, "Improving graph trend filtering with non-convex penalties," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5391–5395.
- [2] M. Newman, *Networks*. Oxford University Press, 2018.
- [3] A. Sandryhaila and J. M. Moura, "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, 2013.
- [4] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, 2013.
- [5] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [6] S. Chen, R. Varma, A. Singh, and J. Kovačević, "Signal representations on graphs: Tools and applications," *arXiv preprint arXiv:1512.05406*, 2015.
- [7] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *IEEE Transactions on Signal Processing*, vol. 63, no. 24, pp. 6510–6523, 2015.
- [8] S. Chen, A. Sandryhaila, J. M. Moura, and J. Kovačević, "Signal recovery on graphs: Variation minimization," *IEEE Transactions on Signal Processing*, vol. 63, no. 17, pp. 4609–4624, 2015.
- [9] D. Romero, M. Ma, and G. B. Giannakis, "Kernel-based reconstruction of graph signals," *IEEE Transactions on Signal Processing*, vol. 65, no. 3, pp. 764–778, 2016.
- [10] A. Elmoataz, O. Lezoray, and S. Boughleux, "Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing," *IEEE Transactions on Image Processing*, vol. 17, no. 7, pp. 1047–1060, 2008.
- [11] M. Belkin, I. Matveeva, and P. Niyogi, "Regularization and semi-supervised learning on large graphs," in *International Conference on Computational Learning Theory*. Springer, 2004, pp. 624–638.
- [12] X. Zhu, Z. Ghahramani, and J. D. Lafferty, "Semi-supervised learning using gaussian fields and harmonic functions," in *Proceedings of the 20th International Conference on Machine Learning*. AAAI Press, 2003, pp. 912–919.
- [13] Y.-X. Wang, J. Sharpnack, A. J. Smola, and R. J. Tibshirani, "Trend filtering on graphs," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 3651–3691, 2016.
- [14] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky, " $\ell_1$  Trend Filtering," *SIAM Review*, vol. 51, no. 2, pp. 339–360, 2009.
- [15] P. Bühlmann and S. van de Geer, *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Science & Business Media, 2011.
- [16] J. Fan and R. Li, "Variable selection via nonconcave penalized likelihood and its oracle properties," *Journal of the American Statistical Association*, vol. 96, no. 456, pp. 1348–1360, 2001.
- [17] C.-H. Zhang, "Nearly unbiased variable selection under minimax concave penalty," *The Annals of Statistics*, vol. 38, no. 2, pp. 894–942, 2010.
- [18] P.-L. Loh and M. J. Wainwright, "Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 559–616.
- [19] P.-L. Loh, "Statistical consistency and asymptotic normality for high-dimensional robust M-estimators," *The Annals of Statistics*, vol. 45, no. 2, pp. 866–896, 2017.
- [20] C.-H. Zhang and T. Zhang, "A general theory of concave regularization for high-dimensional sparse estimation problems," *Statistical Science*, vol. 27, no. 4, pp. 576–593, 2012.
- [21] P. Breheny and J. Huang, "Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection," *The Annals of Applied Statistics*, vol. 5, no. 1, pp. 232–253, 2011.
- [22] S. Ma and J. Huang, "A concave pairwise fusion approach to subgroup analysis," *Journal of the American Statistical Association*, vol. 112, no. 517, pp. 410–423, 2017.
- [23] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," *IEEE Transactions on Signal Processing*, vol. 62, no. 11, pp. 2879–2893, 2014.
- [24] D. Hallac, J. Leskovec, and S. Boyd, "Network lasso: Clustering and optimization in large graphs," in *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 387–396.
- [25] A. Jung, A. O. Hero III, A. Mara, and S. Jahromi, "Semi-supervised learning via sparse label propagation," *arXiv preprint arXiv:1612.01414*, 2016.
- [26] A. Jung, N. Tran, and A. Mara, "When is network lasso accurate?" *Frontiers in Applied Mathematics and Statistics*, vol. 3, pp. 1–11, 2018.
- [27] N. Tran, S. Basirian, and A. Jung, "When is network lasso accurate: The vector case," *arXiv preprint arXiv:1710.03942*, 2017.
- [28] E. Mammen and S. van de Geer, "Locally adaptive regression splines," *The Annals of Statistics*, vol. 25, no. 1, pp. 387–413, 1997.
- [29] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: nonlinear phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [30] T. F. Chan, S. Osher, and J. Shen, "The digital TV filter and nonlinear denoising," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 231–241, 2001.
- [31] F. Mahmood, N. Shahid, U. Skoglund, and P. Vandergheynst, "Adaptive graph-based total variation for tomographic reconstructions," *IEEE Signal Processing Letters*, vol. 25, no. 5, pp. 700–704, 2018.
- [32] R. J. Tibshirani, "Adaptive piecewise polynomial estimation via trend filtering," *The Annals of Statistics*, vol. 42, no. 1, pp. 285–323, 2014.
- [33] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused lasso," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 1, pp. 91–108, 2005.
- [34] J. Sharpnack, A. Singh, and A. Rinaldo, "Sparsity of the edge lasso over graphs," in *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, vol. 22. PMLR, 2012, pp. 1028–1036.
- [35] Z. Harchaoui and C. Lvy-Leduc, "Multiple change-point estimation with a total variation penalty," *Journal of the American Statistical Association*, vol. 105, no. 492, pp. 1480–1493, 2010.
- [36] J.-C. Hütter and P. Rigollet, "Optimal rates for total variation denoising," in *29th Annual Conference on Learning Theory*, vol. 49. PMLR, 2016, pp. 1115–1146.

