

Scalable Inference of Sparsely-changing Markov Random Fields with Strong Statistical Guarantees

Salar Fattahi¹ and Andrés Gómez²

¹Industrial and Operations Engineering, University of Michigan

²Industrial & Systems Engineering, University of Southern California

February 2, 2021

Abstract

In this paper, we study the problem of inferring time-varying Markov random fields (MRF), where the underlying graphical model is both sparse and changes sparsely over time. Most of the existing methods for the inference of time-varying MRFs rely on the *regularized maximum likelihood estimation* (MLE), that typically suffer from weak statistical guarantees and high computational time. Instead, we introduce a new class of constrained optimization problems for the inference of sparsely-changing MRFs. The proposed optimization problem is formulated based on the exact ℓ_0 regularization, and can be solved in near-linear time and memory. Moreover, we show that the proposed estimator enjoys a provably small estimation error. As a special case, we derive sharp statistical guarantees for the inference of sparsely-changing Gaussian MRFs (GMRF) in the high-dimensional regime, showing that such problems can be learned with as few as one sample per time. Our proposed method is extremely efficient in practice: it can accurately estimate sparsely-changing graphical models with more than 500 million variables in less than one hour.

1 Introduction

Contemporary systems are comprised of massive numbers of interconnected components that interact according to a hierarchy of complex, unknown, and time-varying topologies. For example, with billions of neurons and hundreds of thousands of voxels, the human brain is considered as one of the most complex physiological networks [32, 17, 25, 27, 21].

The temporal behavior of today’s interconnected systems, such as those mentioned above, can be captured via *time-varying Markov random fields* (MRF). Time-varying MRFs are associated with a temporal sequence of undirected *Markov graphs* $\mathcal{G}_t(V, E_t)$, where V and E_t are the set of nodes and edges in the graph at time t . The node set V represents the random variables in the model, while the edge set E_t captures the conditional dependency between these variables at time t . A popular approach for the inference of graphical models is based on the so-called *maximum-likelihood estimation* (MLE): to obtain a model based on which the observed data is most probable to occur [38].

Despite being known as theoretically powerful tools [19, 34], MLE-based methods suffer from several fundamental *drawbacks* which render them impractical in realistic settings. First, they often suffer from notoriously high computational cost in massive problems, where the number of variables to be inferred is in the order of millions, or more. Second, they fail to efficiently incorporate *prior* structural information into their estimation procedure. For instance, it is well-known that most large-scale systems exhibit *sparse* interactions amongst their components, which can be captured via *sparsely-changing MRFs*, and estimated

using sparsity-promoting *regularizers* (such as ℓ_0 penalty). However, due to the inherent computational complexity of the ℓ_0 -regularized MLE, most of the existing methods inevitably resort to *relaxed* or weaker variants of such regularization (such as ℓ_1 penalty), thereby suffering from inferior statistical guarantees. The aforementioned drawbacks of MLE estimators are further compounded in time-varying settings, since parameters need to be estimated for *each* time period, resulting in a dramatic increase in the size of the problems.

In this work we propose an alternative to MLE estimators, which explicitly incorporates the ℓ_0 -penalties but is also tractable and scales to massive instances.

Notations. The i^{th} element of a time-series vector v_t is denoted as $v_{t,i}$. For a vector v , the notation $v_{i:j}$ is used to denote the subvector of v from index i to j . For a vector v , the notations $\|v\|_\infty$, $\|v\|_2$, $\|v\|_0$ denote the ℓ_∞ norm, ℓ_2 norm, and the number of nonzero elements, respectively. Moreover, for a matrix M , the notations $\|M\|_2$, $\|M\|_\infty$, $\|M\|_{1/1}$, $\|M\|_{\infty/\infty}$ refer to the induced 2-norm, induced ∞ -norm, ℓ_1/ℓ_1 norm, and ℓ_∞/ℓ_∞ norm, respectively. Moreover, we define $\|M\|_{\text{off}} = \|M\|_{1/1} - \sum_{i=1}^d |M_{ii}|$. For a vector v and matrix M , the notations $\text{supp}(v)$ and $\text{supp}(M)$ are defined as the sets of their nonzero elements. Given two sequences $f(n)$ and $g(n)$ indexed by n , the notation $f(n) \lesssim g(n)$ implies that there exists a constant $C < \infty$ that satisfies $f(n) \leq Cg(n)$. Finally, the notation $f(n) \asymp g(n)$ implies that $f(n) \lesssim g(n)$ and $g(n) \lesssim f(n)$.

Due to space restrictions, all proofs are deferred to the supplementary file.

