
LogSpecT: Feasible Graph Learning Model from Stationary Signals with Recovery Guarantees

Shangyuan Liu

The Chinese University of Hong Kong
shangyuanliu@link.cuhk.edu.hk

Linglingzhi Zhu

The Chinese University of Hong Kong
llzzhu@se.cuhk.edu.hk

Anthony Man-Cho So

The Chinese University of Hong Kong
manchoso@se.cuhk.edu.hk

Abstract

Graph learning from signals is a core task in graph signal processing (GSP). A significant subclass of graph signals called the stationary graph signals that broadens the concept of stationarity of data defined on regular domains to signals on graphs is gaining increasing popularity in the GSP community [10, 20, 25]. The most commonly used model to learn graphs from these stationary signals is *SpecT* [31], which forms the foundation for nearly all the subsequent, more advanced models. Despite its strengths, the practical formulation of the model, known as *rSpecT*, has been identified to be susceptible to the choice of hyperparameters. More critically, it may suffer from infeasibility as an optimization problem. In this paper, we introduce the first condition that ensures the infeasibility of *rSpecT* and design a novel model called *LogSpecT*, along with its practical formulation *rLogSpecT* to overcome this issue. Contrary to *rSpecT*, our novel practical model *rLogSpecT* is always feasible. Furthermore, we provide recovery guarantees of *rLogSpecT* from modern optimization tools related to epi-convergence, which could be of independent interest and significant for various learning problems. To demonstrate the practical advantages of *rLogSpecT*, a highly efficient algorithm based on the linearized alternating direction method of multipliers (L-ADMM) that allows closed-form solutions for each subproblem is proposed with convergence guarantees. Extensive numerical results on both synthetic and real networks not only corroborate the stability of our proposed methods, but also highlight their comparable and even superior performance than existing models.

1 Introduction

Learning with graphs has proved its relevance in many practical areas, such as life science [36, 37], signal processing [13, 14, 38, 39], and financial engineering [1, 21], to just name a few. However, there are many cases that the graphs are not readily prepared and only the data closely related to the graphs can be observed. Hence, a core task in Graph Signal Processing (GSP) is to learn the underlying graph topology based on the interplay between data and graphs [22].

Among these commonly used data properties, an assumption that is referred to as the graph signal stationarity [10] is gaining attention [30, 20, 25]. Graph signal stationarity is the extension of the notion of stationarity from the regular space/time domain to the irregular domain (i.e. graphs). The latter is a traditional hypothesis in signal processing to capture a special type of statistical relationship between samples of a temporal signal. It is a cornerstone of many signal processing methods. Although the graph signal stationarity property is proposed more from a theoretic end, the

techniques based on this idealistic assumption pave the way for the development of the following more advanced and applicable models [35, 34]. Furthermore, several works have shown that some real datasets can be approximately viewed as stationary or partially explained by the stationarity assumption. For instance, [25] revealed that the well-known USPS dataset and the CMUPIE set of cropped faces exhibit near stationarity. [10] found that certain weather data could be explained by stationary graph signals. Additionally, [31] highlighted the relevance of the stationarity assumption in learning protein structure. [43] applied the stationary assumption to US Senate Roll Calls and achieved good performance.

The predominant methods to process stationary graph signals and learn topology under the stationarity assumption are the spectral-template-based models. The start of this line of works is [31], which proposed a vanilla model called *SpecT* [31] to learn graphs from stationary graph signals. Many extensions of this fundamental model have been made since then to accommodate the more realistic and complicated settings, e.g. joint inference of multiple graphs [32, 27], learning with streaming graph signals [33], learning with hidden nodes [4], learning from both smooth and stationary graph signals [5]. In practice, *SpecT* requires the unknown data covariance matrix. Hence, a robust formulation called *rSpecT* is proposed [31], which replaces the unknown covariance matrix with its estimate and introduces a hyperparameter to reflect the estimation inaccuracy. Unfortunately, the model is sensitive to this hyperparameter and improper tuning of it may jeopardize the model performance or lead to model infeasibility.

The current approach to selecting an appropriate value is heuristic. It finds the smallest one that allows for a feasible solution. This method has two shortcomings. On one hand, It is computationally costly. The smallest value is the exact solution to a second-order conic programming, which is solved approximately with computational methods. Any approximate solution that is smaller than the exact one will make *rSpecT* infeasible. On the other hand, such an approach lacks theoretical guarantees as how the smallest value behaves as the number of observations increases is unknown. This issue makes the existing recovery guarantees invalid. To the best of our knowledge, only a few works have focused on how to understand the performance of the robust formulation *rSpecT*. [24] empirically showed that the feasible set of the robust formulation approximates the original one as more samples are collected. [23] provided the first recovery guarantee for *rSpecT*. Their work exploits the fact that the objective is the ℓ_1 -norm and the constraints are linear [44]. However, the conditions needed in their work are not only restrictive but also hard to check, as they require the full rankness of a large-scale matrix related to random samples. Another model [34] tries to circumvent the model infeasibility issue by turning a constraint into a penalty. However, this approach introduces another hyperparameter that is neither easy to tune nor amenable to theoretical analysis. Also, this hyperparameter should go to infinity heuristically when more and more samples are collected, which may exacerbate model instability.

In this paper, we propose an alternative formulation of *SpecT* that can bypass the aforementioned issues. We first provide a condition that ensures the infeasibility of the fundamental model *rSpecT*. Then, we propose a novel formulation called *LogSpecT* and its practical formulation *rLogSpecT* based on the spectral template and incorporating a log barrier. The inclusion of a log barrier as a regularizer on the degrees is inspired by its extensive application in graph learning from smooth graph signals [15, 16, 19, 47] and the discussion in [7] emphasizing that restricting the degrees to preclude isolated nodes is crucial to learning graphs with special structures. It is worth noting that the approach in [7], which directly imposes constraints, may exacerbate the infeasibility issue in our problem. Although *rLogSpecT* still requires a parameter to reflect the estimation error of the unknown covariance matrix, it no longer suffers from the infeasibility issue associated with the improper choice of the parameter. Furthermore, we provide the decaying rate of the parameter that can guarantee the recovery of *rLogSpecT*. This not only aids in the parameter tuning, but also establishes a connection between *rLogSpecT* and *LogSpecT* through non-asymptotic recovery analysis. These theoretical results are derived using contemporary optimization tools related to epi-convergence [29]. Unlike the current guarantees for *rSpecT* that are built on an ℓ_1 analysis model [44], our approach based on epi-convergence is not confined to the specific types of optimization problems (e.g. the combination of the ℓ_1 -norm and log barrier in the objective function) and consequently admits broader applications that can be of independent interest.

In the algorithmic aspect, we design a linearized alternating direction method of multipliers (L-ADMM) to solve *rLogSpecT*. The subproblems of L-ADMM admit closed-form solutions and can be implemented efficiently due to the linearization technique. Also, we provide the convergence result of the proposed method. Finally, we conduct extensive experiments on both synthetic data and

real networks. Notably, the infeasibility issue of rSpecT frequently occurs in experiments on real networks, thereby emphasizing its practical significance beyond theoretical concerns. Furthermore, we observe that even when the classical models are set to be feasible, our novel models (LogSpecT and rLogSpecT) exhibit improved accuracy and stability, particularly in the case of synthetic BA graphs. This empirically illustrates that although our models are initially proposed to circumvent the infeasibility issue, they are superior alternatives to classical methods in the general cases.

1.1 Notation

The notation we use in this paper is standard. We use $[m]$ to denote the set $\{1, 2, \dots, m\}$ for any positive integer m . Let the Euclidean space of all real matrices be equipped with the inner product $\langle \mathbf{X}, \mathbf{Y} \rangle := \text{tr}(\mathbf{X}^\top \mathbf{Y})$ for any matrices \mathbf{X}, \mathbf{Y} and denote the induced Frobenius norm by $\|\cdot\|_F$ (or $\|\cdot\|_2$ when the argument is a vector). For any $\mathbf{X} \in \mathbb{R}^{m \times n}$, we use $\mathbb{B}(\mathbf{X}, \rho)$ to denote the closed Euclidean ball centering at \mathbf{X} with radius ρ . Let $\|\cdot\|$ be the operator norm, $\|\mathbf{X}\|_{1,1} := \sum_{i,j} |\mathbf{X}_{ij}|$, $\|\mathbf{X}\|_{\infty,\infty} := \max_{i,j} |\mathbf{X}_{ij}|$, and let $\text{diag}(\mathbf{X})$ be vector formed by the diagonal entries of \mathbf{X} . For a column vector \mathbf{x} , let $\text{Diag}(\mathbf{x})$ be the diagonal matrix whose diagonal elements are given by \mathbf{x} . Given a closed and convex set \mathcal{C} , we use $\Pi_{\mathcal{C}}(\mathbf{w}) := \text{argmin}_{\mathbf{v} \in \mathcal{C}} \|\mathbf{v} - \mathbf{w}\|_2$ to denote the projection of the point \mathbf{w} onto the set \mathcal{C} . We use $\mathbf{1}$ (resp. $\mathbf{0}$) to denote an all-one vector (resp. all-zero vector) whose dimension will be clear from the context. For a set \mathcal{S} and any real number α , let $\alpha\mathcal{S} := \{\alpha x \mid x \in \mathcal{S}\}$, $\mathcal{S}\mathbf{1} := \{\mathbf{x}\mathbf{1} \mid \mathbf{x} \in \mathcal{S}\}$, and $\iota_{\mathcal{S}}(\cdot)$ be the indicator function of \mathcal{S} . For two non-empty and compact sets \mathcal{X} and \mathcal{Y} , the distance between them is defined as $\text{dist}(\mathcal{X}, \mathcal{Y}) := \sup_{\mathbf{X} \in \mathcal{X}} \inf_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{X} - \mathbf{Y}\|_F$ ¹.

2 Preliminaries

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph, where $\mathcal{V} = [m]$ is the set of nodes and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges. Let \mathbf{S} be the weight matrix associated with the graph \mathcal{G} , where \mathbf{S}_{ij} represents the weight of the edge between nodes i and j . In this paper, we consider undirected graphs without self-loops. The set of adjacency matrices for such graphs is $\mathcal{S} := \{\mathbf{S} \in \mathbb{R}^{m \times m} \mid \mathbf{S} = \mathbf{S}^\top, \text{diag}(\mathbf{S}) = \mathbf{0}, \mathbf{S} \geq \mathbf{0}\}$. Suppose that the adjacency matrix \mathbf{S} admits the eigen-decomposition $\mathbf{S} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, where $\mathbf{\Lambda}$ is a diagonal matrix and \mathbf{U} is an orthogonal matrix. A graph filter is a linear operator $h : \mathbb{R}^{m \times m} \rightarrow \mathbb{R}^{m \times m}$ defined as $h(\mathbf{S}) = \sum_{i=0}^p h_i \mathbf{S}^i = \mathbf{U} \left(\sum_{i=0}^p h_i \mathbf{\Lambda}^i \right) \mathbf{U}^\top$, where $p > 0$ is the order of the graph filter and $\{h_i\}_{i=0}^p$ are the filter coefficients. According to the convention, we have $\mathbf{S}^0 = \mathbf{I}$.

A graph signal can be represented by a vector $\mathbf{x} \in \mathbb{R}^m$, where the i -th element x_i is the signal value associated with node i . A signal \mathbf{x} is said to be stationary if it is generated from

$$\mathbf{x} = h(\mathbf{S})\mathbf{w}, \quad (1)$$

where \mathbf{w} satisfies $\mathbb{E}[\mathbf{w}] = \mathbf{0}$ and $\mathbb{E}[\mathbf{w}\mathbf{w}^\top] = \mathbf{I}$. Simple calculations give that the covariance matrix of \mathbf{x} , which is denoted by \mathbf{C}_∞ , shares the same eigenvectors with \mathbf{S} . Hence, we have the constraint

$$\mathbf{C}_\infty \mathbf{S} = \mathbf{S} \mathbf{C}_\infty. \quad (2)$$

Based on this, the following fundamental model SpecT is proposed to learn graphs from stationary signals without the knowledge of graph filters [31, 34]:

$$\min_{\mathbf{S}} \|\mathbf{S}\|_{1,1}, \text{ s.t. } \mathbf{C}_\infty \mathbf{S} = \mathbf{S} \mathbf{C}_\infty, \mathbf{S} \in \mathcal{S} \cap \left\{ \mathbf{S} \in \mathbb{R}^{m \times m} : \sum_{j=1}^m \mathbf{S}_{1j} = \mathbf{1} \right\}, \quad (\text{SpecT})$$

where the constraint $\sum_{j=1}^m \mathbf{S}_{1j} = \mathbf{1}$ is used to preclude the trivial optimal solution $\mathbf{S}^* = \mathbf{0}$. When \mathbf{C}_∞ is unknown and only n i.i.d samples of $\{\mathbf{x}_i\}$ from (1) are available, the robust formulation, which is based on the estimate \mathbf{C}_n of \mathbf{C}_∞ and called rSpecT, is used:

$$\min_{\mathbf{S}} \|\mathbf{S}\|_{1,1}, \text{ s.t. } \|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F \leq \delta, \mathbf{S} \in \mathcal{S} \cap \left\{ \mathbf{S} \in \mathbb{R}^{m \times m} : \sum_{j=1}^m \mathbf{S}_{1j} = \mathbf{1} \right\}. \quad (\text{rSpecT})$$

For this robust formulation, the recovery guarantee is studied empirically in [24] and theoretically in [23] under some conditions that are restrictive and hard to check.

¹It reduces to the classic point-to-set (resp. point-to-point) distance when \mathcal{Y} (resp. \mathcal{X} and \mathcal{Y}) is a singleton.

3 Infeasibility of rSpecT and Novel Models

Even though rSpecT has gained much popularity, there are few works that discuss the choice of the hyperparameter δ and how it affects the model feasibility. In this section, we present a condition under which rSpecT is always infeasible and then propose an alternative formulation. To motivate our results, let us consider the following 2-node example.