- [37] A. S. Dalalyan, M. Hebiri, and J. Lederer, “On the prediction performance of the lasso,” *Bernoulli*, vol. 23, no. 1, pp. 552–581, 2017.
- [38] K. Lin, J. Sharpnack, A. Rinaldo, and R. J. Tibshirani, “Approximate Recovery in Changepoint Problems, from  $\ell_2$  Estimation Error Rates,” *arXiv preprint arXiv:1606.06746*, 2016.
- [39] A. Jung and M. Hulsebos, “The network nullspace property for compressed sensing of big data over networks,” *Frontiers in Applied Mathematics and Statistics*, vol. 4, p. 9, 2018.
- [40] S. A. van de Geer and P. Bühlmann, “On the conditions used to prove oracle results for the lasso,” *Electronic Journal of Statistics*, vol. 3, pp. 1360–1392, 2009.
- [41] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2477–2488, 2005.
- [42] J. Chen and X. Huo, “Theoretical results on sparse representations of multiple-measurement vectors,” *IEEE Transactions on Signal Processing*, vol. 54, no. 12, pp. 4634–4643, 2006.
- [43] Y. C. Eldar, P. Kuppinger, and H. Bolcskei, “Block-sparse signals: Uncertainty relations and efficient recovery,” *IEEE Transactions on Signal Processing*, vol. 58, no. 6, pp. 3042–3054, 2010.
- [44] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, “Online Learning for Matrix Factorization and Sparse Coding,” *Journal of Machine Learning Research*, vol. 11, pp. 19–60, Mar. 2010.
- [45] Y. Chen, N. M. Nasrabadi, and T. D. Tran, “Hyperspectral Image Classification Using Dictionary-Based Sparse Representation,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 10, pp. 3973–3985, Oct. 2011.
- [46] Y. Li and Y. Chi, “Off-the-Grid Line Spectrum Denoising and Estimation With Multiple Measurement Vectors,” *IEEE Transactions on Signal Processing*, vol. 64, no. 5, pp. 1257–1269, Mar. 2016.
- [47] Y. C. Eldar and M. Mishali, “Robust Recovery of Signals From a Structured Union of Subspaces,” *IEEE Transactions on Information Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [48] M. E. Davies and Y. C. Eldar, “Rank Awareness in Joint Sparse Recovery,” *IEEE Transactions on Information Theory*, vol. 58, no. 2, pp. 1135–1146, Feb. 2012.
- [49] P.-L. Loh and M. J. Wainwright, “Support recovery without incoherence: A case for nonconvex regularization,” *The Annals of Statistics*, vol. 45, no. 6, pp. 2455–2482, 2017.
- [50] L. Chen and Y. Gu, “The convergence guarantees of a non-convex approach for sparse recovery,” *IEEE Transactions on Signal Processing*, vol. 62, no. 15, pp. 3754–3767, 2014.
- [51] R. Chartrand and V. Staneva, “Restricted isometry properties and non-convex compressive sensing,” *Inverse Problems*, vol. 24, no. 3, pp. 20–35, 2008.
- [52] K. Ji, J. Tan, Y. Chi, and J. Xu, “Learning latent features with pairwise penalties in matrix completion,” *arXiv preprint arXiv:1802.05821*, 2018.
- [53] M. Belkin and P. Niyogi, “Laplacian eigenmaps and spectral techniques for embedding and clustering,” in *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*. MIT Press, 2001, pp. 585–591.
- [54] —, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, 2003.
- [55] P. P. Talukdar and K. Crammer, “New regularized algorithms for transductive learning,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2009, pp. 442–457.
- [56] R. J. Tibshirani, *The Solution Path of the Generalized Lasso*. Stanford University, 2011.
- [57] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, pp. 267–288, 1996.
- [58] C.-H. Zhang and J. Huang, “The sparsity and bias of the lasso selection in high-dimensional linear regression,” *The Annals of Statistics*, vol. 36, no. 4, pp. 1567–1594, 2008.
- [59] F. Chung and M. Radcliffe, “On the spectra of general random graphs,” *The Electronic Journal of Combinatorics*, vol. 18, no. 1, p. 215, 2011.
- [60] F. R. Chung and F. C. Graham, *Spectral Graph Theory*. American Mathematical Society, 1997.
- [61] A. Blum, J. Hopcroft, and R. Kannan, “Foundations of data science,” *Vorabversion eines Lehrbuchs*, 2016.
- [62] A. Lubotzky, R. Phillips, and P. Sarnak, “Ramanujan graphs,” *Combinatorica*, vol. 8, no. 3, pp. 261–277, 1988.
- [63] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [64] J. Huang, P. Breheny, and S. Ma, “A selective review of group selection in high-dimensional models,” *Statistical science: a review journal of the Institute of Mathematical Statistics*, vol. 27, no. 4, 2012.
- [65] J. Bergstra, D. Yamins, and D. D. Cox, “Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures,” in *Proceedings of the 30th International Conference on Machine Learning*. JMLR.org, 2013, pp. I–115–I–123.
- [66] M. Defferrard, L. Martin, R. Pena, and N. Perraudin, “PyGSP: Graph Signal Processing in Python.” [Online]. Available: <https://github.com/epfl-lts2/pygsp/>
- [67] A. Hagberg, P. Swart, and D. S. Chult, “Exploring network structure, dynamics, and function using NetworkX,” in *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, 2008, pp. 11–16.
- [68] G. Boeing, “OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks,” *Computers, Environment and Urban Systems*, vol. 65, pp. 126–139, 2017.
- [69] Taxi and Limousine Commission (TLC), “2011 yellow taxi trip data.” [Online]. Available: <https://data.cityofnewyork.us/Transportation/2011-Yellow-Taxi-Trip-Data/uwyp-dntv>
- [70] Socrata API, “2011 yellow taxi trip data.” [Online]. Available: <https://dev.socrata.com/foundry/data.cityofnewyork.us/uwyp-dntv>
- [71] “2011 Pride Guide.” [Online]. Available: [https://issuu.com/nycpride/docs/prideguide\\_9\\_rev2\\_low-res](https://issuu.com/nycpride/docs/prideguide_9_rev2_low-res)
- [72] D. Dua and C. Graff, “UCI machine learning repository,” 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [73] M. J. Wainwright, *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019, vol. 48.
- [74] P. Tseng, “Convergence of a block coordinate descent method for nondifferentiable minimization,” *Journal of Optimization Theory and Applications*, vol. 109, no. 3, pp. 475–494, 2001.