1.1 Warm-up: Regularized MLE for Sparsely-changing Gaussian MRFs

To illustrate the fundamental drawbacks of the regularized MLE, first we consider the class of time-varying Gaussian MRFs (GMRF) with sparsely-changing structure. The parameters of time-varying GMRFs can be inferred by estimating a sequence of inverse covariance matrices (also known as precision matrices) $\{\Theta_t\}_{t=0}^T$, where both Θ_t and $\Theta_t - \Theta_{t-1}$ are sparse. Given these precision matrices, the edge set of the Markov matrix \mathcal{G}_t coincides with the off-diagonal nonzero elements of Θ_t [40]. The sparse precision matrices can be estimated via the following regularized MLE, also known as *time-varying Graphical Lasso (GL)* [15, 11]:

$$\begin{aligned} \{\hat{\Theta}_t\}_{t=0}^T = \arg \min_{\Theta_t} \sum_{t=0}^T & \left(\langle \Theta_t, \hat{\Sigma}_t \rangle - \log \det(\Theta_t) \right) \\ & + \gamma_1 \sum_{t=0}^T \|\Theta_t\|_{\text{off}} + \gamma_2 \sum_{t=1}^T \|\Theta_t - \Theta_{t-1}\|_{1/1} \end{aligned} \quad (1a)$$

$$\text{s.t. } \Theta_t \succ 0 \quad t = 0, 1, \dots, T \quad (1b)$$

where $\hat{\Sigma}_t \in \mathbb{R}^{d \times d}$ is the sample covariance matrix. Without loss of generality and to streamline the presentation, we assume that the samples have zero mean. The next example shows that the time-varying GL may lead to poor estimates.

Example 1. Consider a scenario where $\{\Theta_t\}_{t=0}^4 \in \mathbb{R}^{25 \times 25}$ are randomly generated symmetric and sparse matrices. At each time $t = 0, \dots, 4$, the precision matrix Θ_t has exactly 100 off-diagonal elements with value one in its upper-triangular part, and the remaining off-diagonal entries are set to zero. Moreover, the diagonal entries $\Theta_{t,ii}$ are chosen as $1 + \sum_{j \neq i} \Theta_{t,ij}$. At every time, 10 nonzero off-diagonal elements are changed to zero, and 10 zero elements are set to one. To generate the sample covariance matrix $\hat{\Sigma}_t$, each element of the true inverse covariance matrix Σ_t^{-1} is perturbed with a noise value, uniformly chosen from $[-0.1, 0.1]$. Figure 1 (left) illustrates a heatmap of the mismatch error, i.e., the total number of mismatches in the sparsity patterns of the true and estimated precision matrices and their differences, for different values




Figure 1: (left) The heatmap of the mismatch error. (right) The true and estimated nonzero elements of the precision matrix.

of the regularization coefficients. It can be seen that after an exhaustive search over the regularization coefficient space, the best achievable mismatch error is in the order of 70. Thus, the estimated parameters reveal little information about the true structure of the time-varying GMRF. Moreover, Figure 1 (right) depicts the concatenation of the nonzero elements in the true precision matrices (dashed red line), and their corresponding values in the estimated matrices (blue curve) at time $t = 1$. It can be seen that, even when the sparsity pattern of the elements is correctly recovered, the estimated nonzero entries are “shrunk” toward zero, incurring a substantial bias.

The above example shows the inferior statistical performance of the time-varying GL as an instance of a regularized MLE method for sparsely-changing GMRFs. In addition to its subpar statistical performance, time-varying GL suffers from expensive computational complexity: a typical numerical solver for the time-varying GL has a *per-iteration* complexity in the order of $\mathcal{O}(Td^3)$, even if it is tailored to a specific class of problems [15, 28, 23]. Solvers with such computational complexity may fall short of practical use in the large-scale settings.

2 Proposed Approach

The proposed framework is based on exact solutions to a class of *tractable* discrete ℓ_0 -problems, thus circumventing bias and other drawbacks of the standard ℓ_1 -approximations, while guaranteeing the scalability of the proposed method. As a general framework, we study the optimization problem:

$$\min (1 - \gamma) \sum_{t=0}^T \|\theta_t\|_0 + \gamma \sum_{t=1}^T \|\theta_t - \theta_{t-1}\|_0 \quad (2a)$$

$$\text{s.t. } \|\theta_t - \tilde{F}^*(\hat{\mu}_t)\|_{\infty/\infty} \leq \lambda_t \quad t = 0, 1, \dots, T \quad (2b)$$

where the optimal solutions $\{\hat{\theta}_t\}_{t=0}^T$ with $\hat{\theta}_t \in \mathbb{R}^p$ are the estimates of the unknown *canonical parameters* of the sparsely-changing MRF, $\{\hat{\mu}_t\}_{t=0}^T$ are the so-called *empirical moment parameters*, and $\tilde{F}^*(\cdot)$ is an *approximate backward mapping* of the model; see §4 for formal definitions. In the context of time-varying GMRFs, the unknown canonical parameters correspond to the vectorization of the upper-triangular part of the precision matrices with $p = d(d+1)/2$. Moreover, the mean parameters correspond to the sample covariance matrices obtained directly from the data, and the approximate backward mapping is a *proxy* of the precision matrix. In particular, we use the *inverse of a soft-thresholded weighted sample covariance matrix* as a proxy; see §5.1.

Our first result establishes a deterministic guarantee on the estimation error of the optimal solution to (2).