Example 3.1. Consider a simple graph containing 2 nodes. Then, the set given by the second constraint of rSpecT is $\mathcal{S} \cap \{\mathbf{S} \in \mathbb{R}^{m \times m} : \sum_{j=1}^m \mathbf{S}_{1j} = 1\} = [0, 1; 1, 0]$, which is a singleton. Suppose that the sample covariance matrix is $\mathbf{C}_n = [h_{11}, h_{12}; h_{12}, h_{22}]$. Then, the constraint $\|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F \leq \delta$ is reduced to $2(h_{11} - h_{22})^2 \leq \delta^2$. Hence, when $h_{11} \neq h_{22}$ and $\delta < \sqrt{2}|h_{11} - h_{22}|$, rSpecT has no feasible solution.

Before delving into the general case, we introduce a linear operator $\mathbf{B} \in \mathbb{R}^{m^2 \times m(m-1)/2}$ that maps the vector $\mathbf{x} \in \mathbb{R}_+^{(m-1)m/2}$ to the vectorization of an adjacency matrix $\mathbf{S} \in \mathcal{S}$ of a simple, undirected graph:

$$(\mathbf{B}\mathbf{x})_{m(i-1)+j} = \begin{cases} x_{i-j+\frac{i-1}{2}(2m-j)}, & \text{if } i > j, \\ 0, & \text{if } i = j, \\ x_{j-i+\frac{i-1}{2}(2m-i)}, & \text{if } i < j, \end{cases} \quad (3)$$

where $i, j \in [m]$. We also define $\mathbf{A}_n := \mathbf{I} \otimes \mathbf{C}_n - \mathbf{C}_n \otimes \mathbf{I} \in \mathbb{R}^{m^2 \times m^2}$, so that the first constraint of rSpecT can be rewritten as $\|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F \leq \delta \iff \|\mathbf{A}_n \text{vec}(\mathbf{S})\|_2 \leq \delta$, where $\text{vec}(\cdot)$ is the vectorization operator. We now give a condition that guarantees the infeasibility of rSpecT.

Theorem 3.2. Consider the linear system

$$\mathbf{A}_n \mathbf{B} \mathbf{y} = \mathbf{0}, \quad \mathbf{y} \leq \mathbf{0}, \quad \sum_{i=1}^{m-1} y_i \neq \mathbf{0}, \quad (4)$$

where $\mathbf{y} \in \mathbb{R}^{\frac{m(m-1)}{2}}$. If (4) has no feasible solution, then there exists a $\bar{\delta}_n > 0$ such that rSpecT is infeasible for all $\delta_n \in [0, \bar{\delta}_n]$.

Remark 3.3. From Theorem 3.2 we can infer that for any fixed n , the linear system (4) has no solution when $\mathbf{A}_n \mathbf{B}$ has full column rank (e.g., Example 3.1). This leads to the infeasibility of rSpecT with $\delta_n \in [0, \bar{\delta}_n]$. As remarked in [31], one should tackle the feasibility issue of rSpecT with caution.

The failure of rSpecT (SpecT) lies in the existence of the constraint $(\mathbf{S}\mathbf{1})_1 = 1$, which is used to preclude the trivial solution $\mathbf{S}^* = \mathbf{0}$. For this reason, we resort to an alternative approach to bypassing the trivial solution. When the graphs are assumed to have no isolated node, the log barrier is commonly applied [15]. In these cases, the zero solution is naturally precluded. This observation inspires us to propose the following novel formulation, which combines the log barrier with the spectral template (2) to learn graphs without isolated nodes from stationary signals:

$$\begin{aligned} \min_{\mathbf{S}} \quad & \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) \\ \text{s.t.} \quad & \mathbf{S} \in \mathcal{S}, \quad \mathbf{S} \mathbf{C}_\infty = \mathbf{C}_\infty \mathbf{S}, \end{aligned} \quad (\text{LogSpecT})$$

where $\|\cdot\|_{1,1}$ is a convex relaxation of $\|\cdot\|_0$ promoting graph sparsity, $\mathbf{1}^\top \log(\mathbf{S}\mathbf{1})$ is the penalty to guarantee the nonexistence of isolated nodes, and α is the tuning parameter. As can be seen from the following proposition, the hyperparameter α in LogSpecT only affects the scale of edge weights instead of the graph connectivity structure.

Proposition 3.4. Let $\text{Opt}(\mathbf{C}, \alpha)$ be the optimal solution set of LogSpecT with input covariance matrix \mathbf{C} and parameter $\alpha > 0$. Then, for any $\gamma > 0$, it follows that

$$\text{Opt}(\mathbf{C}, \alpha) = \gamma \text{Opt}(\mathbf{C}, \alpha/\gamma) = \alpha \text{Opt}(\mathbf{C}, 1).$$

Remark 3.5. The result of Proposition 3.4 spares us from tuning the hyperparameter α when we are coping with binary graphs. In fact, certain normalization will eliminate the impact of different values of α and preserve the connectivity information. Hence, we may simply set $\alpha = 1$ in implementation.

Note that the true covariance matrix \mathbf{C}_∞ in LogSpecT is usually unknown and an estimate \mathbf{C}_n from n i.i.d samples is available. To tackle the estimation inaccuracy, we introduce the following robust

formulation:

$$\begin{aligned} \min_{\mathbf{S}} \quad & \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) \\ \text{s.t.} \quad & \mathbf{S} \in \mathcal{S}, \quad \|\mathbf{S}\mathbf{C}_n - \mathbf{C}_n\mathbf{S}\|_F^2 \leq \delta_n^2. \end{aligned} \tag{rLogSpecT}$$

This formulation substitutes \mathbf{C}_∞ with \mathbf{C}_n and relaxes the equality constraint to an inequality constraint with a tolerance threshold δ_n . Contrary to rSpecT, we prove that rLogSpecT is always feasible.

Proposition 3.6. *For $\delta_n > 0$ and \mathbf{C}_n with any fixed n , rLogSpecT always has a nontrivial feasible solution.*

4 Recovery Guarantee of rLogSpecT

In this section, we investigate the non-asymptotic behavior of rLogSpecT when more and more i.i.d samples are collected. For the sake of brevity, we denote $f := \|\cdot\|_{1,1} - \alpha \mathbf{1}^\top \log(\cdot \mathbf{1})$. The theoretical recovery guarantee is as follows:

Theorem 4.1. *If $\delta_n \geq 2\rho\|\mathbf{C}_n - \mathbf{C}_\infty\|$ with given $\rho := \max\{f^*, \alpha m, \alpha m(\log \alpha - 1)\}$, then there exist constants $c_1, \tilde{c}_1, \tilde{c}_2 > 0$ such that*

$$(i) |f_n^* - f^*| \leq \varepsilon_n; \quad (ii) \text{dist}(\mathcal{S}^{n,*}, \mathcal{S}_{2\varepsilon_n}^*) \leq \varepsilon_n; \quad (iii) \text{dist}(\mathcal{S}^{n,*}\mathbf{1}, \mathcal{S}_0^*\mathbf{1}) \leq \tilde{c}_1\varepsilon_n + \tilde{c}_2\sqrt{\varepsilon_n},$$

where $\varepsilon_n := c_1(\delta_n + 2\rho\|\mathbf{C}_n - \mathbf{C}_\infty\|)$, $\mathcal{S}^{n,*}$ is the optimal solution set of rLogSpecT, f^* (resp. f_n^*) denotes the optimal value of LogSpecT (resp. rLogSpecT), and

$$\mathcal{S}_\varepsilon^* := \{\mathbf{S} \in \mathcal{S} \mid \mathbf{S}\mathbf{C}_\infty = \mathbf{C}_\infty\mathbf{S}, f(\mathbf{S}) \leq f^* + \varepsilon\}$$

is the ε -suboptimal solution set of LogSpecT.

Remark 4.2. (i) Compared with the conclusion (ii) in the Theorem 4.1, the conclusion (iii) links the optimal solution sets of rLogSpecT and LogSpecT instead of the sub-optimal solutions.

(ii) A byproduct of the proof (see Appendix E.2 for details) shows that the optimal node degree vector set ($\mathcal{S}_0^*\mathbf{1}$) is a singleton. However, there is no guarantee of the uniqueness of \mathbf{S}^* itself. We will discuss the impact of such non-uniqueness on model performance in Section 6.4.

Remark 4.3. *The proof of Theorem 4.1 relies on two important optimization concepts: Epi-convergence and truncated Hausdorff distance. Epi-convergence is closely related to the asymptotic solution behavior of approximate minimization problems, and the truncated Hausdorff distance is used to characterize the epi-convergence non-asymptotically. With the help of the Kenmochi condition, which allows us to explicitly calculate the truncated Hausdorff distance, we are able to study the non-asymptotic behavior of the optimal value and optimal solutions of the models. We refer the reader to Appendix E and also [28, Chapter 7], [29, Chapter 6.J] for more details.*

Corollary 4.4. *Under the assumptions in Theorem 4.1, it follows that if $\delta_n \rightarrow 0$, then*

$$\lim_{n \rightarrow \infty} \text{dist}(\mathcal{S}^{n,*}, \mathcal{S}_0^*) = 0. \tag{5}$$

Remark 4.5. *From the strong law of large numbers, if we choose $\delta_n = \mathcal{O}(\|\mathbf{C}_n - \mathbf{C}_\infty\|)$, then almost surely one has $\lim_{n \rightarrow \infty} \|\mathbf{C}_n - \mathbf{C}_\infty\| = 0$ and consequently $\delta_n \rightarrow 0$. Hence, Corollary 4.4 shows that (5) holds almost surely.*

For the remaining part of this section, we study the choice of δ_n under certain statistical assumptions. A large number of distributions (e.g., Gaussian distributions, exponential distributions and any bounded distributions) can be covered by the sub-Gaussian distribution (cf. [41, Section 2.5]), whose formal definition is as follows.

Definition 4.6 (Sub-Gaussian Distributions). *The probability distribution of a random vector \mathbf{w} is called sub-Gaussian if there are $C, v > 0$ such that for every $t > 0$, $\mathbb{P}(\|\mathbf{w}\|_2 > t) \leq Ce^{-vt^2}$.*

Consider the case that \mathbf{w} in the generative model (1) follows a sub-Gaussian distribution. The following result is adapted from [40, Proposition 2.1].

Lemma 4.7. *Suppose that \mathbf{w} in the generative model (1) follows a sub-Gaussian distribution. Then, \mathbf{x} follows a sub-Gaussian distribution, and with probability larger than $1 - \mathcal{O}(\frac{1}{n})$, one has $\|\mathbf{C}_n - \mathbf{C}_\infty\| \leq \mathcal{O}(\sqrt{\log n/n})$.*

Equipped with non-asymptotic results in Theorem 4.1, we choose δ_n with a specific decaying rate.

Corollary 4.8. *If the input w in (1) follows a sub-Gaussian distribution and $\delta_n = \mathcal{O}(\sqrt{\log n/n})$, then the assumptions in Theorem 4.1 hold with probability larger than $1 - \mathcal{O}(\frac{1}{n})$.*

Corollary 4.8 together with Theorem 4.1 illustrates the non-asymptotic convergence of optimal function value/suboptimal solution set/optimal node degree vector of rLogSpecT from n i.i.d samples to the ideal model LogSpecT with the convergence rate $\mathcal{O}(\varepsilon_n)/\mathcal{O}(\varepsilon_n)/\mathcal{O}(\sqrt{\varepsilon_n})$, where $\varepsilon_n = \mathcal{O}(|\mathbf{C}_n - \mathbf{C}_\infty|) \leq \mathcal{O}(\sqrt{\log n/n})$ for sub-Gaussian ω . This matches the convergence rate $\mathcal{O}(\sqrt{\log n/n})$ of classic spectral template models, e.g., Proposition 2 in [31] and Theorem 2 in [23] for the SpecT model, which shows that LogSpecT is also competitive on recovery rate.

5 Linearized ADMM for Solving rLogSpecT

In this section, we design a linearized ADMM algorithm to solve rLogSpecT that admits closed-form solutions for the subproblems. The ADMM-type algorithms have been successfully applied to tackle various graph learning tasks [46, 42].

Firstly, we reformulate rLogSpecT such that it fits the ADMM scheme:

$$\begin{aligned} \min_{\mathbf{S}, \mathbf{Z}, \mathbf{q}} \quad & \langle \mathbf{S}, \mathbf{1}\mathbf{1}^\top \rangle - \alpha \mathbf{1}^\top \log \mathbf{q} + \iota_{\mathcal{S}}(\mathbf{S}) + \iota_{\mathbb{B}(\mathbf{0}, \delta_n)}(\mathbf{Z}) \\ \text{s.t.} \quad & \mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n = \mathbf{Z}, \quad \mathbf{q} = \mathbf{S}\mathbf{1}. \end{aligned} \quad (6)$$

The augmented Lagrangian function of problem (6) is

$$\begin{aligned} L(\mathbf{S}, \mathbf{Z}, \mathbf{q}, \boldsymbol{\Lambda}, \boldsymbol{\lambda}_2) = & \langle \mathbf{1}\mathbf{1}^\top, \mathbf{S} \rangle - \alpha \mathbf{1}^\top \log \mathbf{q} + \iota_{\mathcal{S}}(\mathbf{S}) + \iota_{\mathbb{B}(\mathbf{0}, \delta_n)}(\mathbf{Z}) + \langle \boldsymbol{\Lambda}, \mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n - \mathbf{Z} \rangle \\ & + \boldsymbol{\lambda}_2^\top (\mathbf{q} - \mathbf{S}\mathbf{1}) + \frac{\rho}{2} \|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n - \mathbf{Z}\|_F^2 + \frac{\rho}{2} \|\mathbf{q} - \mathbf{S}\mathbf{1}\|_2^2. \end{aligned}$$

Update of primal variables \mathbf{Z} and \mathbf{q} Since \mathbf{Z} and \mathbf{q} are separable given \mathbf{S} , we update the primal variables in two blocks: \mathbf{S} and (\mathbf{Z}, \mathbf{q}) . More specifically, the update of (\mathbf{Z}, \mathbf{q}) in the k -th iteration is

$$(\mathbf{Z}^{(k+1)}, \mathbf{q}^{(k+1)}) := \operatorname{argmin}_{(\mathbf{Z}, \mathbf{q})} L(\mathbf{S}^{(k)}, \mathbf{Z}, \mathbf{q}, \boldsymbol{\Lambda}^{(k)}, \boldsymbol{\lambda}_2^{(k)}), \quad (7)$$

which admits the following closed-form solution:

Proposition 5.1. *The update (7) can be explicitly rewritten as*

$$\mathbf{Z}^{(k+1)} = \min \left\{ 1, \frac{\delta_n}{\|\tilde{\mathbf{Z}}\|_F} \right\} \tilde{\mathbf{Z}}, \quad \mathbf{q}^{(k+1)} = \frac{\tilde{\mathbf{q}} + \sqrt{\tilde{\mathbf{q}}^2 + 4\alpha/\rho\mathbf{1}}}{2},$$

where $\tilde{\mathbf{Z}} = \mathbf{C}_n \mathbf{S}^{(k)} - \mathbf{S}^{(k)} \mathbf{C}_n + \frac{\boldsymbol{\Lambda}^{(k)}}{\rho}$, $\tilde{\mathbf{q}} = \mathbf{S}^{(k)} \mathbf{1} - \frac{1}{\rho} \boldsymbol{\lambda}_2^{(k)}$.