## APPENDIX

### A. Proof of Theorem 1

*Proof.* We denote  $\mathbf{D} = \mathbf{\Delta}^{(k+1)}$ . Define  $R$  as the row space of  $\mathbf{D}$ , and  $R^\perp$  the null space. Let  $\mathcal{P}_R = \mathbf{D}^\dagger \mathbf{D}$ , the projection onto  $R$ , and  $\|\mathbf{x}\|_R = \|\mathcal{P}_R \mathbf{x}\|_2$ . Additionally,  $\mathcal{P}_{R^\perp} = \mathbf{I} - \mathbf{D}^\dagger \mathbf{D}$ , the projection onto  $R^\perp$ . Since  $\hat{\boldsymbol{\beta}}$  is a stationary point of  $f(\boldsymbol{\beta})$ , it follows that

$$\mathbf{0} \in \nabla_{\boldsymbol{\beta}} f(\boldsymbol{\beta})|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}} = (\hat{\boldsymbol{\beta}} - \mathbf{y}) + \nabla_{\boldsymbol{\beta}} g(\mathbf{D}\boldsymbol{\beta})|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}}. \quad (23)$$

By the chain rule,  $\nabla_{\boldsymbol{\beta}} g(\mathbf{D}\boldsymbol{\beta})|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}} = \{\mathbf{D}^\top \mathbf{z} : \mathbf{z} \in \nabla_{\mathbf{x}} g(\mathbf{x})|_{\mathbf{x}=\mathbf{D}\hat{\boldsymbol{\beta}}}\}$ . Then by (23), there exists  $\mathbf{z} \in \nabla_{\mathbf{x}} g(\mathbf{x})|_{\mathbf{x}=\mathbf{D}\hat{\boldsymbol{\beta}}}$ , such that

$$\mathbf{0} = (\hat{\boldsymbol{\beta}} - \mathbf{y}) + \mathbf{D}^\top \mathbf{z}.$$