Theorem 1 (Estimation error and sparsistency). *Suppose that $0 < \gamma < 1$. For every $t = 0, \dots, T$, define \mathcal{S}_t as the set of indices corresponding to the nonzero elements of the true canonical parameter θ_t^* . Similarly, for every $t = 1, \dots, T$, define \mathcal{D}_t as the set of indices corresponding to the nonzero elements of $\theta_t^* - \theta_{t-1}^*$.*

Assume that

- $\|\theta_t^* - \tilde{F}^*(\hat{\mu}_t)\|_\infty < \lambda_t, \forall 0 \leq t \leq T,$
- $2\lambda_t \leq \min_{i \in \mathcal{S}_t} |\theta_{t,i}^*|, \forall 0 \leq t \leq T,$
- $2\lambda_t + 2\lambda_{t-1} \leq \min_{i \in \mathcal{D}_t} |\theta_{t,i}^* - \theta_{t-1,i}^*|, \forall 0 \leq t \leq T.$

Then, the following statements hold for every $0 \leq t \leq T$:

- (Sparsistency) *We have $\text{supp}(\hat{\theta}_t) = \text{supp}(\theta_t^*)$ and $\text{supp}(\hat{\theta}_t - \hat{\theta}_{t-1}) = \text{supp}(\theta_t^* - \theta_{t-1}^*)$.*
- (Estimation error) *We have*

$$\|\hat{\theta}_t - \theta_t^*\|_\infty \leq 2\lambda_t, \quad \|\hat{\theta}_t - \theta_t^*\|_2 \leq 2\sqrt{|\mathcal{S}_t|}\lambda_t.$$

Theorem 1 presents a set of conditions under which the proposed estimation method achieves sparsistency and small element-wise estimation error. The first condition entails that the true canonical parameter θ_t^* is a feasible solution to (2). The second and third conditions imply that there is a non-negligible gap between the zero and nonzero elements of the true parameters and their temporal changes. Such assumptions are crucial for excluding false negatives from the support of the estimated parameters.

Theorem 1 holds for the general class of sparsely-changing MRFs, provided that an accurate approximation of the backward mapping is available. Such backward mappings have been widely studied for different classes of MRFs, such as GMRFs [41, 13] and Discrete MRFs [37]. We focus on the class of sparsely-changing GMRFs, and show how Theorem 1 can be used to provide end-to-end sample complexity bounds on the inference of sparsely-changing GMRFs in the high-dimensional settings, see §5.1.

Theorem 2 (Computational complexity). *The optimization problem (2) can be solved to optimality in at most $\mathcal{O}(pT^2)$ time and memory on a single thread.*

Theorem 2 shows that the optimization problem (2) can be solved efficiently in practice despite its non-convex nature. Our solution method for (2) relies on the *element-wise decomposability* of (2). In particular, we decompose (2) into smaller subproblems over different coordinates of $\{\theta_t\}_{t=0}^T$. Then, we show that the optimal solution to each subproblem can be obtained by solving a shortest path problem on an auxiliary *weighted directed acyclic graph (DAG)*. The details of our solution method, as well as some improvements on its runtime are presented in §6. Moreover, our algorithm is easily parallelizable, leading to better runtimes in practice.

Example 1 (continued). *Figure 2 depicts the performance of the proposed method, compared to that of the regularized MLE with $\gamma_1 = 0.06$ and $\gamma_2 = 0.08$ (corresponding to the smallest mismatch error) for the instances generated in Example 1. The regularization parameter γ is set to $1/3$. Each element of the true inverse covariance matrix Σ_t^{-1} is perturbed with a noise value, uniformly chosen from $[-\alpha, \alpha]$ for a varying scalar α , and λ_t is set to α . Figure 2 (left) demonstrates that the proposed method enjoys a significantly smaller mismatch error, for different levels of approximation in the backward mapping. In particular, the proposed estimator achieves a zero mismatch error, even for a fairly large approximation error ($\|\theta_t^* - \tilde{F}^*(\hat{\mu}_t)\|_\infty \leq 0.6$). On the other hand, Figure 2 (right) shows that the synthetic bias caused by the regularized MLE is alleviated via the proposed method.*

3 Related Works

We now summarize the works most relevant to our results.

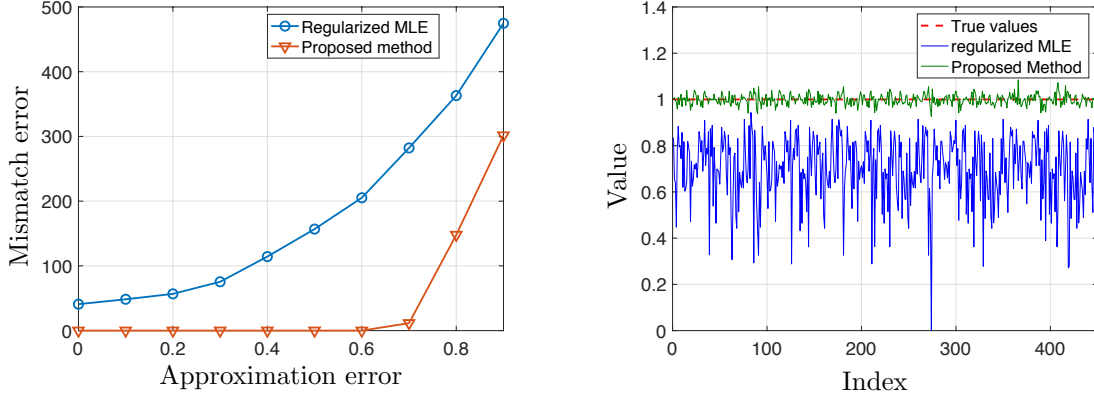


Figure 2: (left) The heatmap of the mismatch error. (right) The true and estimated nonzero elements of the precision matrix.