Update of primal variables \mathbf{S} The ADMM update of \mathbf{S} does not admit a closed-form solution. Hence, we solve the linearized version \hat{L}_k of the augmented Lagrangian function L when all the variables except \mathbf{S} are fixed. It is given by

$$\hat{L}_k(\mathbf{S}) = \langle \mathbf{C}^{(k)}, \mathbf{S} \rangle + \frac{\rho\tau}{2} \|\mathbf{S} - \mathbf{S}^{(k)}\|_F^2 + \iota_{\mathcal{S}}(\mathbf{S}),$$

where $\mathbf{C}^{(k)} := \mathbf{1}\mathbf{1}^\top + \mathbf{C}_n \boldsymbol{\Lambda}^{(k)} - \boldsymbol{\Lambda}^{(k)} \mathbf{C}_n - \boldsymbol{\lambda}_2^{(k)} \mathbf{1}^\top + \rho(\mathbf{S}^{(k)} \mathbf{C}_n^2 + \mathbf{C}_n^2 \mathbf{S}^{(k)} - 2\mathbf{C}_n \mathbf{S}^{(k)} \mathbf{C}_n + \mathbf{Z}^{(k+1)} \mathbf{C}_n - \mathbf{C}_n \mathbf{Z}^{(k+1)} - \mathbf{q}^{(k+1)} \mathbf{1}^\top + \mathbf{S}^{(k)} \mathbf{1}\mathbf{1}^\top)$ and τ is the parameter for linearization. Therefore, the update of \mathbf{S} in the k -th iteration is

$$\mathbf{S}^{(k+1)} := \operatorname{argmin}_{\mathbf{S} \in \mathcal{S}} \hat{L}_k(\mathbf{S}) = \Pi_{\mathcal{S}}(\mathbf{S}^{(k)} - \mathbf{C}^{(k)}/\rho\tau), \quad (8)$$

where $\Pi_{\mathcal{S}}(\mathbf{X})$ is the projection of \mathbf{X} onto the closed and convex set \mathcal{S} and can be computed by

$$(\Pi_{\mathcal{S}}(\mathbf{X}))_{ij} = \begin{cases} \frac{1}{2} \max\{0, X_{ij} + X_{ji}\}, & \text{if } i \neq j, \\ 0, & \text{if } i = j. \end{cases}$$

Update of dual variables and parameters Finally, we update the dual variables Λ and λ_2 via gradient ascent steps of the augmented Lagrangian function:

$$\Lambda^{(k+1)} := \Lambda^{(k)} + \rho(C_n \mathbf{S}^{(k+1)} - \mathbf{S}^{(k+1)} C_n - \mathbf{Z}^{(k+1)}), \quad (9)$$

$$\lambda_2^{(k+1)} := \lambda_2^{(k)} + \rho(\mathbf{q}^{(k+1)} - \mathbf{S}^{(k+1)} \mathbf{1}). \quad (10)$$

Together with the updating rule for the augmented parameter ρ , we complete the whole algorithm. The details are presented in Appendix F.

Complexity of each iteration To illustrate the efficiency of L-ADMM, we calculate the complexity of each subproblem. In the update of \mathbf{Z} and \mathbf{q} , the most time consuming part is to calculate $C_n \mathbf{S}^{(k)}$. When the number of nodes m is larger than the sample size n , we rewrite $C_n \mathbf{S}^{(k)}$ as $(1/n \mathbf{X})(\mathbf{X}^\top \mathbf{S}^{(k)})$, where $\mathbf{X} \in \mathbb{R}^{m \times n}$ is the data matrix. In this case, the complexity is in the order of $\mathcal{O}(nm^2)$. When m is smaller than n , we first calculate C_n and then multiply it with $\mathbf{S}^{(k)}$. This gives the complexity $\mathcal{O}(nm^2)$. Similarly, the complexity needed to update \mathbf{S} is $\mathcal{O}(nm^2 + mn^2)$. The update of the dual variables can be decomposed as the summation of components calculated when updating the primal variables. In summary, the total complexity of each iteration is $\mathcal{O}(nm^2 + mn^2)$.

Convergence of L-ADMM For the convergence analysis of L-ADMM, we treat (\mathbf{Z}, \mathbf{q}) as one variable, and the two constraints $C_n \mathbf{S} - \mathbf{S} C_n = \mathbf{Z}$ and $\mathbf{q} = \mathbf{S} \mathbf{1}$ can be written into a single one:

$$[\mathbf{I} \otimes C_n - C_n \otimes \mathbf{I}; \mathbf{I} \otimes \mathbf{1}^\top] \text{vec}(\mathbf{S}) = [\text{vec}(\mathbf{Z}); \mathbf{q}].$$

Then we can apply a two-block proximal ADMM (i.e., L-ADMM) to the problem by alternately updating \mathbf{S} (see details in Appendix F.4 for this part) and (\mathbf{Z}, \mathbf{q}) . Consequently, the convergence result in [45, Theorem 4.2] can be invoked directly to derive the following theorem.

Theorem 5.2. *If $\tau > m + \|C_n \otimes \mathbf{I}_m - \mathbf{I}_m \otimes C_n\|^2$, then $\lim_{k \rightarrow \infty} f(\mathbf{S}^{(k)}) = f_n^*$.*

6 Experiments

In this section, we evaluate our proposed algorithm and models via numerical experiments. The experiments are conducted on both synthetic and real networks. We apply the standard metrics [18]: F-measure, Precision, and Recall to assess the quality of learned graphs. The source code is available at: <https://github.com/StevenSYL/NeurIPS2023-LogSpecT>.

6.1 Data Generation

Random Graphs. In the experiments, we consider two types of synthetic graphs, namely, the Erdős-Rényi (ER) graph [9] and the Barabasi-Albert model graph (BA) [2]. The ER graphs are generated by placing an edge between each pair of nodes independently with probability $p = 0.2$ and the weight on each edge is set to 1. The BA graphs are generated by having two connected nodes initially and then adding new nodes one at a time, where each new node is connected to exactly one previous node that is randomly chosen with a probability proportional to its degree at the time.

Graph Filters and Signals. Three graph filters are used in the experiments. The first one is the low-pass graph filter (lowpass-EXP) $h(\mathbf{S}) = \exp(\frac{1}{2} \mathbf{S})$. The second one is the high-pass graph filter (highpass-EXP) $h(\mathbf{S}) = \exp(-\mathbf{S})$. The last one is a quadratic graph filter (QUA) $h(\mathbf{S}) = \mathbf{S}^2 + \mathbf{S} + \mathbf{I}$. Note that the quadratic graph filter is neither low-pass nor high-pass. We will also use random graph filters developed from the above ones for experiments on real networks. They will be introduced when applied. The stationary graph signals are generated from the generative model (1). The input graph signal $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$ is a random vector following the normal distribution.

6.2 How to Infer Binary Graphs

Many practical tasks require binary graphs instead of weighted graphs. However, the results from LogSpec and rLogSpecT are generally weighted. In this section, we tackle the issue of how to convert a weighted graph to a binary one. Firstly, we normalize each edge in the graph \mathbf{W} by dividing it over the maximal weight, which yields a collection of values ranging from 0 to 1. Our task is then to select a threshold ε to round these values to 0 or 1 and construct a binary graph \mathbf{W}^* accordingly. Mathematically, this procedure is captured by the formula:

$$(W^*)_{ij} = \begin{cases} 1, & \text{if } W_{ij} / \max\{W_{i'j'}\} \geq \varepsilon, \\ 0, & \text{if } W_{ij} / \max\{W_{i'j'}\} < \varepsilon. \end{cases}$$

To choose the threshold ε , we either use a training-based strategy or a searching-based strategy. The training-based one searches the best ε on k graphs and applies the value to the newly learned graph. The searching-based one simply searches for ε that yields the best performance for a graph.

6.3 Convergence of L-ADMM

We present the convergence performance of L-ADMM in this section. Since there does not literally exist customized algorithm for rLogSpecT, we compare our L-ADMM with CVXPY [8]. The solver used in CVXPY is MOSEK. We conduct the experiments on 100-node BA graphs with 10000 stationary graph signals generated from low-pass EXP graph filter. The parameter δ is set as $10\sqrt{\log n/n}$. For L-ADMM algorithm, the target accuracy is set as 10^{-6} and the initialization is set as zero. The running time for L-ADMM takes around 20 seconds while the solver takes over 110 seconds. The primal residual and dual residual of L-ADMM in each iteration are plot in Figure 1. The results corroborate the efficiency of our proposed algorithm. Furthermore, the locally linear convergence can be observed from the experiment. This implies that L-ADMM converges with fast speed to high accuracy. It remains an open question to prove the property theoretically.

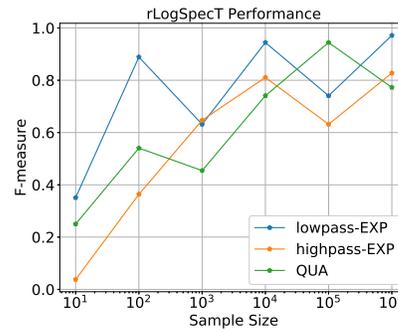
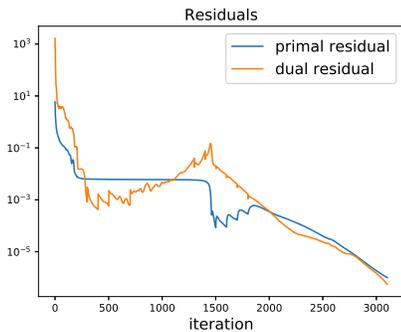


Figure 1: L-ADMM residual on 100-node graph Figure 2: rLogSpecT with $\delta_n = 0.2\sqrt{\log n/n}$

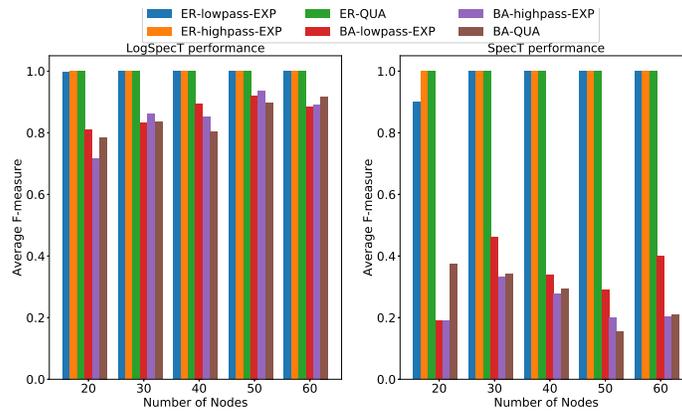


Figure 3: Comparison of LogSpecT & SpecT on Synthetic Data

6.4 Experiments on Synthetic Networks

To evaluate the efficacy of LogSpecT and rLogSpecT, we compare them with SpecT and rSpecT on synthetic data. Notice that the prior of no isolated nodes has been incorporated in practice by adding restrictions on the smallest degree. This approach introduces an additional hyperparameter and may exacerbate the infeasibility issue in robust formulation. Also, no recovery guarantees are developed for this trick. For these reasons, we still compare LogSpecT & rLogSpecT with SpecT & rSpecT.

Performance of LogSpecT. We first compare the performance of the ideal models: LogSpecT and SpecT on the two types of random graphs with three graph filters. We conduct experiments on graphs with nodes ranging from 20 to 60 and use the searching-based strategy to set threshold ε^* . After 10 repetitions, we report the average results in Figure 3. The left column presents the performance of LogSpecT and the right column presents that of SpecT. Firstly, we observe that on ER graphs, both models achieve good performance. They nearly recover the graphs perfectly. However, the performance on BA graphs differs. In this case, LogSpecT can work efficiently. The learned graphs from LogSpecT enjoy much higher F-measure values than those from SpecT. Secondly, the comparison between different graph filters shows that different graph filters have few impacts on recovery performance for both LogSpecT and SpecT. Finally, we observe that the outputs of LogSpecT on BA graphs tend to possess higher F-measure when the number of nodes increases. This suggests that LogSpecT may behave better on larger networks. However, such phenomena cannot be observed from SpecT results.

Performance of rLogSpecT. Since LogSpecT persuasively outperforms SpecT, we only present the performance of rLogSpecT when more samples are collected. As we have mentioned before, the optimal solution to LogSpecT is not necessarily unique. Thus, we do not expect to be able to show how rLogSpecT’s optimal solution converges to LogSpecT’s optimal solution. Moreover, the non-uniqueness may jeopardize the performance of rLogSpecT intuitively. However, the following experiment on BA graphs shows that the learned graphs from rLogSpecT with the binary-graph-transforming strategy tend to approach the ground truth when enough samples are collected.