In particular,  $\forall \boldsymbol{\beta} \in \mathbb{R}^n$ , we have

$$\boldsymbol{\beta}^\top (\mathbf{y} - \hat{\boldsymbol{\beta}}) = (\mathbf{D}\boldsymbol{\beta})^\top \mathbf{z}, \quad (24)$$

and, specializing to  $\hat{\boldsymbol{\beta}}$ ,

$$\hat{\boldsymbol{\beta}}^\top (\mathbf{y} - \hat{\boldsymbol{\beta}}) = (\mathbf{D}\hat{\boldsymbol{\beta}})^\top \mathbf{z}. \quad (25)$$

Subtract (25) from (24), and use the definition of subgradient to get  $\forall \boldsymbol{\beta} \in \mathbb{R}^n$ ,

$$\begin{aligned} \boldsymbol{\beta}^\top (\mathbf{y} - \hat{\boldsymbol{\beta}}) - \hat{\boldsymbol{\beta}}^\top (\mathbf{y} - \hat{\boldsymbol{\beta}}) &= (\mathbf{D}\boldsymbol{\beta} - \mathbf{D}\hat{\boldsymbol{\beta}})^\top \mathbf{z} \\ &\leq g(\mathbf{D}\boldsymbol{\beta}) - g(\mathbf{D}\hat{\boldsymbol{\beta}}). \end{aligned} \quad (26)$$

By the measurement model  $\mathbf{y} = \boldsymbol{\beta}^* + \boldsymbol{\epsilon}$  and the polarization equality, i.e.  $2\mathbf{a}^\top \mathbf{b} = \|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2 - \|\mathbf{a} - \mathbf{b}\|_2^2$ , the left-hand

side of (26) can be rewritten as

$$\begin{aligned}
& \boldsymbol{\beta}^\top (\mathbf{y} - \widehat{\boldsymbol{\beta}}) - \widehat{\boldsymbol{\beta}}^\top (\mathbf{y} - \widehat{\boldsymbol{\beta}}) \\
&= (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}})^\top (\boldsymbol{\beta}^* - \widehat{\boldsymbol{\beta}}) + \boldsymbol{\epsilon}^\top (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}) \\
&= \frac{1}{2} \|\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}\|_2^2 + \frac{1}{2} \|\boldsymbol{\beta}^* - \widehat{\boldsymbol{\beta}}\|_2^2 - \frac{1}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + \boldsymbol{\epsilon}^\top (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}).
\end{aligned} \tag{27}$$

Combining (26) and (27) gives us  $\forall \boldsymbol{\beta} \in \mathbb{R}^n$

$$\begin{aligned}
& \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 \\
& \leq \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + 2\boldsymbol{\epsilon}^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + 2g(\mathbf{D}\boldsymbol{\beta}) - 2g(\mathbf{D}\widehat{\boldsymbol{\beta}}).
\end{aligned} \tag{28}$$

Let us first consider  $\boldsymbol{\epsilon}^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ . From Hölder's inequality,

$$\begin{aligned}
\boldsymbol{\epsilon}^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) &= (\mathbf{D}^\dagger \mathbf{D}\boldsymbol{\epsilon})^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + (\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon})^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \\
&\leq \|(\mathbf{D}^\dagger)^\top \boldsymbol{\epsilon}\|_\infty \|\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})\|_1 + \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2.
\end{aligned} \tag{29}$$

By standard tail bounds for independent Gaussian random variables, we have with probability at least  $1 - \delta$ ,

$$\|(\mathbf{D}^\dagger)^\top \boldsymbol{\epsilon}\|_\infty \leq \sigma \zeta_k \sqrt{2 \log\left(\frac{er}{\delta}\right)}. \tag{30}$$

Additionally, recognize that  $\|\boldsymbol{\epsilon}\|_{\mathbb{R}^\perp}^2$  is a chi-squared random variable with  $C_G$  degrees of freedom. We can then invoke the one-sided tail bound for chi-squared random variables (c.f. [73, Example 2.5]) such that for any  $0 \leq t \leq 1$ ,

$$P(\|\boldsymbol{\epsilon}\|_{\mathbb{R}^\perp}^2 \geq \sigma^2 C_G (1+t)) \leq \exp\left(-\frac{C_G t^2}{8}\right).$$

Consequently, with probability at least  $1 - \delta$ ,

$$\|\boldsymbol{\epsilon}\|_{\mathbb{R}^\perp}^2 \leq \sigma^2 \left(C_G + 2\sqrt{2C_G \log(1/\delta)}\right). \tag{31}$$

The inequalities (31) and (30) hold simultaneously with probability at least  $1 - 2\delta$ . Then, using  $\lambda \|\mathbf{x}\|_1 \leq g(\mathbf{x}) + \frac{\mu}{2} \|\mathbf{x}\|_2^2$  and  $\lambda = \sigma \zeta_k \sqrt{2 \log\left(\frac{er}{\delta}\right)} \geq \|(\mathbf{D}^\dagger)^\top \boldsymbol{\epsilon}\|_\infty$ , we can bound (29) further as

$$\begin{aligned}
\boldsymbol{\epsilon}^\top (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) &\leq \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 + \lambda \|\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})\|_1 \\
&\leq \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 + g(\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})) + \frac{\mu}{2} \|\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})\|_2^2.
\end{aligned}$$