Inference of time-varying MRF. In addition to the time-varying graphical Lasso introduced in §1.1, a recent line of works have studied the inference of smoothly changing GMRFs [43, 14], where a kernel averaging technique combined with Graphical Lasso is used to estimate the smoothly-changing precision matrices. However, these methods do not leverage the prior information about the sparsity of the parameter differences. With the goal of addressing this deficiency, several works have studied the inference of sparsely-changing MRF (also known as sparse *differential networks*) [39, 42, 24]. However, the main drawback of these methods is that they only estimate the parameter differences, and their theoretical guarantees are restricted to problems with two time steps ($T = 0, 1$).

Sparsity-promoting optimization. Optimization problems with ℓ_0 terms are often deemed to be intractable, and approximations are solved instead. Perhaps the most popular approach is the *fused lasso* [30, 35, 33, 36], which calls for replacing terms $\|\theta_t\|_0$ and $\|\theta_t - \theta_{t-1}\|_0$ with their ℓ_1 -approximations. Nonetheless, such approximations result in subpar statistical performance when compared with exact ℓ_0 methods [18, 26].

Exact or near-optimal methods for optimization problems of the form

$$\min_{\theta \in [\ell, u]^p} \|\theta\|_0 + \sum_{i=1}^p \sum_{j=i+1}^p g_{ij}(\theta_i - \theta_j), \quad (3)$$

for given one-dimensional functions $g_{ij} : \mathbb{R} \rightarrow \mathbb{R}$, have also been studied in the literature. If functions g_{ij} are convex, then problem (3) admits pseudo-polynomial time algorithm [3, 8]. Moreover, convex relaxations that deliver near-optimal solutions for (3) were proposed for the special case of convex quadratic g functions [7]; if, additionally, we have $\ell = 0$ and $u = \infty$, then problem (3) is in fact solvable in strongly polynomial time [6]. On the other hand, problem (3) is much more challenging for non-convex g : if $g(x) = \mathbb{1}\{x \neq 0\}$, as is the case in (2), then problem (3) is NP-hard even if the term $\|\theta\|_0$ is dropped from the objective [16]. Nonetheless, as we show in this paper, problem (3) can be solved quite efficiently in the context of time-varying MRFs, where $g_{ij}(x) = 0$ whenever $j > i + 1$.

4 On Time-varying MRFs and Exponential Families

A large class of time-varying MRFs can be expressed as sequences of exponential distributions, defined as:

$$\mathbb{P}(X_t; \theta_t) = \exp \{ \langle \theta_t, \phi(X_t) \rangle - A(\theta_t) \} \quad \forall 0 \leq t \leq T, \quad (4)$$

where $\theta_t \in \Omega \in \mathbb{R}^p$ is the *canonical parameter* of the exponential distribution at time t , the function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is the *sufficient statistics*, and $A : \mathbb{R}^p \rightarrow \mathbb{R}$ is the *log-partition* function, which is used to

normalize the distribution. Special classes of time-varying MRFs that can be represented as instances of (4) include time-varying Gaussian MRFs (GMRFs) and Discrete MRFs (DMRFs), corresponding to multivariate Gaussian and discrete random processes, respectively. Due to the equivalence between time-varying MRFs and exponential families, Markov graphs can be systematically obtained from the canonical parameters $\{\theta_t\}_{t=0}^T$ [38]. For instance, the canonical parameters in time-varying GMRFs at time t correspond to the tuple of time-varying precision matrix Θ_t and mean vector η_t . Moreover, the edge set of the Markov matrix \mathcal{G}_t coincides with the off-diagonal nonzero elements of Θ_t [40].

An alternative parameterization of exponential families is via *mean* or *moment parameters*, i.e., the moments of the sufficient statistics $\mu_t(\theta_t) = \mathbb{E}_{\theta_t}[\phi(X_t)] \in \mathcal{M}$. Given the canonical parameters $\{\theta_t\}_{t=0}^T$, the mean parameters $\{\mu_t(\theta_t)\}_{t=0}^T$ can be obtained via the *forward mapping* $F : \Omega \rightarrow \mathcal{M}$, where $\mu_t(\theta_t) = F(\theta_t) = \nabla A(\theta_t)$. The conjugate (or Fenchel) duality can be used to define the *backward mapping* $F^* : \mathcal{M} \rightarrow \Omega$ with $\theta_t(\mu_t) = F^*(\mu_t) = \nabla A^*(\mu_t)$, where A^* is the conjugate dual of the log-partition function. In practice, the true mean parameters are rarely available, and should be replaced by their *empirical* versions $\{\hat{\mu}_t\}_{t=0}^T$, where $\hat{\mu}_t = \frac{1}{N_t} \sum_{i=1}^{N_t} \phi(X_t^{(i)})$, and $\{X_t^{(i)}\}_{i=1}^{N_t}$ is the sequence of available data samples at time t .