In this experiment, the sample size n is chosen from 10 to 10^6 and δ_n is set as $0.2\sqrt{\log n/n}$. We rely on the training-based strategy to obtain the best threshold ε^* from 10 randomly chosen training graphs. We then calculate the F-measure of the learned binary graph in the testing set. The result is reported in Figure 2. It shows that rLogSpecT works for all three graph filters on BA graphs and tends to possess higher F-measure values when more and more signals are collected. This indicates that similar to LogSpecT, rLogSpecT is efficient for different types of graph filters on BA graphs, including the high-pass ones. For more experiments and discussions, we refer readers to Appendix G.

6.5 Experiments on Real Networks

In this set of experiments, we compare the performance of LogSpecT (resp. rLogSpecT) with SpecT (resp. rSpecT) and other methods from the statistics and GSP communities on *Protein* database and *Reddit* database from [6]. The *Protein* database is a Bioinformatics dataset, where nodes are secondary structure elements (SSEs) and there is an edge between two nodes if they are neighbors in the amino-acid sequence or in 3D space. The *Reddit* database is a social network dataset. In this dataset, each graph corresponds to an online discussion thread, where nodes correspond to users, and there is an edge between two nodes if at least one of them responds to another’s comments. We choose graphs in databases whose numbers of nodes are smaller than 50 and generate stationary graph signals from the generative model (1), with the graph filtering chosen as $H(\mathcal{S}) = \exp \frac{1}{2}\mathcal{S}$. Our experiments include in total 871 testing graphs in the *Protein* database and 309 in *Reddit* database.

Infeasibility of rSpecT. For these real networks, we first check whether the infeasibility issue encountered by rSpecT is significant. To this end, we adopt the random graph filters $t_1\mathcal{S}^2 + t_2\mathcal{S} + t_3\mathbf{I}$, where $t_i, i = 1, 2, 3$ are random variables following a Gaussian distribution with $\mu = 0$ and $\sigma = 2$. Then, we calculate the smallest δ_n such that rSpecT is feasible. If the smallest $\delta_n > 0$, rSpecT is likely to be encountered with the infeasibility issue. The results with different numbers of graph signals observed are shown in Table 1.

Table 1: Likelihood of infeasibility

	Protein			Reddit		
sample size	10	100	1000	10	100	1000
frequency	0.980	0.999	0.994	0.984	1	1
mean of δ_{\min}	38.514	26.012	14.813	1094.788	830.642	531.006

The experiment results present a decrease in the mean value of δ_{\min} when more samples are used. However, a high frequency of the infeasibility issue occurring in these real datasets cannot be mitigated by the increasing sample size. These observations highlight the necessity of the careful treatment for choosing δ in rSpecT again.

Performance of different graph learning models. In this section, the stationary graph signals are generated by the low-pass-EXP filter. We choose the famous statistical method called (thresholded) correlation [17] and the first GSP method that applies the log barrier to graph inference [15] as the baselines. The optimal threshold for the correlation method is selected from 0.1 to 0.6 and we search in $\{0.01, 0.1, 1, 10, 100, 1000\}$ to obtain the best hyperparameter in Kalofolias' model. The parameter δ_n in rLogSpecT is set as $10\sqrt{\log n/n}$ and in rSpecT it is set as the smallest value that allows a feasible solution [31]. We also rely on the searching-based strategy to convert the learned weighted graphs from (r)SpecT & (r)LogSpecT.

The results of the ideal models with the true covariance matrix C_∞ applied are collected in Figure 4. We observe that on the real graphs, LogSpecT achieves the best performance on average (median represented by the red line). Also, compared with SpecT, LogSpecT performs more stably. We remark that since the graph signals are not necessarily smooth, Kalofolias' model cannot provide guaranteed performance, especially on the Reddit networks.

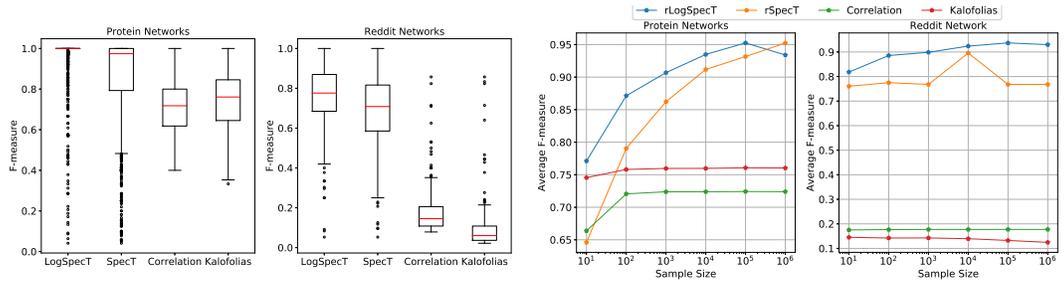


Figure 4: Performance of LogSpecT and SpecT Figure 5: Asymptotic behavior of rLogSpecT and rSpecT.

Figure 5 compares the performance of four different methods when different numbers of signals are observed.² When the sample size increases, the models focusing on stationarity property can recover the graphs more accurately while correlation method and Kalofolias' method fail to exploit the signal information. This can also be inferred from the experiment results in Figure 4 since the models fail to achieve good performance from full information, let alone from partial information. The experiment also shows that when a fixed number of samples are observed, the learned graphs from rLogSpecT approximate the ground truth better than rSpecT. This further corroborates the superiority of rLogSpecT on graph learning from stationary signals.

7 Conclusion

In this paper, we directly tackle the infeasibility issue encountered in graph learning from stationary graph signals [31] and propose the first condition that guarantees the model infeasibility. To overcome this, we propose an efficient alternative. The recovery guarantees of its robust formulation are analyzed with advanced optimization tools, which may find broader applications in learning tasks. Compared with current literature [23], these theoretical results require less stringent conditions. We also design an L-ADMM algorithm that allows for efficient implementation and theoretical convergence. Extensive experiments on both synthetic and real data are conducted in this paper. The results show that our proposed model can significantly outperform the existing ones.

We believe this work represents an important step beyond the fundamental model SpecT. Its general formulation allows for the transfer of the current extensions made on SpecT. Testing our proposed models with these extensions is one future direction. Also, we notice that although the recovery guarantees for robust formulations are clear, the estimation performance analysis for the ideal-case models (i.e. SpecT and LogSpecT) is still incomplete. Investigating the exact recovery conditions is another future direction.

²The model in [34] is a substitute of rSpecT to approximate SpecT and its hyperparameter is hard to determine. Hence, we omit the performance of that model.

Acknowledgments and Disclosure of Funding

This work was supported in part by the Hong Kong Research Grants Council (RGC) General Research Fund (GRF) Project CUHK 14203920.

References

- [1] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.
- [2] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [3] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.
- [4] Andrei Buciulea, Samuel Rey, Cristobal Cabrera, and Antonio G Marques. Network reconstruction from graph-stationary signals with hidden variables. In *Proceedings of the 53rd Asilomar Conference on Signals, Systems, and Computers*, pages 56–60. IEEE, 2019.
- [5] Andrei Buciulea, Samuel Rey, and Antonio G Marques. Learning graphs from smooth and graph-stationary signals with hidden variables. *IEEE Transactions on Signal and Information Processing over Networks*, 8:273–287, 2022.
- [6] Samir Chowdhury and Tom Needham. Generalized spectral clustering via gromov-wasserstein learning. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, pages 712–720. PMLR, 2021.
- [7] Jose Vinicius de Miranda Cardoso, Jiayi Ying, and Daniel Palomar. Graphical models in heavy-tailed markets. *Advances in Neural Information Processing Systems*, 34:19989–20001, 2021.
- [8] Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, 17(1):2909–2913, 2016.
- [9] Paul Erdős, Alfréd Rényi, et al. On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1):17–60, 1960.
- [10] Benjamin Girault. Stationary graph signals using an isometric graph translation. In *Proceedings of the 23rd European Signal Processing Conference (EUSIPCO 2015)*, pages 1516–1520. IEEE, 2015.
- [11] Yiran He and Hoi-To Wai. Detecting central nodes from low-rank excited graph signals via structured factor analysis. *IEEE Transactions on Signal Processing*, 70:2416–2430, 2022.
- [12] Alan J Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952.
- [13] Alexander Jung, Alfred O Hero III, Alexandru Cristian Mara, Saeed Jahromi, Ayelet Heimowitz, and Yonina C Eldar. Semi-supervised learning in network-structured data via total variation minimization. *IEEE Transactions on Signal Processing*, 67(24):6256–6269, 2019.
- [14] Alexander Jung and Nguyen Tran. Localized linear regression in networked data. *IEEE Signal Processing Letters*, 26(7):1090–1094, 2019.
- [15] Vassilis Kalofolias. How to learn a graph from smooth signals. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS 2016)*, pages 920–929. PMLR, 2016.
- [16] Vassilis Kalofolias, Andreas Loukas, Dorina Thanou, and Pascal Frossard. Learning time varying graphs. In *Proceedings of 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2017)*, pages 2826–2830. IEEE, 2017.

- [17] Eric D. Kolaczyk. *Statistical Analysis of Network Data: Methods and Models*. Springer-Verlag, New York, NY, USA, 2009.
- [18] D. Christopher Manning, Hinrich Schütze, and Prabhakar Raghavan. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [19] Hermina Petric Maretic and Pascal Frossard. Graph laplacian mixture model. *IEEE Transactions on Signal and Information Processing over Networks*, 6:261–270, 2020.
- [20] Antonio G Marques, Santiago Segarra, Geert Leus, and Alejandro Ribeiro. Stationary graph processes and spectral estimation. *IEEE Transactions on Signal Processing*, 65(22):5911–5926, 2017.
- [21] Gautier Marti, Frank Nielsen, Mikołaj Bińkowski, and Philippe Donnat. A review of two decades of correlations, hierarchies, networks and clustering in financial markets. *Progress in Information Geometry*, pages 245–274, 2021.
- [22] Gonzalo Mateos, Santiago Segarra, Antonio G Marques, and Alejandro Ribeiro. Connecting the dots: Identifying network structure via graph signal processing. *IEEE Signal Processing Magazine*, 36(3):16–43, 2019.
- [23] Madeline Navarro, Yuhao Wang, Antonio G Marques, Caroline Uhler, and Santiago Segarra. Joint inference of multiple graphs from matrix polynomials. *Journal of Machine Learning Research*, 23(76):1–35, 2022.
- [24] Bastien Padeloup, Vincent Gripon, Grégoire Mercier, Dominique Pastor, and Michael G Rabbat. Characterization and inference of graph diffusion processes from observations of stationary signals. *IEEE transactions on Signal and Information Processing over Networks*, 4(3):481–496, 2017.
- [25] Nathanaël Perraudin and Pierre Vandergheynst. Stationary signal processing on graphs. *IEEE Transactions on Signal Processing*, 65(13):3462–3477, 2017.
- [26] Raksha Ramakrishna, Hoi-To Wai, and Anna Scaglione. A user guide to low-pass graph signal processing and its applications: Tools and applications. *IEEE Signal Processing Magazine*, 37(6):74–85, 2020.
- [27] Samuel Rey, Andrei Buciualea, Madeline Navarro, Santiago Segarra, and Antonio G Marques. Joint inference of multiple graphs with hidden variables from stationary graph signals. In *Proceedings of 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2022)*, pages 5817–5821. IEEE, 2022.
- [28] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [29] Johannes O Royset and Roger JB Wets. *An Optimization Primer*. Springer, 2021.
- [30] Santiago Segarra, Sundeep Prabhakar Chepuri, Antonio G Marques, and Geert Leus. Statistical graph signal processing: Stationarity and spectral estimation. *Cooperative and Graph Signal Processing*, pages 325–347, 2018.
- [31] Santiago Segarra, Antonio G Marques, Gonzalo Mateos, and Alejandro Ribeiro. Network topology inference from spectral templates. *IEEE Transactions on Signal and Information Processing over Networks*, 3(3):467–483, 2017.
- [32] Santiago Segarra, Yuhao Wang, Caroline Uhler, and Antonio G Marques. Joint inference of networks from stationary graph signals. In *Proceedings of the 51st Asilomar Conference on Signals, Systems, and Computers*, pages 975–979. IEEE, 2017.
- [33] Rasoul Shafipour, Abolfazl Hashemi, Gonzalo Mateos, and Haris Vikalo. Online topology inference from streaming stationary graph signals. In *2019 IEEE Data Science Workshop (DSW)*, pages 140–144. IEEE, 2019.
- [34] Rasoul Shafipour and Gonzalo Mateos. Online topology inference from streaming stationary graph signals with partial connectivity information. *Algorithms*, 13(9):228, 2020.