Together with  $\|\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})\|_2^2 \leq \|\mathbf{D}\|_2^2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2$ , we can upper bound (28) as

$$\begin{aligned}
& \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 \\
& \leq \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + 2\|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 + \mu \|\mathbf{D}\|_2^2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 \\
& \quad + 2g(\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})) + 2g(\mathbf{D}\boldsymbol{\beta}) - 2g(\mathbf{D}\widehat{\boldsymbol{\beta}}).
\end{aligned} \tag{32}$$

Note that for any set  $T$ ,  $g(\mathbf{x}) = \sum_{i \in T} \rho(x_i) + \sum_{j \in T^c} \rho(x_j) = g((\mathbf{x})_T) + g((\mathbf{x})_{T^c})$ . Therefore, using the triangle inequality and subadditivity and symmetry of  $\rho$ ,

$$\begin{aligned}
& g(\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})) + g(\mathbf{D}\boldsymbol{\beta}) - g(\mathbf{D}\widehat{\boldsymbol{\beta}}) \\
& \leq g((\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}))_T) + g((\mathbf{D}\boldsymbol{\beta})_{T^c}) + g((\mathbf{D}\widehat{\boldsymbol{\beta}})_{T^c}) \\
& \quad + g(\mathbf{D}\boldsymbol{\beta}) - g((\mathbf{D}\widehat{\boldsymbol{\beta}})_T) - g((\mathbf{D}\widehat{\boldsymbol{\beta}})_{T^c}) \\
& = g((\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}))_T) + 2g((\mathbf{D}\boldsymbol{\beta})_{T^c}) + g((\mathbf{D}\boldsymbol{\beta})_T) - g((\mathbf{D}\widehat{\boldsymbol{\beta}})_T) \\
& \leq 2g((\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}))_T) + 2g((\mathbf{D}\boldsymbol{\beta})_{T^c}).
\end{aligned} \tag{33}$$

We bound (33) further by the compatibility factor,

$$\begin{aligned}
g((\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}))_T) &\leq \lambda \|(\mathbf{D}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}))_T\|_1 \\
&\leq \lambda \sqrt{|T|} \kappa_T^{-1} \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2.
\end{aligned} \tag{34}$$

Now combining (32), (33), and (34), we then have

$$\begin{aligned}
\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 &\leq \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + 4g((\mathbf{D}\boldsymbol{\beta})_{T^c}) \\
& \quad + 2 \left( \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 + 2\lambda \sqrt{|T|} \kappa_T^{-1} \right) \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \\
& \quad + \mu \|\mathbf{D}\|_2^2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2.
\end{aligned}$$

Apply Young's inequality, which is  $2ab \leq a^2/\epsilon + \epsilon b^2$  for  $\epsilon > 0$ , with  $a = \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 + 2\lambda \sqrt{|T|} \kappa_T^{-1}$ ,  $b = \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2$ , and  $\epsilon = 1 - \mu \|\mathbf{D}\|_2^2 > 0$ , we have

$$\begin{aligned}
& 2 \left( \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 + 2\lambda \sqrt{|T|} \kappa_T^{-1} \right) \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \\
& \leq \frac{1}{\epsilon} \left( \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2 + 2\lambda \sqrt{|T|} \kappa_T^{-1} \right)^2 + \epsilon \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 \\
& \leq \frac{2}{(1 - \mu \|\mathbf{D}\|_2^2)} \left( \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2^2 + 4\lambda^2 |T| \kappa_T^{-2} \right) \\
& \quad + (1 - \mu \|\mathbf{D}\|_2^2) \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2.
\end{aligned} \tag{35}$$

Therefore,

$$\begin{aligned}
& \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 \\
& \leq \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + 4g((\mathbf{D}\boldsymbol{\beta})_{T^c}) + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2 \\
& \quad + \frac{2}{(1 - \mu \|\mathbf{D}\|_2^2)} \left( \|\mathcal{P}_{\mathbb{R}^\perp} \boldsymbol{\epsilon}\|_2^2 + 4\lambda^2 |T| \kappa_T^{-2} \right).
\end{aligned} \tag{36}$$

Cancel  $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2^2$  on both sides, apply the infimum over  $\boldsymbol{\beta}$  and plug in the bounds (31) to get

$$\begin{aligned}
\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 &\leq \inf_{\boldsymbol{\beta}} \left\{ \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + 4g((\mathbf{D}\boldsymbol{\beta})_{T^c}) \right\} \\
& \quad + \frac{2\sigma^2}{(1 - \mu \|\mathbf{D}\|_2^2)} \left( C_G + 2\sqrt{2C_G \log\left(\frac{1}{\delta}\right)} + \frac{8\zeta_k^2 |T|}{\kappa_T^2} \log\left(\frac{er}{\delta}\right) \right).
\end{aligned}$$