5 Time-varying GMRFs

Theorem 1 presents a set of deterministic conditions under which the estimates from (2) enjoy zero mismatch error and small estimation error. However, the formulation of (2) is contingent upon the availability of an accurate backward mapping, and a choice of λ_t that satisfies the conditions of Theorem 1. In this section, we show how to efficiently design sample-efficient approximate backward mappings, and select λ_t accordingly for the class of sparsely-changing GMRFs. Moreover, we use the deterministic conditions of Theorem 1 to arrive at a non-asymptotic probabilistic guarantee for the inference of time-varying GMRFs under different prior knowledge on their temporal behavior, such as sparsity and smoothness.

5.1 Sparsely-changing GMRFs

Consider a multivariate zero-mean Gaussian process $\{X_t\}_{t=0}^T$ with distribution

$$\mathbb{P}(X_t) = \exp \left\{ -\frac{1}{2} \langle \Theta_t, X_t X_t^\top \rangle + \langle \eta_t, X_t \rangle - A(\mu_t, \Theta_t) \right\} \quad (5)$$

for $t = 0, 1, \dots, T$ where, without loss of generality, we assumed that the mean is zero. Note that (5) is a special case of the exponential family (4). The canonical parameter, i.e., the precision matrix Θ_t , belongs to the domain $\Omega = \{M \in \mathbb{R}^{d \times d} : M \succ 0\}$ with the effective dimension of $p = d(d+1)/2$. Suppose that at any given time t , a sequence of data samples $\{X_t^{(i)}\}_{i=1}^{N_t}$ is collected from (5). Therefore, the inference of time-varying GMRFs reduces to estimating the time-varying precision matrix Θ_t from the data samples $\{X_t^{(i)}\}_{i=1}^{N_t}$. The forward and backward mappings take the following closed-form expressions: $F(\Theta_t) = \Theta_t^{-1}$ and $F^*(\Sigma_t) = \Sigma_t^{-1}$, where Σ_t is the true covariance matrix at time t . A common choice of approximate backward mapping is $\tilde{F}^*(\hat{\Sigma}_t) = \hat{\Sigma}_t^{-1}$, where $\hat{\Sigma}_t = \frac{1}{N_t} \sum_{i=1}^{N_t} X_t^{(i)} X_t^{(i)\top}$ is the sample covariance matrix. However, in the high-dimensional settings where $p \gg N_t$, this approximate backward mapping is not well-defined, since the sample covariance matrix is highly rank-deficient and not invertible.

To address this issue, [41] propose a *proxy* backward mapping for high-dimensional settings: consider the soft-thresholding operator $\text{ST}_\nu(M) : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$, where $\text{ST}_\nu(M)_{ij} = M_{ij} - \text{sign}(M_{ij}) \min\{|M_{ij}|, \nu\}$ if $i \neq j$, and $\text{ST}_\nu(M)_{ij} = M_{ij}$ if $i = j$. Based on this definition, we choose the approximate backward mapping as $\tilde{F}^*(\hat{\Sigma}_t) = [\text{ST}_\nu(\hat{\Sigma}_t)]^{-1}$. [41] show that $\tilde{F}^*(\hat{\Sigma}_t)$ is well-defined, even in the high dimension, with an appropriate choice of the threshold ν .

We now make two assumptions.

Assumption 1 (Bounded norm). *There exist constant numbers $\kappa_1 < \infty$, $\kappa_2 > 0$, and $\kappa_3 < \infty$ such that*

$$\|\Theta_t\|_\infty \leq \kappa_1, \quad \inf_{w: \|w\|_\infty=1} \|\Sigma_t w\|_\infty \geq \kappa_2, \quad \|\Sigma_t\|_{\infty/\infty} \leq \kappa_3,$$

for every $t = 0, \dots, T$.

Assumption 1 implies that the true covariance matrices and their inverses have bounded norms.

Assumption 2 (Weak sparsity). *Given any $t = 0, \dots, T$, the covariance matrix Σ_t satisfies $\max_i \sum_{j=1}^d |[\Sigma_t]_{ij}|^q \leq s(q, d)$, for some function $s: \mathbb{R} \times \mathbb{Z}_+ \rightarrow \mathbb{R}$ and $0 \leq q < 1$.*

Assuming that $s(0, d) \leq s \ll d$, Assumption 2 reduces to the covariance matrix being sparse. Moreover, in many cases, a sparse inverse covariance matrix leads to weakly sparse covariance matrices. For instance, if Θ_t has a banded structure with small bandwidth, then it is known that the elements of $\Sigma_t = \Theta_t^{-1}$ enjoy exponential decay away from the main diagonal elements [12, 20]. Under such circumstances, one can verify that $s(q, d) \leq \frac{C}{1-\rho^q}$ for some constant $C > 0$ and $\rho < 1$. More generally, a similar statement holds for a class of inverse covariance matrices whose support graphs have large average path length [9, 10]; a large class of inverse covariance matrices with row- and column-sparse structures satisfy this condition.

To streamline the presentation, define $\Theta_t^{\min} = \min_{(i,j) \in \mathcal{S}_t} |[\Theta_t]_{ij}|$ and $\Delta\Theta_t^{\min} = \min_{(i,j) \in \mathcal{D}_t} |[\Theta_t - \Theta_{t-1}]_{ij}|$, where \mathcal{S}_t and \mathcal{D}_t are defined as in Theorem 1. We assume that Θ_t^{\min} and $\Delta\Theta_t^{\min}$ are independent of d and T .