- [35] Rasoul Shafipour, Santiago Segarra, Antonio G Marques, and Gonzalo Mateos. Network topology inference from non-stationary graph signals. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5870–5874. IEEE, 2017.
- [36] Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, 2011.
- [37] Oliver Stegle, Sarah A Teichmann, and John C Marioni. Computational and analytical challenges in single-cell transcriptomics. *Nature Reviews Genetics*, 16(3):133–145, 2015.
- [38] Yuichi Tanaka and Yonina C Eldar. Generalized sampling on graphs with subspace and smoothness priors. *IEEE Transactions on Signal Processing*, 68:2272–2286, 2020.
- [39] Yuichi Tanaka, Yonina C Eldar, Antonio Ortega, and Gene Cheung. Sampling signals on graphs: From theory to applications. *IEEE Signal Processing Magazine*, 37(6):14–30, 2020.
- [40] Roman Vershynin. How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability*, 25(3):655–686, 2012.
- [41] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.
- [42] Xiaolu Wang, Chaorui Yao, Haoyu Lei, and Anthony Man-Cho So. An efficient alternating direction method for graph learning from smooth signals. In *Proceedings of 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2021)*, pages 5380–5384. IEEE, 2021.
- [43] Chenyue Zhang, Yiran He, and Hoi-To Wai. Product graph learning from multi-attribute graph signals with inter-layer coupling. In *Proceedings of 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2023)*, pages 1–5. IEEE, 2023.
- [44] Hui Zhang, Ming Yan, and Wotao Yin. One condition for solution uniqueness and robustness of both l_1 -synthesis and l_1 -analysis minimizations. *Advances in Computational Mathematics*, 42(6):1381–1399, 2016.
- [45] Xiaoqun Zhang, Martin Burger, and Stanley Osher. A unified primal-dual algorithm framework based on bregman iteration. *Journal of Scientific Computing*, 46(1):20–46, 2011.
- [46] Licheng Zhao, Yiwei Wang, Sandeep Kumar, and Daniel P Palomar. Optimization algorithms for graph laplacian estimation via admm and mm. *IEEE Transactions on Signal Processing*, 67(16):4231–4244, 2019.
- [47] Shuai Zheng, Zhenfeng Zhu, Zhizhe Liu, Zhenyu Guo, Yang Liu, Yuchen Yang, and Yao Zhao. Multi-modal graph learning for disease prediction. *IEEE Transactions on Medical Imaging*, 41(9):2207–2216, 2022.

A Organization of the Appendix

The appendix includes the missing proofs, detailed discussions of some argument in the main body and more numerical experiments. We organize the appendix as follows:

- The proof of infeasibility condition (Theorem 3.2) is provided in Section B.
- Explanations on conditions derived in Theorem 3.2 are included in Section C.
- The proof of properties of the proposed model (r)LogSpecT (Proposition 3.4 & 3.6) is given in Section D and some additional properties are discussed.
- The truncated Hausdorff distance based proof details of Theorem 4.1 and Corollary 4.4 are given in Section E.
- Details of L-ADMM and its convergence analysis are in Section F.
- Additional experiments and discussions on synthetic and real data are included in Section G.

B Proof of Theorem 3.2

Since the linear system (4) has no solution, we know from Farkas' lemma that the following system has solutions:

$$\begin{cases} \begin{bmatrix} \mathbf{I}_{m-1} & \mathbf{0}_{\frac{(m-1)(m-2)}{2}} \end{bmatrix} \mathbf{B}^\top \mathbf{A}_n^\top \mathbf{x} < \mathbf{0}_{(m-1) \times 1}, \\ \begin{bmatrix} \mathbf{0}_{\frac{(m-1)(m-2)}{2} \times (m-1)} & \mathbf{I}_{\frac{(m-1)(m-2)}{2}} \end{bmatrix} \mathbf{B}^\top \mathbf{A}_n^\top \mathbf{x} \leq \mathbf{0}_{\frac{(m-1)(m-2)}{2} \times 1}. \end{cases} \quad (11)$$

Let $\mathbf{x}^* \in \mathbb{R}^{m^2}$ be a solution to (11). Denote $\mathbf{x}_+ := \max\{\mathbf{x}^*, \mathbf{0}\}$, $\mathbf{x}_- := \max\{-\mathbf{x}^*, \mathbf{0}\}$. Then, there exists $c \in (0, 1]$ such that

$$\mathbf{B}^\top \mathbf{A}_n^\top (\mathbf{x}_+ - \mathbf{x}_-) + c \mathbf{1}_{m^2}^\top (\mathbf{x}_+ + \mathbf{x}_-) [\mathbf{1}_{m-1}; \mathbf{0}_{\frac{(m-1)(m-2)}{2}}] \leq \mathbf{0}.$$

Define $y := -\mathbf{1}_{m^2}^\top (\mathbf{x}_+ + \mathbf{x}_-)$, $z := c \mathbf{1}_{m^2}^\top (\mathbf{x}_+ + \mathbf{x}_-)$ and set $\bar{\delta} = c$. For all $\delta \in [0, \bar{\delta})$, $(\mathbf{x}_+, \mathbf{x}_-, y, z)$ is a solution to the following linear system:

$$\begin{cases} \mathbf{B}^\top \mathbf{A}_n^\top (\mathbf{x}_+ - \mathbf{x}_-) + z [\mathbf{1}_{m-1}; \mathbf{0}_{\frac{(m-1)(m-2)}{2}}] \leq \mathbf{0}, \\ \mathbf{1}_{m^2}^\top (\mathbf{x}_+ + \mathbf{x}_-) + y \leq 0, \\ \delta y + z > 0, \\ \mathbf{x}_+, \mathbf{x}_-, -y \geq \mathbf{0}. \end{cases}$$

Again, from Farkas' lemma, this implies that the following linear system does not have a solution:

$$\begin{cases} \mathbf{A}_n \mathbf{B} \mathbf{s} + t \mathbf{1}_{m^2} \geq \mathbf{0}, \\ \mathbf{A}_n \mathbf{B} \mathbf{s} - t \mathbf{1}_{m^2} \leq \mathbf{0}, \\ t \leq \delta, \\ \begin{bmatrix} \mathbf{1}_{m-1} & \mathbf{0}_{\frac{(m-1)(m-2)}{2}} \end{bmatrix} \mathbf{s} = \mathbf{1}, \end{cases} \quad (12)$$

where $\mathbf{s} \in \mathbb{R}^{m(m-1)/2}$ and $t \in \mathbb{R}$. Since (12) is equivalent to:

$$\begin{cases} \|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_{\infty, \infty} \leq \delta, \\ (\mathbf{S} \mathbf{1})_1 = 1, \\ \mathbf{S} \in \mathcal{S}, \end{cases} \quad (13)$$

the above argument indicates that (13) does not have a solution. Suppose rSpecT has a feasible solution \mathbf{S}' , then

$$\|\mathbf{C}_n \mathbf{S}' - \mathbf{S}' \mathbf{C}_n\|_{\infty, \infty} \leq \|\mathbf{C}_n \mathbf{S}' - \mathbf{S}' \mathbf{C}_n\|_F \leq \delta.$$

Hence, \mathbf{S}' is also a solution to (13). However, (13) does not have a solution. We can conclude that rSpecT is infeasible in this case.

C Explanations on Sufficient Conditions in Theorem 3.2

We elaborate more on the infeasibility condition that $\mathbf{A}_n \mathbf{B}$ has full column rank. An application of the condition is Example 3.1. Specifically, we know that in this case,

$$\mathbf{B} = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix} \quad \text{and} \quad \mathbf{A}_n = \begin{pmatrix} 0 & h_{12} & -h_{12} & 0 \\ h_{12} & h_{22} - h_{11} & 0 & -h_{12} \\ -h_{12} & 0 & h_{11} - h_{22} & h_{12} \\ 0 & -h_{12} & h_{12} & 0 \end{pmatrix}.$$

This implies that

$$\mathbf{A}_n \mathbf{B} = \begin{pmatrix} 0 \\ h_{22} - h_{11} \\ h_{11} - h_{22} \\ 0 \end{pmatrix}.$$

Hence, when $h_{11} \neq h_{22}$, $\mathbf{A}_n \mathbf{B}$ has full column rank. This means that when δ is small enough (from Example 3.1 we know $\tilde{\delta} = \sqrt{2}|h_{11} - h_{22}|$), the model rSpecT is infeasible.

D Proofs of Properties of (r)LogSpecT

D.1 Proof of Proposition 3.4

Since the constraint set \mathcal{S} is a cone, it follows that for all $\gamma > 0$, $\gamma \mathcal{S} = \mathcal{S}$. Then, we know that

$$\begin{aligned} \text{Opt}(\mathbf{C}, \alpha) &= \underset{\mathbf{S} \in \mathcal{S}, \mathbf{C}\mathbf{S} = \mathbf{S}\mathbf{C}}{\text{argmin}} \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) \\ &= \gamma \cdot \underset{\gamma \mathbf{S} \in \mathcal{S}, \mathbf{C}\gamma \mathbf{S} = \gamma \mathbf{S}\mathbf{C}}{\text{argmin}} \|\gamma \mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\gamma \mathbf{S}\mathbf{1}) \\ &= \gamma \cdot \underset{\mathbf{S} \in \frac{1}{\gamma} \mathcal{S}, \mathbf{C}\mathbf{S} = \mathbf{S}\mathbf{C}}{\text{argmin}} \gamma \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) \\ &= \gamma \cdot \underset{\mathbf{S} \in \mathcal{S}, \mathbf{C}\mathbf{S} = \mathbf{S}\mathbf{C}}{\text{argmin}} \|\mathbf{S}\|_{1,1} - \frac{\alpha}{\gamma} \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) \\ &= \gamma \text{Opt}(\mathbf{C}, \alpha/\gamma), \end{aligned}$$

where the third equality is from the basic calculus rule of the logarithm function. Set $\gamma = \alpha$ and then $\text{Opt}(\mathbf{C}, \alpha) = \alpha \text{Opt}(\mathbf{C}, 1)$, which completes the proof.

D.2 Proof of Proposition 3.6

The proof will be conducted by constructing a feasible solution for rLogSpecT. Recall that $\mathbf{A}_n = \mathbf{I} \otimes \mathbf{C}_n - \mathbf{C}_n \otimes \mathbf{I}$ and the matrix $\mathbf{B} \in \mathbb{R}^{m^2 \times m(m-1)/2}$ that maps a non-negative vector to the vectorization of a valid adjacency matrix. Let $\mathbf{S} = \min\left\{\frac{\delta}{\|\mathbf{A}_n \mathbf{B} \mathbf{s}\|_2}, 1\right\} \cdot \text{mat}(\mathbf{B} \mathbf{s})$ with $\mathbf{s} \in \mathbb{R}^{(m-1)m/2}$ being a non-negative vector, where $\text{mat}(\cdot)$ is the matricization operator. Note that

$$\text{vec}(\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n) = (\mathbf{I} \otimes \mathbf{C}_n - \mathbf{C}_n \otimes \mathbf{I}) \text{vec}(\mathbf{S}) = \mathbf{A}_n \text{vec}(\mathbf{S}).$$

Then, we know that

$$\|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F = \|\text{vec}(\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n)\|_2 = \min\left\{\frac{\delta}{\|\mathbf{A}_n \mathbf{B} \mathbf{s}\|_2}, 1\right\} \cdot \|\mathbf{A}_n \mathbf{B} \mathbf{s}\|_2 \leq \delta.$$

Thus, the given \mathbf{S} is a feasible solution for rLogSpecT and it completes the proof.

D.3 Properties of optimal solutions and values of (r)LogSpecT

In this section, we further discuss some properties of the optimal solutions/value of the proposed models, which are useful for deriving the recovery guarantee. More specifically, we obtain an upper bound on the optimal solutions (which may not be unique) independent of the sample size n and the inaccuracy parameter δ_n . Also, a lower bound of optimal values follows.

Proposition D.1. *The following statements hold:*

- For an optimal solution \mathbf{S}^* (resp. \mathbf{S}_n^*) to LogSpecT (resp. rLogSpecT with any given sample size n), it follows that

$$\|\mathbf{S}^*\|_{1,1} = \alpha m \text{ and } \|\mathbf{S}_n^*\|_{1,1} \leq \alpha m, \quad \forall \delta_n > 0.$$

- If $\delta_n \geq 2\alpha m \|\mathbf{C}_n - \mathbf{C}_\infty\|$, then

$$\alpha m(1 - \log \alpha) \leq f_n^* \leq f^*, \quad \forall n,$$

where f^* (resp. f_n^*) denotes the optimal value of LogSpecT (resp. rLogSpecT).

For the first statement, let us consider the Karush-Kuhn-Tucker (KKT) conditions of LogSpecT and rLogSpecT . Since the LogSpecT is a convex problem and Slater's condition holds, the KKT conditions are necessary and sufficient for the optimality, i.e., there exists $(\mathbf{\Lambda}_1, \mathbf{\Lambda}_2) \in \mathbb{R}^{m \times m} \times \mathcal{N}_S(\mathbf{S}^*)$ such that

$$\begin{cases} \nabla_S(\|\mathbf{S}^*\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}^* \mathbf{1})) + \mathbf{C}_\infty \mathbf{\Lambda}_1 - \mathbf{\Lambda}_1 \mathbf{C}_\infty + \mathbf{\Lambda}_2 = \mathbf{0}, \\ \mathbf{C}_\infty \mathbf{S}^* = \mathbf{S}^* \mathbf{C}_\infty, \\ \mathbf{S}^* \in \mathcal{S}, \end{cases} \quad (14)$$

where $\mathcal{N}_S(\mathbf{S}^*) := \{\mathbf{N} \in \mathbb{R}^{m \times m} : \sup_{\mathbf{X} \in \mathcal{S}} \langle \mathbf{X} - \mathbf{S}^*, \mathbf{N} \rangle \leq 0\}$ is the normal cone of \mathcal{S} at \mathbf{S}^* , and $\nabla \|\mathbf{S}^*\|_{1,1}$ is well-defined since $\|\cdot\|_{1,1} = \langle \cdot, \mathbf{1} \mathbf{1}^\top \rangle$ at $\mathbf{S}^* \geq 0$, which is differentiable. Taking further calculation gives that

$$\nabla \|\mathbf{S}^*\|_{1,1} = \mathbf{1} \mathbf{1}^\top, \quad (\nabla_S \mathbf{1}^\top \log(\mathbf{S}^* \mathbf{1}))_{ij} = \frac{1}{(\mathbf{S}^* \mathbf{1})_i}.$$

Combining this with (14) by taking inner product of both sides with \mathbf{S}^* , we obtain that

$$\sum_{i,j} (\mathbf{S}^*)_{ij} - \alpha \sum_{i,j} \frac{(\mathbf{S}^*)_{ij}}{(\mathbf{S}^* \mathbf{1})_i} + \langle \mathbf{\Lambda}_1, \mathbf{C}_\infty \mathbf{S}^* - \mathbf{S}^* \mathbf{C}_\infty \rangle + \langle \mathbf{\Lambda}_2, \mathbf{S}^* \rangle = 0. \quad (15)$$