The proof extends for the vector-GTF (10) in a similar manner. We need to replace (29) by

$$\begin{aligned}
\langle \mathbf{E}, \widehat{\mathbf{B}} - \mathbf{B} \rangle &= \langle \mathbf{D}^\dagger \mathbf{D}\mathbf{E}, \widehat{\mathbf{B}} - \mathbf{B} \rangle + \langle \mathcal{P}_{\mathbb{R}^\perp} \mathbf{E}, \widehat{\mathbf{B}} - \mathbf{B} \rangle \\
&\leq \lambda \sum_{\ell=1}^r \|\mathbf{D}_\ell(\widehat{\mathbf{B}} - \mathbf{B})\|_2 + \|\mathcal{P}_{\mathbb{R}^\perp} \mathbf{E}\|_{\mathbb{F}} \|\widehat{\mathbf{B}} - \mathbf{B}\|_{\mathbb{F}},
\end{aligned}$$

where  $\|\mathcal{P}_{\mathbb{R}^\perp} \mathbf{E}\|_{\mathbb{F}}^2 \leq d\sigma^2 \left(C_G + 2\sqrt{2C_G \log(d/\delta)}\right)$  with probability at least  $1 - \delta$ . Similarly, for (34), we use the generalized definition of the compatibility factor  $\kappa_T$ , given as

$$\begin{aligned}
h((\mathbf{D}(\widehat{\mathbf{B}} - \mathbf{B}))_T) &\leq \lambda \sum_{\ell \in T} \|(\mathbf{D}(\widehat{\mathbf{B}} - \mathbf{B}))_\ell\|_2 \\
&\leq \lambda \sqrt{|T|} \kappa_T^{-1} \|\widehat{\mathbf{B}} - \mathbf{B}\|_{\mathbb{F}},
\end{aligned}$$

which will lead to the claimed bound in the theorem.  $\square$

## B. Proof of Proposition 1

*Proof.* By Cauchy-Schwartz inequality, we have

$$\sum_{\ell \in T} \|(\boldsymbol{\Delta}^{(k+1)} \mathbf{B})_\ell\|_2 \leq \sqrt{|T|} \|(\boldsymbol{\Delta}^{(k+1)} \mathbf{B})_T\|_{\mathbb{F}},$$

and note that given two matrices  $\mathbf{U}$  and  $\mathbf{V}$ ,  $(\mathbf{UV})_T = (\mathbf{U})_T \mathbf{V}$  where  $T$  is a subset of rows indices. We also use the fact that  $\|\mathbf{UV}\|_F \leq \|\mathbf{U}\| \|\mathbf{V}\|_F$ . We consider two cases:

- For even  $k$ , we have

$$\begin{aligned} \|(\Delta^{(k+1)} \mathbf{B})_T\|_F &= \|(\Delta)_T \Delta^{(k)} \mathbf{B}\|_F \\ &\leq \|(\Delta)_T\| \|\Delta^{(k)} \mathbf{B}\|_F = \sqrt{\lambda_{\max}((\Delta)_T^\top (\Delta)_T)} \|\Delta^{(k)} \mathbf{B}\|_F. \end{aligned}$$

Note that  $(\Delta)_T$  is equivalent to the incidence matrix of a subgraph with only  $T$  edges, which allows us to bound,

$$\lambda_{\max}((\Delta)_T^\top (\Delta)_T) \leq \max_{(u,v) \in T} \{d_u + d_v\} \leq 2d_{\max}$$

where  $d_i$  is the degree of node  $i$ .

- For odd  $k$ , we have

$$\begin{aligned} \|(\Delta^{(k+1)} \mathbf{B})_T\|_F &= \|(\Delta^\top)_T \Delta^{(k)} \mathbf{B}\|_F \\ &\leq \|(\Delta^\top)_T\| \|\Delta^{(k)} \mathbf{B}\|_F = \sqrt{\lambda_{\max}(\Delta_{T \times T}^{(2)})} \|\Delta^{(k)} \mathbf{B}\|_F, \end{aligned}$$

where  $\Delta_{T \times T}^{(2)} \in \mathbb{R}^{|T| \times |T|}$  is the principal submatrix of  $\Delta^{(2)}$  indexed by  $T$ . By Cauchy's interlacing theorem, the maximum eigenvalue of the submatrix is upper bounded, so

$$\lambda_{\max}(\Delta_{T \times T}^{(2)}) \leq \lambda_{\max}(\Delta^{(2)}) \leq 2d_{\max}.$$