Theorem 3. *Consider a sparsely-changing GMRF and let $\zeta = \max\{\log_d(T+1), 1\}$. Given an arbitrary $\tau > \zeta + 2$, let $N_t \gtrsim s(q, d)^{\frac{2}{1-q}} \tau \log d$. Then, with the approximate backward mapping $\tilde{F}^*(\hat{\Sigma}_t) = [S\mathcal{T}_{\nu_t}(\hat{\Sigma}_t)]^{-1}$, and parameters $\nu_t \asymp \sqrt{\frac{\tau \log d}{N_t}}$ and $\lambda_t \asymp \sqrt{\frac{\tau \log d}{N_t}}$, the estimates $\{\hat{\Theta}_t\}_{t=0}^T$ obtained from (2) satisfy the following statements for all $t = 0, 1, \dots, T$, with probability of at least $1 - 4d^{-\tau+\zeta+2}$.*

- (Sparsistency) We have $\text{supp}(\hat{\Theta}_t) = \text{supp}(\Theta_t^*)$ and $\text{supp}(\hat{\Theta}_t - \hat{\Theta}_{t-1}) = \text{supp}(\Theta_t^* - \Theta_{t-1}^*)$.
- (Estimation error) We have

$$\|\hat{\Theta}_t - \Theta_t^*\|_{\infty/\infty} \lesssim \sqrt{\frac{\tau \log d}{N_t}}, \quad \|\hat{\Theta}_t - \Theta_t^*\|_F \lesssim \sqrt{\frac{\tau |\mathcal{S}_t| \log d}{N_t}}.$$

Theorem 3 is a direct consequence of Theorem 1 and provides, to the best of our knowledge, the first non-asymptotic guarantee on the inference of sparsely-changing GMRFs with an arbitrary length of time horizon T . In particular, it shows that the proposed optimization (2) guarantees small estimation error and zero mismatch error for sparsely-changing GMRFs, provided that N_t scales logarithmically with the dimension of the precision matrices. In the static setting ($T = 0$), the derived bound recovers the existing results on the sample complexity of learning static GMRFs [29, 22, 31].

5.2 Sparsely-and-smoothly-changing GMRFs

In many applications, such as financial markets and motion detection in video frames, the associated graphical model should be learned “on-the-go”, as the data arrives with a continuously changing graphical model. Under such circumstances, one may have access to few (or even one) samples at each time.

Suppose that the precision matrices change smoothly over time. Such smooth changes can be modeled via a continuous function $\Theta(x) : [0, 1] \rightarrow \mathbb{R}^{d \times d}$ with uniformly bounded element-wise second derivatives $[\Theta(x)_{ij}]'' = \frac{d^2 \Theta(x)_{ij}}{dx^2}$, such that $\Theta_t^* = \Theta(t/T)$ [14, 43]. If $\Theta(t) \succeq aI$ for every $t \in [0, 1]$ and some $a > 0$,

then the covariance matrix $\Sigma(t) = \Theta(t)^{-1}$ is well-defined and smooth. Then, the problem of inferring the time-varying GMRF reduces to estimating a sequence of precision matrices $\{\Theta(0), \Theta(1/T), \dots, \Theta(1)\}$ given the samples $X_t \sim \mathcal{N}(0, \Sigma(t/T))$. To alleviate the scarcity of samples, [14, 43] leverage the smoothness of the precision matrices, by taking the weighted average of the samples over time, where the weights are obtained from a nonparametric kernel. In particular, consider the weighted sample covariance matrix $\hat{\Sigma}_t^w$:

$$\hat{\Sigma}_t^w = \sum_{s=0}^t w(s, t) \Sigma_s, \text{ where } w(s, t) = \frac{1}{Th} K\left(\frac{s-t}{Th}\right) \quad (6)$$

and $K(\cdot)$ is a symmetric nonnegative kernel that satisfies a set of mild conditions which hold for most standard kernels, including (truncated) Gaussian kernel. These conditions are delineated in the appendix. The following assumptions are the counterparts of Assumptions 1 and 2 for sparsely-and-smoothly-changing GMRF.

Assumption 3 (Bounded norm). *There exist constant numbers $\kappa_1 < \infty$, $\kappa_2 > 0$, and $\kappa_3 < \infty$ such that*

$$\|\Theta(t/T)\|_\infty \leq \kappa_1, \quad \inf_{w: \|w\|_\infty=1} \|\Sigma(t/T)w\|_\infty \geq \kappa_2, \quad \|\Sigma(t/T)\|_{\infty/\infty} \leq \kappa_3,$$

for every $t \in [0, T]$.

Assumption 4 (Weak sparsity). *The covariance matrix $\Sigma(t/T)$ satisfies $\max_i \sum_{j=1}^d |[\Sigma(t/T)]_{ij}|^q \leq s(q, d)$, for some $0 \leq q < 1$, and every $t \in [0, T]$.*

Under Assumptions 3-4, we present the analog of Theorem 3 for sparsely-and-smoothly-changing GMRFs.