From the structure of \mathcal{S} and the fact that $\mathbf{\Lambda}_2 \in \mathcal{N}_S(\mathbf{S}^*)$, one has that $\langle \mathbf{\Lambda}_2, \mathbf{S}^* \rangle = 0$. Also, note that $\mathbf{C}_\infty \mathbf{S}^* = \mathbf{S}^* \mathbf{C}_\infty$. Hence, the equation (15) can be simplified as the desired result:

$$\|\mathbf{S}^*\|_{1,1} = \sum_{i,j} (\mathbf{S}^*)_{ij} = \alpha \sum_{i,j} \frac{(\mathbf{S}^*)_{ij}}{(\mathbf{S}^* \mathbf{1})_i} = \alpha \sum_{i=1}^m \sum_{j=1}^m \frac{(\mathbf{S}^*)_{ij}}{(\mathbf{S}^* \mathbf{1})_i} = \alpha m.$$

The KKT conditions of rLogSpecT indicate that there exist $\lambda_1 \geq 0$, $\mathbf{\Lambda}_2 \in \mathcal{N}_S(\mathbf{S}_n^*)$ and $\mathbf{Q} \in \partial \|\mathbf{C}_n \mathbf{S}_n^* - \mathbf{S}_n^* \mathbf{C}_n\|_F$ (i.e., the subgradient of the function $\mathbf{S} \mapsto \|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F$ at \mathbf{S}_n^*) such that

$$\begin{cases} \nabla_S(\|\mathbf{S}_n^*\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}_n^* \mathbf{1})) + \lambda_1 \mathbf{Q} + \mathbf{\Lambda}_2 = \mathbf{0}, \\ \lambda_1 (\|\mathbf{C}_n \mathbf{S}_n^* - \mathbf{S}_n^* \mathbf{C}_n\|_F - \delta_n) = 0, \\ \mathbf{S}_n^* \in \mathcal{S}. \end{cases} \quad (16)$$

Moreover, from the definition of the convex subdifferential we know that $0 \geq \|\mathbf{C}_n \mathbf{S}_n^* - \mathbf{S}_n^* \mathbf{C}_n\|_F - \langle \mathbf{Q}, \mathbf{S}_n^* \rangle$. Thus, after taking inner product of both sides of the equation (16) with \mathbf{S}_n^* , it follows that:

$$\begin{aligned} 0 &= \sum_{i,j} (\mathbf{S}_n^*)_{ij} - \alpha m + \lambda_1 \langle \mathbf{Q}, \mathbf{S}_n^* \rangle + \langle \mathbf{\Lambda}_2, \mathbf{S}_n^* \rangle \\ &\geq \sum_{i,j} (\mathbf{S}_n^*)_{ij} - \alpha m + \lambda_1 \|\mathbf{C}_n \mathbf{S}_n^* - \mathbf{S}_n^* \mathbf{C}_n\|_F + \langle \mathbf{\Lambda}_2, \mathbf{S}_n^* \rangle \\ &= \sum_{i,j} (\mathbf{S}_n^*)_{ij} - \alpha m + \lambda_1 \delta_n, \end{aligned}$$

which implies that $\sum_{i,j} (\mathbf{S}_n^*)_{ij} \leq \alpha m - \lambda_1 \delta_n \leq \alpha m$. This completes the proof of the first statement.

For the second statement, we first prove that v_n^* and v^* are larger than $\alpha m(1 - \log \alpha)$. Define the auxiliary function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $g(x) := x - \alpha \log x$ for any $x \in \mathbb{R}_+$, whose minimum is attained at α . Since for any $\mathbf{S} \in \mathcal{S}$,

$$f(\mathbf{S}) = \sum_{i=1}^m g \left(\sum_{j=1}^m S_{ij} \right),$$

where f is the objective in LogSpecT, it follows that

$$f(\mathbf{S}) \geq \sum_{i=1}^m g(\alpha) = \alpha m(1 - \log \alpha).$$

This implies that v_n^* and v^* are larger than $\alpha m(1 - \log \alpha)$. Next, we will show $v_n^* \leq v^*$. Consider any optimal solution \mathbf{S}^* to LogSpecT. We show that it is feasible for rLogSpecT.

$$\begin{aligned} \|\mathbf{C}_n \mathbf{S}^* - \mathbf{S}^* \mathbf{C}_n\| &= \|\mathbf{C}_n \mathbf{S}^* - \mathbf{C}_\infty \mathbf{S}^* + \mathbf{S}^* \mathbf{C}_\infty - \mathbf{S}^* \mathbf{C}_n\| \\ &\leq 2\|\mathbf{S}^*\|_{1,1} \|\mathbf{C}_n - \mathbf{C}_\infty\| \leq 2\alpha m \|\mathbf{C}_n - \mathbf{C}_\infty\| \leq \delta_n, \end{aligned}$$

where the equality comes from $\mathbf{C}_\infty \mathbf{S}^* = \mathbf{S}^* \mathbf{C}_\infty$, the first inequality comes from the fact that $\|\mathbf{X}\mathbf{Y}\| \leq \|\mathbf{X}\|_F \|\mathbf{Y}\| \leq \|\mathbf{X}\|_{1,1} \|\mathbf{Y}\|$, the second one comes from the first statement and the last one is due to $\delta_n \geq 2\alpha m \|\mathbf{C}_n - \mathbf{C}_\infty\|$. Hence, \mathbf{S}^* is feasible for rLogSpecT, which indicates that $v_n^* \leq v^*$. The proof is completed.

E Proof of Theorem 4.1 & Corollary 4.4

E.1 Truncated Hausdorff distance

In this section, we introduce an advanced technique in optimization that is efficient in analyzing the recovery guarantee of robust formulations. Before that, we introduce the concept of truncated Hausdorff distance between two sets.

Definition E.1 (Truncated Hausdorff Distance [29, 6.J]). *For any $\rho \geq 0$, the truncated Hausdorff distance between two sets \mathcal{C} and \mathcal{D} is defined as:*

$$\hat{d}_\rho(\mathcal{C}, \mathcal{D}) = \max\{\text{dist}(\mathcal{C} \cap \mathbb{B}(\mathbf{0}, \rho), \mathcal{D}), \text{dist}(\mathcal{D} \cap \mathbb{B}(\mathbf{0}, \rho), \mathcal{C})\}.$$

It turns out that the distance between the optimum of two minimization problems can be bounded with the truncated Hausdorff distance of the epigraphs under some conditions. The result is captured in the following lemma.

Lemma E.2 ([29, Theorem 6.56]). *Let $\rho \in [0, \infty)$. Suppose that the extended-real-valued lower semicontinuous functions $f, g : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ satisfy*

- $\inf f, \inf g \in [-\rho, \rho]$,
- $\text{argmin } f, \text{argmin } g \subseteq \mathbb{B}(\mathbf{0}, \rho)$.

Then, it follows that

$$|\inf f - \inf g| \leq \hat{d}_\rho(\text{epi } f, \text{epi } g).^3 \tag{17}$$

Suppose further that $\varepsilon \geq 2\hat{d}_\rho(\text{epi } f, \text{epi } g)$, then one has

$$\text{dist}(\mathbf{x}_g^*, \varepsilon\text{-argmin } f) \leq \hat{d}_\rho(\text{epi } f, \text{epi } g), \tag{18}$$

where $\varepsilon\text{-argmin } f$ is the ε -suboptimal solution set of f that is defined as $\varepsilon\text{-argmin } f := \{\mathbf{x} \in \text{dom } f : f(\mathbf{x}) \leq \inf f + \varepsilon\}$, and \mathbf{x}_g^ is a minimizer of g .*

From the above lemma, we know that if two optimization problems are close enough (in the sense of truncated Hausdorff distance), then the optimum of them should be close to each other. Hence, in order to apply this result, we need to bound the truncated Hausdorff distance in an explicit way, which is solved by the following Kenmochi condition.

³For a function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$, its epigraph is defined as $\text{epi } f := \{(\mathbf{x}, y) \mid y \geq f(\mathbf{x})\}$.

Lemma E.3 (Kenmochi Condition [29, Proposition 6.58]). *Let $\rho \in [0, \infty)$. Then, for $f, g : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ with nonempty epigraphs, one has that*

$$\hat{d}_\rho(\text{epi } f, \text{epi } g) = \inf \left\{ \eta > 0 : \begin{array}{l} \inf_{\mathbb{B}(\mathbf{x}, \eta)} g \leq \max\{f(\mathbf{x}), -\rho\} + \eta, \forall \mathbf{x} \in [f \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho) \\ \inf_{\mathbb{B}(\mathbf{x}, \eta)} f \leq \max\{g(\mathbf{x}), -\rho\} + \eta, \forall \mathbf{x} \in [g \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho) \end{array} \right\},$$

where $[f \leq \rho] := \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) \leq \rho\}$.

E.2 Proof of Theorem 4.1

Before presenting the proof, we first introduce the following lemma.

Lemma E.4 (Hoffman's Error Bound [12]). *Consider the set $\mathcal{S} := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$. There exists $C > 0$ such that for any $\mathbf{x} \in \mathbb{R}^n$, one has*

$$\text{dist}(\mathbf{x}, \mathcal{S}) \leq C \cdot \|(\mathbf{A}\mathbf{x} - \mathbf{b})_+\|_2.$$

For the sake of brevity, we denote

$$\begin{aligned} \bar{f}_n(\mathbf{S}) &:= \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) + \iota_{\mathbb{R}_-}(\|\mathbf{C}_n \mathbf{S} - \mathbf{S} \mathbf{C}_n\|_F - \delta_n) + \iota_{\mathcal{S}}(\mathbf{S}), \\ \bar{f}(\mathbf{S}) &:= \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1}) + \iota_{\{0\}}(\|\mathbf{C}_\infty \mathbf{S} - \mathbf{S} \mathbf{C}_\infty\|_F) + \iota_{\mathcal{S}}(\mathbf{S}). \end{aligned}$$

Hence, the optimization problem LogSpecT (resp. rLogSpecT) is equivalent to $\inf \bar{f}$ (resp. $\inf \bar{f}_n$).

Now, we aim to use Lemma E.3 to bound $\hat{d}_\rho(\text{epi } \bar{f}, \text{epi } \bar{f}_n)$. Let $\mathbf{S} \in \mathcal{S} \cap \mathbb{B}(\mathbf{0}, \rho)$ satisfy

$$\bar{f}(\mathbf{S}) \leq \rho \quad \text{and} \quad \mathbf{S} \mathbf{C}_\infty = \mathbf{C}_\infty \mathbf{S}.$$

Then, we know that

$$\|\mathbf{S} \mathbf{C}_n - \mathbf{C}_n \mathbf{S}\|_F \leq 2\|\mathbf{S}\|_F \|\mathbf{C}_n - \mathbf{C}_\infty\| \leq 2\rho \|\mathbf{C}_n - \mathbf{C}_\infty\| \leq \delta_n,$$

and consequently \mathbf{S} is in the domain of \bar{f}_n . Then, it follows that for any $\eta > 0$, we have

$$\inf_{\mathbb{B}(\mathbf{S}, \eta)} \bar{f}_n \leq \bar{f}_n(\mathbf{S}) = \bar{f}(\mathbf{S}) \leq \max\{\bar{f}(\mathbf{S}), -\rho\}, \quad \forall \mathbf{S} \in [\bar{f} \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho). \quad (19)$$

Before verifying the reverse side of the Kenmochi condition, we first consider the non-emptiness of $[\bar{f}_n \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho)$. Since

$$\delta_n \geq 2\rho \|\mathbf{C}_n - \mathbf{C}_\infty\| \geq 2\alpha m \|\mathbf{C}_n - \mathbf{C}_\infty\|,$$

it follows from Proposition D.1 that $\|\mathbf{S}_n^*\|_{1,1} \leq \alpha m \leq \rho$ and $f_n^* \leq f^* \leq \rho$, which implies that $[\bar{f}_n \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho)$ is nonempty. Let $\mathbf{S}_n \in [\bar{f}_n \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho)$. Then, one has that

$$\mathbf{S}_n \in \mathcal{S} \quad \text{and} \quad \|\mathbf{C}_n \mathbf{S}_n - \mathbf{S}_n \mathbf{C}_n\|_F \leq \delta_n.$$

Hence, it follows that

$$\|\mathbf{C}_\infty \mathbf{S}_n - \mathbf{S}_n \mathbf{C}_\infty\| \leq 2\|\mathbf{S}_n\|_F \|\mathbf{C}_\infty - \mathbf{C}_n\| + \|\mathbf{C}_n \mathbf{S}_n - \mathbf{S}_n \mathbf{C}_n\|_F \leq 2\rho \|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n.$$

Also, note that there exists $\beta > 0$ such that $(\mathbf{S}_n \mathbf{1})_i \geq \beta$ for all $i \in [m]$ as $\bar{f}_n \leq \rho$ and $\|\mathbf{S}_n\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}_n \mathbf{1}) \rightarrow \infty$ when $\mathbf{S}_n \rightarrow \mathbf{0}$. Thus, applying Lemma E.4 to the linear system

$$\tilde{\mathcal{S}} := \{\mathbf{S} \in \mathbb{R}^{m \times m} : \mathbf{S} \mathbf{C}_\infty = \mathbf{C}_\infty \mathbf{S}, \mathbf{S} \in \mathcal{S}, (\mathbf{S}\mathbf{1})_i \geq \beta, \forall i \in [m]\}$$

yields that there exists $\tilde{c} > 0$ such that

$$\text{dist}(\mathbf{S}_n, \tilde{\mathcal{S}}) \leq \tilde{c} \cdot (2\rho \|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n).$$