Therefore, for all  $k$ ,  $\|(\Delta^{(k+1)} \mathbf{B})_T\|_F \leq \sqrt{2d_{\max}} \|\Delta^{(k)} \mathbf{B}\|_F$ . To conclude the proof, note

$$\begin{aligned} \sum_{\ell \in T} \|(\Delta^{(k+1)} \mathbf{B})_\ell\|_2 &\leq \sqrt{|T|} \sqrt{2d_{\max}} \|\Delta^{(k)} \mathbf{B}\|_F \\ &\leq \sqrt{|T|} \sqrt{2d_{\max}} \|\Delta^{(k)}\| \|\mathbf{B}\|_F \leq (2d_{\max})^{\frac{k+1}{2}} \sqrt{|T|} \|\mathbf{B}\|_F. \end{aligned}$$

□

### C. Proof of Theorem 3

We show the convergence of Alg. 1 by proving a modified version of [22, Proposition 1]. The superscript  $(m)$  denotes the values of  $\mathbf{B}, \mathbf{Z}, \mathbf{U}$  at the  $m$ th iteration of the loop inside Alg. 1.

**Proposition 2** (Convergence to a feasible solution). *If  $\tau \geq \mu$ , then the primal residual  $r^{(m)} = \|\Delta^{(k+1)} \mathbf{B}^{(m)} - \mathbf{Z}^{(m)}\|_F$  and the dual residual  $s^{(m+1)} = \|\tau(\Delta^{(k+1)})^\top (\mathbf{Z}^{(m+1)} - \mathbf{Z}^{(m)})\|_F$  of Alg. 1 satisfy that  $\lim_{m \rightarrow \infty} r^{(m)} = 0$  and  $\lim_{m \rightarrow \infty} s^{(m)} = 0$ .*

*Proof.* Denote  $\mathbf{D} = \Delta^{(k+1)}$ , and  $\rho_\lambda(\cdot) = \rho(\cdot; \lambda, \gamma)$ . Recall from Assumption 1 (c) that there exists  $\mu > 0$  such that  $\rho_\lambda(\|\mathbf{x}\|_2) + \frac{\mu}{2} \|\mathbf{x}\|_2^2$  is convex. Now consider the Lagrangian  $\mathcal{L}(\mathbf{B}, \mathbf{Z}, \mathbf{U})$  with regard to the  $\ell$ -th row  $\mathbf{z}_\ell$  of  $\mathbf{Z} = [\mathbf{z}_1^\top, \dots, \mathbf{z}_r^\top]^\top$ , assuming all other variables are fixed:

$$\begin{aligned} &\rho_\lambda(\|\mathbf{z}_\ell\|_2) + \frac{\tau}{2} \|\mathbf{z}_\ell - \mathbf{c}_1\|_2^2 + c_2 \\ &= \rho_\lambda(\|\mathbf{z}_\ell\|_2) + \frac{\tau}{2} \|\mathbf{z}_\ell\|_2^2 - \tau \langle \mathbf{z}_\ell, \mathbf{c}_1 \rangle + \frac{\tau}{2} \|\mathbf{c}_1\|_2^2 + c_2 \end{aligned}$$

where  $\mathbf{c}_1$  and  $c_2$  represent terms of  $\mathcal{L}(\mathbf{B}, \mathbf{Z}, \mathbf{U})$  that do not depend on  $\mathbf{z}_\ell$ . With our choice of  $\tau \geq \mu$ , then  $\mathcal{L}(\mathbf{B}, \mathbf{Z}, \mathbf{U})$  is convex with regard to each of  $\mathbf{B}, \mathbf{U}$ , and for each row of  $\mathbf{Z}$ , allowing us to apply [74, Theorem 5.1]. Therefore, Alg. 1 converges to limit points  $\mathbf{B}^*, \mathbf{Z}^*, \mathbf{U}^*$ .

Then it follows that the dual residual  $\lim_{m \rightarrow \infty} s^{(m)} = \|\tau \mathbf{D}^\top (\mathbf{Z}^* - \mathbf{Z}^*)\|_F = 0$ . For the primal residual, notice that the  $\mathbf{U}$  update step in line 10 of Alg. 1 also shows that  $\forall m, t \geq 0$ ,

$$\mathbf{U}^{(m+t)} = \mathbf{U}^{(m)} + \sum_{i=1}^t (\mathbf{D} \mathbf{B}^{(m+i)} - \mathbf{Z}^{(m+i)}).$$

Fixing  $t$  and setting  $m \rightarrow \infty$ , we have

$$\mathbf{U}^* = \mathbf{U}^* + t(\mathbf{D} \mathbf{B}^* - \mathbf{Z}^*)$$

holds  $\forall t \geq 0$ . Hence,  $\mathbf{D} \mathbf{B}^* - \mathbf{Z}^* = \mathbf{0}$ , and therefore  $\lim_{m \rightarrow \infty} r^{(m)} = \|\mathbf{D} \mathbf{B}^* - \mathbf{Z}^*\|_F = 0$ . □