Theorem 4. *Consider a sparsely-and-smoothly-changing GMRF with one sample per time, let $\zeta = \max\{\log_d(T+1), 1\}$, and suppose that the sample covariance matrices are constructed according to (6) with $h \asymp T^{-1/3}$. Given an arbitrary $\tau > \zeta + 2$, let $T \gtrsim s(q, d)^{\frac{3}{1-q}} (\tau \log d)^{3/2}$. Then, with the approximate backward mapping $\tilde{F}^*(\hat{\Sigma}_t) = ST_{\nu_t}(\hat{\Sigma}_t^w)$ and parameters $\nu_t \asymp \frac{\sqrt{\tau \log d}}{T^{1/3}}$ and $\lambda_t \asymp \frac{\sqrt{\tau \log d}}{T^{1/3}}$, the estimates $\{\hat{\Theta}_t\}_{t=0}^T$ obtained from (2) satisfy the following statements for all $T = 0, 1, \dots, T$ with probability of at least $1 - d^{-\tau+\zeta+2}$:*

- (Sparsistency) We have $\text{supp}(\hat{\Theta}_t) = \text{supp}(\Theta_t^*)$ and $\text{supp}(\hat{\Theta}_t - \hat{\Theta}_{t-1}) = \text{supp}(\Theta_t^* - \Theta_{t-1}^*)$.
- (Estimation error) We have

$$\|\hat{\Theta}_t - \Theta_t^*\|_{\infty/\infty} \lesssim \sqrt{\frac{\tau \log d}{T^{1/3}}}, \quad \|\hat{\Theta}_t - \Theta_t^*\|_F \lesssim \sqrt{\frac{\tau |\mathcal{S}_t| \log d}{T^{1/3}}}.$$

Theorem 4 shows how the smoothness assumption on the true covariance matrix can be used to construct the backward mappings using the samples collected during the *entire* time horizon, thereby significantly reducing the sample complexity of learning time-varying GMRFs. In particular, leveraging the smoothness of the covariance matrix can reduce the minimum required number of samples from $\mathcal{O}(T \log d)$ to $\mathcal{O}((\log d)^{1.5})$. On the other hand, Theorem 3 does not impose any lower bound on T , and its estimation error decays faster in terms of the sample size.

6 Solution Method

In this section, we describe the proposed algorithm for solving (2). For the simplicity of notation, we define the lower bound and upper bound vectors l_t and u_t as $l_{t;k} = [\tilde{F}^*(\hat{\mu}_t)]_k - \lambda_t$ and $u_{t;k} = [\tilde{F}^*(\hat{\mu}_t)]_k + \lambda_t$, for every $k = 1, \dots, p$. The following fact plays a key role in our analysis.

Fact 1. An optimal solution $\{\hat{\theta}_t\}_{t=0}^T$ of (2) satisfies for every $k = 1, \dots, p$,

$$\{\hat{\theta}_{t;k}\}_{t=0}^T \in \arg \min_{\{\theta_{t;k}\}_{t=0}^T} (1 - \alpha) \sum_{t=0}^T \mathbb{1}\{\theta_{t;k} \neq 0\} + \alpha \sum_{t=1}^T \mathbb{1}\{\theta_{t;k} - \theta_{t-1;k} \neq 0\} \quad (7a)$$

$$\text{s.t. } l_{t;k} \leq \theta_{t;k} \leq u_{t;k} \quad \forall 0 \leq t \leq T \quad (7b)$$

Fact 1 implies that (2) decomposes into the smaller subproblems (7). Therefore, our main focus is devoted to solving each subproblem independently. To further simplify the notation, we drop the subscript k from (7), whenever it is chosen arbitrarily. Let $\text{OPT}_{i \rightarrow j}(\alpha)$ denote, for a given k , the truncated problem from index i to j with the regularization coefficient α , defined as

$$f_{i \rightarrow j}^* = \min_{\{\theta_t\}_{t=i}^j} (1 - \alpha) \sum_{t=i}^j \mathbb{1}\{\theta_t \neq 0\} + \alpha \sum_{t=i+1}^j \mathbb{1}\{\theta_t - \theta_{t-1} \neq 0\} \quad (8a)$$

$$\text{subject to } l_t \leq \theta_t \leq u_t \quad \forall i \leq t \leq j. \quad (8b)$$

Let the objective function for a candidate solution θ be denoted as $f_{i \rightarrow j}(\theta)$; by convention, we let $f_{i \rightarrow j}(\theta) = 0$ whenever $j < i$. Moreover, the optimal objective value and the set of optimal solutions to $\text{OPT}_{i \rightarrow j}(\alpha)$ are respectively denoted as $f_{i \rightarrow j}^*$ and $\mathcal{X}_{i \rightarrow j}^*$. Similarly, $\hat{\theta}_{i \rightarrow j} \in \mathcal{X}_{i \rightarrow j}^*$ is used to denote an optimal solution to $\text{OPT}_{i \rightarrow j}(\alpha)$. We omit the subscript $i \rightarrow j$ whenever $i = 0$ and $j = T$. The t^{th} feasible interval is defined as $\Delta_t = [l_t, u_t]$. Accordingly, the notation $\Delta_{t \rightarrow s}^\cap$ refers to $\Delta_{t \rightarrow s}^\cap \stackrel{\text{def}}{=} \Delta_t \cap \Delta_{t+1} \cap \dots \cap \Delta_s$.