Hence, there exists $\tilde{\mathbf{S}}$ in the domain of \bar{f} such that

$$\|\mathbf{S}_n - \tilde{\mathbf{S}}\|_F \leq \tilde{c} \cdot (2\rho \|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n) \quad \text{and} \quad (\tilde{\mathbf{S}}\mathbf{1})_i \geq \beta, \quad \forall i \in [m].$$

Since the function $\mathbf{S} \mapsto \|\mathbf{S}\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}\mathbf{1})$ is locally Lipschitz continuous when $(\mathbf{S}\mathbf{1})_i \geq \beta$, there exists $L > 0$ such that

$$\begin{aligned} \bar{f}(\tilde{\mathbf{S}}) &= \|\tilde{\mathbf{S}}\|_{1,1} - \alpha \mathbf{1}^\top \log(\tilde{\mathbf{S}}\mathbf{1}) \leq \|\mathbf{S}_n\|_{1,1} - \alpha \mathbf{1}^\top \log(\mathbf{S}_n\mathbf{1}) + L\|\mathbf{S}_n - \tilde{\mathbf{S}}\|_F \\ &= \bar{f}_n(\mathbf{S}_n) + L\|\mathbf{S}_n - \tilde{\mathbf{S}}\|_F \\ &\leq \bar{f}_n(\mathbf{S}_n) + L\tilde{c} \cdot (2\rho\|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n). \end{aligned}$$

Setting $c_1 \geq \max\{1, L\} \cdot \tilde{c}$, one can obtain that for any $\mathbf{S}_n \in [\bar{f}_n \leq \rho] \cap \mathbb{B}(\mathbf{0}, \rho)$

$$\inf_{\mathbb{B}(\mathbf{S}_n, \eta)} \bar{f} \leq \bar{f}(\tilde{\mathbf{S}}) \leq \bar{f}_n(\mathbf{S}_n) + c_1 \cdot (2\rho\|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n) \leq \max\{\bar{f}_n(\mathbf{S}_n), -\rho\} + \eta, \quad (20)$$

where $\eta := c_1 \cdot (2\rho\|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n)$. Combining inequality (19) and (20), we can conclude that

$$\hat{d}_\rho(\text{epi } \bar{f}, \text{epi } \bar{f}_n) \leq c_1 \cdot (2\rho\|\mathbf{C}_\infty - \mathbf{C}_n\| + \delta_n). \quad (21)$$

In order to derive the conclusion (i) and (ii), it remains to check the requirements in Lemma E.2. Since $\rho \geq \alpha m$, the first statement of Proposition D.1 shows that the optimal solutions to $\inf \bar{f}$ and $\inf \bar{f}_n$ lie in $\mathbb{B}(\mathbf{0}, \rho)$. Since $\rho \geq f^*$ and $-\rho \leq \alpha m(1 - \log \alpha)$, the second statement of the proposition shows that $\inf \bar{f}, \inf \bar{f}_n \in [-\rho, \rho]$. Hence, applying Lemma E.2 completes the proof of the first two statements.

To prove conclusion (iii), we first make the following two claims:

- (a) $\mathcal{S}_0^*\mathbf{1}$ is a singleton, whose element is denoted by $\mathbf{S}^*\mathbf{1}$,
- (b) For any $\bar{\varepsilon} \in [0, \infty)$, there exists a $\delta(\bar{\varepsilon}) > 0$ such that for all $0 \leq \varepsilon \leq \bar{\varepsilon}$ and $\mathbf{S}_\varepsilon \in \mathcal{S}_\varepsilon^*$, one has that

$$\|\mathbf{S}_\varepsilon\mathbf{1} - \mathbf{S}^*\mathbf{1}\|_2 \leq \delta(\bar{\varepsilon}) \cdot \sqrt{\varepsilon}. \quad (22)$$

Granting these and with the help of Theorem 4.1 (ii), we can derive that for all $\mathbf{S}_n^* \in \mathcal{S}^{n,*}$

$$\begin{aligned} \text{dist}(\mathbf{S}_n^*\mathbf{1}, \mathbf{S}_0^*\mathbf{1}) &= \|\mathbf{S}_n^*\mathbf{1} - \mathbf{S}^*\mathbf{1}\|_2 \leq \|\mathbf{S}_n^*\mathbf{1} - \mathbf{S}_{2\varepsilon_n}\mathbf{1}\|_2 + \|\mathbf{S}_{2\varepsilon_n}\mathbf{1} - \mathbf{S}^*\mathbf{1}\|_2 \\ &\leq \sqrt{m} \text{dist}(\mathbf{S}_n^*, \mathbf{S}_{2\varepsilon_n}^*) + \|\mathbf{S}_{2\varepsilon_n}\mathbf{1} - \mathbf{S}^*\mathbf{1}\|_2 \\ &\leq \tilde{c}_1\varepsilon_n + \tilde{c}_2\sqrt{\varepsilon_n}, \end{aligned}$$

where \tilde{c}_1, \tilde{c}_2 are positive constants, and $\mathbf{S}_{2\varepsilon_n} \in \mathcal{S}_{2\varepsilon_n}^*$ satisfies $\|\mathbf{S}_n^* - \mathbf{S}_{2\varepsilon_n}\|_F = \text{dist}(\mathbf{S}_n^*, \mathcal{S}_{2\varepsilon_n}^*)$ (whose existence is guaranteed since $\mathcal{S}_\varepsilon^*$ is convex and compact). Hence,

$$\text{dist}(\mathcal{S}^{n,*}\mathbf{1}, \mathbf{S}_0^*\mathbf{1}) \leq \tilde{c}_1\varepsilon_n + \tilde{c}_2\sqrt{\varepsilon_n}.$$

To proceed, it remains to prove the claims. Define an auxiliary function $h : \mathbb{R}^m \rightarrow \mathbb{R}$ as $h(\mathbf{x}) = \sum_{i=1}^m x_i - \alpha \sum_{i=1}^m \log x_i$ for each $\mathbf{x} \in \mathbb{R}_+^m$. Consider the following optimization problem:

$$\begin{aligned} \min_{\mathbf{x}} \quad & h(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x} \in \{\mathbf{S}\mathbf{1} \in \mathbb{R}^m \mid \mathbf{S} \text{ that is feasible for LogSpecT}\}. \end{aligned} \quad (23)$$

For the sake of brevity, denote the ε -suboptimal solution set of (23) as $\mathcal{H}_\varepsilon^*$. In the remaining part, we will first show that $\mathcal{S}_\varepsilon^*\mathbf{1} = \mathcal{H}_\varepsilon^*$ and then, by the strict convexity of h , the desired two claims hold.

The first step is to show that the optimal function value of the problem (23) satisfies $h^* = f^*$. Since it is obvious that $\tilde{\mathbf{x}} = \mathbf{S}^*\mathbf{1}$ is feasible for (23), $h^* \leq h(\tilde{\mathbf{x}}) = f(\mathbf{S}^*) = f^*$. Suppose to the contrary that $h^* < f^*$, from the fact that the objective function is coercive and continuous and the feasible set is closed, there exists $\tilde{\mathbf{S}}$ such that it is feasible for LogSpecT and $\mathbf{x}^* = \tilde{\mathbf{S}}\mathbf{1}$, where \mathbf{x}^* is an optimal solution to (23). Since $h^* = h(\mathbf{x}^*) = h(\tilde{\mathbf{S}}\mathbf{1}) = f(\tilde{\mathbf{S}})$, this contradicts the fact that $f(\tilde{\mathbf{S}}) \geq f^*$. Hence, $h^* = f^*$. Next, we will show that $\mathcal{S}_\varepsilon^*\mathbf{1} = \mathcal{H}_\varepsilon^*$. Consider any ε -suboptimal solution $\mathbf{S} \in \mathcal{S}_\varepsilon^*$, i.e.,

$$h(\mathbf{S}\mathbf{1}) = f(\mathbf{S}) \leq f^* + \varepsilon = h^* + \varepsilon.$$

Hence, $\mathbf{S}\mathbf{1} \in \mathcal{H}_\varepsilon^*$ and it implies that $\mathcal{S}_\varepsilon^*\mathbf{1} \subseteq \mathcal{H}_\varepsilon^*$. On the other hand, for any ε -suboptimal solution $\mathbf{x} \in \mathcal{H}_\varepsilon^*$, there exists \mathbf{S} that is feasible for LogSpecT such that $\mathbf{x} = \mathbf{S}\mathbf{1}$. Thus,

$$f(\mathbf{S}) = h(\mathbf{x}) \leq h^* + \varepsilon = f^* + \varepsilon.$$

This implies that $\mathcal{S} \in \mathcal{S}_\varepsilon^*$ and consequently $\mathcal{H}_\varepsilon^* \subseteq \mathcal{S}_\varepsilon^* \mathbf{1}$. Hence, $\mathcal{H}_\varepsilon^* = \mathcal{S}_\varepsilon^* \mathbf{1}$.

Since h is strictly convex, its optimal solution set \mathcal{H}_0^* is a singleton. Then, $\mathcal{S}_0^* \mathbf{1} = \mathcal{H}_0^*$ is a singleton, which proves the first claim. For the second claim, we know that for any $\mathcal{S}_\varepsilon \in \mathcal{S}_\varepsilon^*$ there exists $\mathbf{x}_\varepsilon \in \mathcal{H}_\varepsilon^*$ such that

$$\|\mathcal{S}_\varepsilon \mathbf{1} - \mathcal{S}^* \mathbf{1}\|_2 = \|\mathbf{x}_\varepsilon - \mathbf{x}^*\|_2, \quad (24)$$

where $\mathbf{x}^* \in \mathcal{H}_0^*$. The coerciveness of h asserts that \mathbf{x}_ε and \mathbf{x}^* are bounded. This together with the fact that h is strongly convex on any bounded set, illustrates that there exists $\mu > 0$ such that

$$h(\mathbf{x}_\varepsilon) \geq h(\mathbf{x}^*) + \langle \nabla h(\mathbf{x}^*), \mathbf{x}_\varepsilon - \mathbf{x}^* \rangle + \frac{1}{\mu} \|\mathbf{x}_\varepsilon - \mathbf{x}^*\|_2^2 \geq h(\mathbf{x}^*) + \frac{1}{\mu} \|\mathbf{x}_\varepsilon - \mathbf{x}^*\|_2^2, \quad (25)$$

where the second inequality comes from the global optimality of \mathbf{x}^* . Combining (24) and (25) gives that

$$\|\mathcal{S}_\varepsilon \mathbf{1} - \mathcal{S}^* \mathbf{1}\|_2 = \|\mathbf{x}_\varepsilon - \mathbf{x}^*\|_2 \leq \sqrt{\mu(h(\mathbf{x}_\varepsilon) - h(\mathbf{x}^*))} \leq \sqrt{\mu\varepsilon}.$$

This completes the proof of the claims.

E.3 Proof of Corollary 4.4

Suppose to the contrary that there exists a sequence $\{\mathcal{S}_n^*\}_n$, where the n th element is an optimal solution to rLogSpecT with sample size n , such that

$$\text{dist}(\mathcal{S}_n^*, \mathcal{S}_0^*) \not\rightarrow 0.$$

From Proposition D.1, we know that $\{\mathcal{S}_n^*\}_n$ is bounded, and consequently, has a convergent subsequence. Without loss of generality, we may assume that the sequence itself is convergent and the limiting point is \mathcal{S}^* . Note that

$$\|\mathcal{C}_n \mathcal{S}_n^* - \mathcal{S}_n^* \mathcal{C}_n\|_F \leq \delta_n, \quad \mathcal{C}_n \rightarrow \mathcal{C}_\infty \quad \text{and} \quad \delta_n \rightarrow 0.$$

Hence, $\mathcal{C}_\infty \mathcal{S}^* = \mathcal{S}^* \mathcal{C}_\infty$. This indicates that \mathcal{S}^* is feasible for LogSpecT. Then, from Theorem 4.1, we know that $f(\mathcal{S}_n^*) = f_n^* \rightarrow f^*$, which leads to $f(\mathcal{S}^*) = f^*$ since $f(\cdot) = \|\cdot\|_{1,1} - \alpha \mathbf{1}^\top \log(\cdot \mathbf{1})$ is continuous. Together with the fact that \mathcal{S}^* is feasible, we conclude that \mathcal{S}^* is an optimal solution to LogSpecT. This further implies that $\text{dist}(\mathcal{S}_n^*, \mathcal{S}_0^*) \rightarrow 0$, which is a contradiction.

E.4 Proof of Lemma 4.7

Recall the generative model (1). Since \mathbf{w} follows a sub-Gaussian distribution, it can be shown that for every $t > 0$,

$$\mathbb{P}(\|\mathbf{x}\|_2 > t) \leq \mathbb{P}\left(\|\mathbf{w}\|_2 > \frac{t}{\|\mathcal{H}(\mathcal{S})\|}\right) \leq C e^{-v't^2},$$

for some positive constant v' , which means that \mathbf{x} also follows a sub-Gaussian distribution. Thus, due to the sub-Gaussian property, $\|\mathcal{C}_n - \mathcal{C}_\infty\|$ can be explicitly bounded by the following lemma.

Lemma E.5 ([40, Proposition 2.1]). *Consider sub-Gaussian, identical, independent random vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^m$ with $n > m$. Then for all $\varepsilon > 0$, it follows that*

$$\mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top - \mathbb{E}[\mathbf{x} \mathbf{x}^\top]\right\|_2 \leq \varepsilon\right) \geq 1 - 2e^{2m - l\varepsilon^2 n},$$

for some constant $l > 0$.

Setting $\varepsilon^2 = (4/l) \log(2n)m/n$, Lemma E.5 indicates that with high probability (lower bounded by $1 - n^{-1}$),

$$\|\mathcal{C}_n - \mathcal{C}_\infty\| \leq \mathcal{O}\left(\sqrt{\frac{\log n}{n}}\right).$$

F Derivations of L-ADMM and Convergence Analysis

This section includes the details of L-ADMM for rLogSpecT.