This proposition shows that the algorithm in the limit achieves primal and dual feasibility, and that the Augmented Lagrangian in (21) with  $\mathbf{Z}^*$  and  $\mathbf{U}^*$  becomes the original GTF formulation in (10).  $\mathbf{B}$  that is produced by every iteration of Alg. 1 is a stationary point of (21) with fixed  $\mathbf{Z}$  and  $\mathbf{U}$ . As a result,  $\mathbf{B}^*$  is a stationary point of (10).



**Rohan Varma** is a Ph.D. student in Electrical and Computer Engineering from Carnegie Mellon University and is advised by Jelena Kovačević. He received a M.Sc in Electrical and Computer Engineering in 2016 from Carnegie Mellon University and a B.Sc in Electrical Engineering and Computer Science, B.A Economics and B.A Statistics in 2014 from U.C Berkeley. Varma was the recipient of the 2019 IEEE Signal Processing Society Young Author Best Paper Award. His research interests include graph signal processing, graph neural networks and more generally using graphs to augment machine learning tasks.




**Harlin Lee** is a Ph.D. candidate in Electrical and Computer Engineering at Carnegie Mellon University. She received her M.Eng. (with concentration in AI) and B.S. degrees in 2017 and 2016 respectively, both in Electrical Engineering and Computer Science from MIT. Her research interests include graph regularization, unsupervised learning, and clinical data analysis.



**Jelena Kovačević** (S'88–M'91–SM'96–F'02) received the Dipl. Electr. Eng. degree from the EE Department, University of Belgrade, Yugoslavia, in 1986, and the M.S. and Ph.D. degrees from Columbia University, New York, in 1988 and 1991, respectively. From 1991–2002, she was with Bell Labs, Murray Hill, NJ. She was a co-founder and Technical VP of xWaveforms, based in New York City and an Adjunct Professor at Columbia University. In 2003, she joined Carnegie Mellon University, where she was Hamerschlag University Professor

and Head of Electrical and Computer Engineering, Professor of Biomedical Engineering, and was the Director of the Center for Bioimage Informatics at Carnegie Mellon University. Since 2018, she is the William R. Berkley Professor and Dean of the Tandon School of Engineering at New York University, New York City. Her research interests include wavelets, frames, graphs, and applications to bioimaging and smart infrastructure.

Dr. Kovačević coauthored the books *Wavelets and Subband Coding* (Prentice Hall, 1995) and *Foundations of Signal Processing* (Cambridge University Press, 2014), a top-10 cited paper in the *Journal of Applied and Computational Harmonic Analysis*, and the paper for which A. Mojsilović received the Young Author Best Paper Award. Her paper on multidimensional filter banks and wavelets was selected as one of the Fundamental Papers in Wavelet Theory. She received the Belgrade October Prize in 1986, the E.I. Jury Award at Columbia University in 1991, and the 2010 CIT Philip L. Dowd Fellowship Award from the College of Engineering at Carnegie Mellon University and the 2016 IEEE SPS Technical Achievement Award. She is a past Editor-in-Chief of the *IEEE Transactions on Image Processing*, served as a guest co-editor on a number of special issues and is/was on the editorial boards of several journals. She was a regular member of the NIH Microscopic Imaging Study Section and served as a Member-at-Large of the IEEE Signal Processing Society Board of Governors. She is a past Chair of the IEEE Signal Processing Society Bio Imaging and Signal Processing Technical Committee.



**Yuejie Chi** (S'09–M'12–SM'17) received the Ph.D. degree in Electrical Engineering from Princeton University in 2012, and the B.E. (Hon.) degree in Electrical Engineering from Tsinghua University, Beijing, China, in 2007. She was with The Ohio State University from 2012 to 2017. Since 2018, she is an Associate Professor with the department of Electrical and Computer Engineering at Carnegie Mellon University, where she holds the Robert E. Doherty Early Career Professorship. Her research interests include signal processing, statistical inference, machine learning, large-scale optimization, and their applications in data science, inverse problems, imaging, and sensing systems.

She is a recipient of the PECASE Award, NSF CAREER Award, AFOSR and ONR Young Investigator Program Awards, Ralph E. Powe Junior Faculty Enhancement Award, Google Faculty Research Award, IEEE Signal Processing Society Young Author Best Paper Award and the Best Paper Award at the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). She has served as an Elected Member of the SPTM, SAM and MLSP Technical Committees of the IEEE Signal Processing Society. She currently serves as an Associate Editor of *IEEE Trans. on Signal Processing*.