F.1 Proof of Proposition 5.1

Note that the minimization problem (7) is separable for \mathbf{Z} and \mathbf{q} , and can be split into two subproblems:

$$\min_{\mathbf{Z} \in \mathbb{B}(\mathbf{0}, \delta_n)} \|\mathbf{C}_n \mathbf{S}^{(k)} - \mathbf{S}^{(k)} \mathbf{C}_n + \mathbf{\Lambda}^{(k)} / \rho - \mathbf{Z}\|_F^2, \quad (26)$$

$$\min_{\mathbf{q}} -\alpha \mathbf{1}^\top \log \mathbf{q} + \lambda_2^{(k)\top} (\mathbf{q} - \mathbf{S}^{(k)} \mathbf{1}) + \frac{\rho}{2} \|\mathbf{q} - \mathbf{S}^{(k)} \mathbf{1}\|_2^2. \quad (27)$$

For problem (26), the optimal solution is the projection of $\mathbf{C}_n \mathbf{S}^{(k)} - \mathbf{S}^{(k)} \mathbf{C}_n + \mathbf{\Lambda}^{(k)} / \rho$ onto $\mathbb{B}(\mathbf{0}, \delta_n)$, which is given by

$$\mathbf{Z}^{(k+1)} = \min \left\{ 1, \frac{\delta_n}{\|\tilde{\mathbf{Z}}\|_F} \right\} \tilde{\mathbf{Z}} \quad \text{with} \quad \tilde{\mathbf{Z}} = \mathbf{C}_n \mathbf{S}^{(k)} - \mathbf{S}^{(k)} \mathbf{C}_n + \mathbf{\Lambda}^{(k)} / \rho.$$

For problem (27), the first-order optimality condition gives

$$-\alpha \mathbf{1} / \mathbf{q} + \lambda_2^{(k)} + \rho (\mathbf{q} - \mathbf{S}^{(k)} \mathbf{1}) = 0.$$

This together with the fact that the objective function is convex implies that

$$\mathbf{q}^{(k+1)} = \frac{\tilde{\mathbf{q}} + \sqrt{\tilde{\mathbf{q}}^2 + 4\alpha/\rho \mathbf{1}}}{2} \quad \text{with} \quad \tilde{\mathbf{q}} = \frac{1}{\rho} (\rho \mathbf{S}^{(k)} \mathbf{1} - \lambda_2^{(k)}).$$

F.2 Calculation of $\Pi_{\mathcal{S}}(\cdot)$

The projection of \mathbf{X} to \mathcal{S} can be calculated via an optimization problem:

$$\begin{aligned} \min_{\mathbf{S}} \quad & \|\mathbf{X} - \mathbf{S}\|_F^2 \\ \text{s.t.} \quad & \mathbf{S}^\top = \mathbf{S}, \\ & S_{ii} = 0, \quad i = 1, 2, \dots, m, \\ & S_{ij} \geq 0, \quad \forall i, j, \end{aligned}$$

which is equivalent to

$$\begin{aligned} \min \quad & \sum_{i < j} ((X_{ij} - S_{ij})^2 + (X_{ji} - S_{ij})^2) \\ \text{s.t.} \quad & S_{ij} \geq 0, \quad \forall i < j, \\ & S_{ii} = 0, \quad \forall i. \end{aligned}$$

Hence

$$(\Pi_{\mathcal{S}}(\mathbf{X}))_{ij} = \begin{cases} \frac{1}{2} \max\{0, X_{ij} + X_{ji}\}, & i \neq j, \\ 0, & i = j. \end{cases}$$

F.3 Stopping criterion and updating rule of ρ

We follow the procedures in [3] to update ρ in each iteration. Similarly, we define the primal residual and dual residual as follows:

$$\begin{aligned} p_{\text{res}}^{(k+1)} &= \sqrt{\|\mathbf{Z}^{(k+1)} - \mathbf{C}_n \mathbf{S}^{(k+1)} + \mathbf{S}^{(k+1)} \mathbf{C}_n\|_F^2 + \|\mathbf{q}^{(k+1)} - \mathbf{S}^{(k+1)} \mathbf{1}\|_2^2}, \\ d_{\text{res}}^{(k+1)} &= \rho^{(k)} \left(\mathbf{C}_n (\mathbf{S}^{(k+1)} - \mathbf{S}^{(k)}) - (\mathbf{S}^{(k+1)} - \mathbf{S}^{(k)}) \mathbf{C}_n + \mathbf{1}^\top (\mathbf{S}^{(k+1)} - \mathbf{S}^{(k)}) \mathbf{1} \right). \end{aligned}$$

The aim of updating ρ is to control the decaying speed of p_{res} and d_{res} such that their difference is not too large. To this end, we update ρ adaptively following the scheme:

$$\rho^{(k+1)} := \begin{cases} 2\rho^{(k)}, & \text{if } p_{\text{res}}^{(k+1)} > 5d_{\text{res}}^{(k+1)}, \\ \rho^{(k)}/2, & \text{if } d_{\text{res}}^{(k+1)} > 5p_{\text{res}}^{(k+1)}, \\ \rho^{(k)}, & \text{otherwise.} \end{cases}$$

When p_{res} and d_{res} are both smaller than the threshold $\varepsilon = 10^{-5}$, we stop the algorithm.

F.4 Convergence analysis

Define $\mathbf{D} := \text{Diag}(\mathbf{1}_m^\top, \dots, \mathbf{1}_m^\top) \in \mathbb{R}^{m \times m^2}$. Then, \mathbf{D} satisfies $\mathbf{D}\text{vec}(\mathbf{S}) = \mathbf{S}\mathbf{1}$ and $\|\mathbf{D}^\top \mathbf{D}\| = m$. Denote

$$\mathbf{Q} := \tau \mathbf{I} - \mathbf{D}^\top \mathbf{D} - \mathbf{A}_n^\top \mathbf{A}_n.$$

Then the linearized ADMM update (8) of \mathbf{S} can be written as:

$$\min_{\mathbf{S}} L(\mathbf{S}) + \frac{\rho}{2} \|\text{vec}(\mathbf{S}) - \text{vec}(\mathbf{S}^{(k)})\|_{\mathbf{Q}},$$

where $\|x\|_{\mathbf{Q}} := x^\top \mathbf{Q} x$. Since $\tau > m + \|\mathbf{A}_n\|^2$, we know that \mathbf{Q} is positively definite. Consequently, by treating (\mathbf{Z}, \mathbf{q}) as one variable, we can apply Theorem 4.2 in [45] and directly obtain the result.

G More Experiments and Discussions on Synthetic and Real Data

G.1 Influence of graph filters

To make a fair comparison between rSpecT and rLogSpecT, we test rSpecT on BA graphs with the same graph filters and the results are reported in Figure 6. It is obvious that rSpecT fails in these cases and cannot benefit from the increase in sample size. This is reasonable since SpecT fails on BA graphs as indicated in Figure 3, let alone the approximation formulation rSpecT.

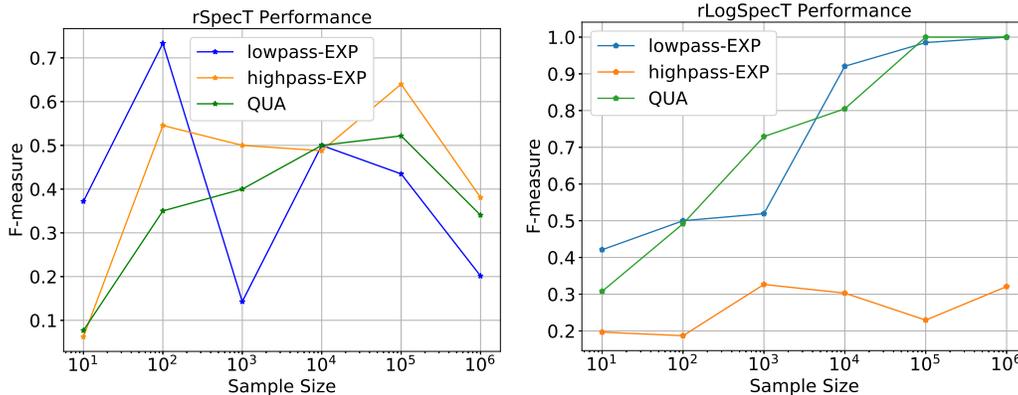


Figure 6: Performance of rSpecT on BA graphs. Figure 7: rLogSpecT on ER graphs with $\delta_n = 20\sqrt{\log n/n}$.

We further test rLogSpecT on ER graphs with different numbers of signals observed. The parameter δ_n is set as $20\sqrt{\log n/n}$ and the results are reported in Figure 7. The figure shows that for graph filters that are not high-pass, rLogSpecT can achieve nearly perfect recovery when the sample size is large enough. Also, compared with the performance on BA graphs, rLogSpecT works better on ER graphs. This observation is in accordance with the conclusion from Figure 3 that LogSpecT performs better on ER graphs than BA ones. We further notice that the difference between the low-pass graph filter and the high-pass one is huge. To check the conjecture that rLogSpecT generally performs better on low-pass graph filters, we choose different graph filters $\exp(t\mathbf{S})$ with t ranging from -2 to 2 and conduct the experiments on ER graphs. When the graph shifting operator is the adjacency matrix, the positive low-pass parameter t corresponds to low-pass graph filters and the negative t corresponds to the high-pass ones [26, 11]. We omit the case when $t = 0$ since this filter does not contain any graph information (note that $\exp(0\mathbf{S}) = \mathbf{I}$).

We then repeat the experiments for 50 times and report the average results in Figure 8. The comparison between the performance of low-pass graph filters and high-pass graph filters indicates that the low-pass graph filters generally outperforms the high-pass ones. A closer look at the results shows that the performance grows faster when the absolute value of t is smaller. And eventually, the graph filter with smaller absolute value of t prevails. This observation is interesting since Figure 3 indicates that the choice of graph filters has few impacts on the model performance. One explanation is that both low-pass graph filters and high-pass graph filters attenuate some frequencies of the graph and the larger absolute value of t leads to the more loss of information carried by finite signals.

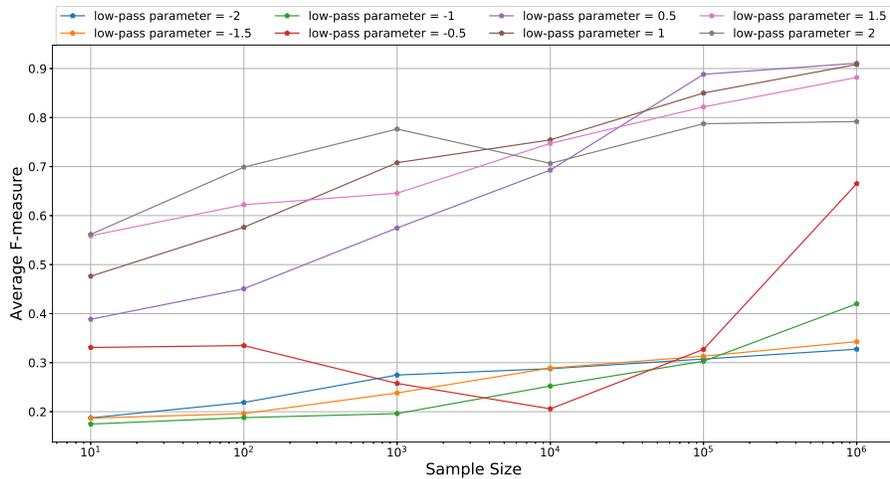


Figure 8: Effect of Low-Pass Parameter: different performance of graph filters $\exp(tS)$ with t ranging from -2 to 2 .

G.2 Experiments on USPS dataset

USPS dataset is a handwritten digit dataset. As shown in [25], it is nearly stationary with respect to the nearest neighbour graph. This dataset collects images of handwritten digits. In the experiment, each pixel is viewed as a node and the value on the pixel is view as the graph signal. We follow the approach in [25] to construct the 20 nearest neighbour graph, on which the data is verified to be nearly stationary. More specifically, we pick the 1296 digit 1 images. The weights between two nodes are decided by the Gaussian radial basis function. It turns out that the stationarity measure $s := \|\text{diag}(\mathbf{U}^\top \mathbf{C}_n \mathbf{U})\|_2 / \|\mathbf{U}^\top \mathbf{C}_n \mathbf{U}\|_F$ equals 0.78 in this dataset. Here \mathbf{U} is the eigen-matrix of the constructed graph and \mathbf{C}_n is the covariance matrix of the observed graph signals. We view the 20 nearest neighbour graph as the groundtruth and compare the learned graph from rLogSpecT with it. We present the subgraphs consisting of the top 10 nodes with the largest degrees in the two graphs respectively in Figure 9. The F-measure of these two subgraphs is 0.96. This result corroborates the efficiency of our proposed rLogSpecT model.

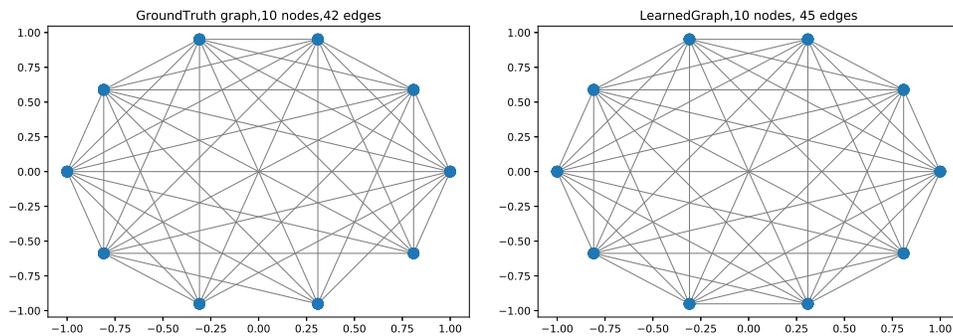


Figure 9: Graph Learning on USPS dataset. (**Left**: subgraph of the symmetric 20 nearest neighbour graph. **Right**: subgraph of the learned graph from rLogSpecT. The F-measure is 0.96)