6.1 Special case: $\alpha = 1$

As the first step, we consider the special case $\alpha = 1$, and provide an efficient algorithm (Algorithm 1) for solving $\text{OPT}_{0 \rightarrow T}(1)$, where the sparsity is only promoted on the parameter differences (and not on the individual parameters). As will be shown later, Algorithm 1 will be used as a subroutine in our proposed algorithm for the general case $0 < \alpha < 1$. At a high level, Algorithm 1 recursively performs the following operations: at any given time τ , the algorithm looks into the future to find a nonempty interval that is feasible for the longest possible time δ . Then, it sets the subvector $\theta_{\tau:\delta}$ to an arbitrarily chosen element from this nonempty interval.

Algorithm 1 Greedy(l, u, τ, T)

- 1: **Output:** Solution $\theta_{\tau \rightarrow T}^{\text{Greedy}}$, the objective value $f_{\tau \rightarrow T}^{\text{Greedy}}$ to $\text{OPT}_{\tau \rightarrow T}(1)$, and the index set Γ of maximal nonempty intervals
 - 2: Find largest δ such that $\Delta_{\tau \rightarrow \delta}^\cap \neq \emptyset$;
 - 3: Set $\theta_{\tau:\delta}^{\text{Greedy}} = \eta$ for some $\eta \in \Delta_{\tau \rightarrow \delta}^\cap$;
 - 4: **if** $\delta \leq T - 1$ **then**
 - 5: Execute Greedy($l, u, \delta + 1, T$);
 - 6: **end if**
 - 7: Set $f_{\tau \rightarrow T}^{\text{Greedy}} = \sum_{t=\tau+1}^T \mathbb{1}\{\theta_t^{\text{Greedy}} - \theta_{t-1}^{\text{Greedy}} \neq 0\}$;
 - 8: Set $\Gamma \leftarrow \Gamma \cup \{\delta\}$;
 - 9: **Return** $\{\theta_t^{\text{Greedy}}\}_{t=\tau}^T$, $f_{\tau \rightarrow T}^{\text{Greedy}}$, and Γ ;
-

Proposition 1. Greedy($l, u, 0, T$) returns an optimal solution $\{\theta_t^{\text{Greedy}}\}_{t=0}^T$ to $\text{OPT}_{0 \rightarrow T}(1)$. Moreover, the truncated solution $\{\theta_t^{\text{Greedy}}\}_{t=0}^j$ is optimal for $\text{OPT}_{0 \rightarrow j}(1)$.

6.2 General case: $0 < \alpha < 1$

Now, we present our main algorithm for the general case $0 < \alpha < 1$. To this goal, we first present the following definition.

Definition 1. The set $\mathcal{Z}_{i \rightarrow j} = \{i, i+1, \dots, j\}$ is called a **zero-feasible sequence** if $l_k \leq 0 \leq u_k$ for every $k \in \mathcal{Z}_{i \rightarrow j}$. Moreover, the zero-feasible sequence $\mathcal{Z}_{i \rightarrow j}$ is called **maximal** if it is not strictly contained within another zero-feasible sequence.

Let $\mathcal{Z}_{i_1 \rightarrow j_1}, \mathcal{Z}_{i_2 \rightarrow j_2}, \dots, \mathcal{Z}_{i_Z \rightarrow j_Z}$ be the set of all maximal zero-feasible sequences such that $0 \leq i_1 \leq j_1 < i_2 \leq j_2 < \dots < i_Z \leq j_Z \leq T$, where Z is the number of maximal zero-feasible sequences. If $Z = 0$, i.e., there is no zero-feasible sequence, then it is easy to see that $\sum_{t=0}^T \mathbb{1}\{\theta_t \neq 0\} = T+1$ for every feasible solution, and hence, $\text{Greedy}(l, u, 0, T)$ leads to an optimal solution to (8). Another special case is when $i_1 = 0$ and $j_1 = T$, i.e., zero is always feasible. In this case, the optimal solution is $\hat{\theta} = 0$. Therefore, without loss of generality, suppose that $Z > 0$ and either $i_1 \neq 0$ or $j_1 \neq T$. As will be shown later, the zero-feasible sequences play an important role in characterizing the optimal solution of (8).

Our goal now is to obtain an optimal solution to (8) by solving a shortest path problem over a *weighted directed acyclic graph* (DAG) whose nodes correspond to the maximal zero-feasible sequences. In particular, consider a weighted DAG \mathcal{G} with the vertex set $V = \{0, 1, \dots, Z, Z+1\}$, where the vertices k and l are connected via a directed arc (k, l) if $k < l$. Moreover, for every arc (k, l) , the weight $W(k, l) = 0$ if $(k, l) = (0, 1), i_1 = 0$ or $(k, l) = (Z, Z+1), j_Z = T$, and

$$W(k, l) = (1 - \alpha)(i_l - j_k - 1) + \alpha f_{j_k+1 \rightarrow i_l-1}^{\text{Greedy}} + \alpha \mathbb{1}\{k \neq 0\} + \alpha \mathbb{1}\{l \neq T+1\} \quad (9)$$

otherwise, where we define $j_0 = -1$ and $i_{T+1} = T+1$.

Algorithm 2 Algorithm for solving (8)

- 1: **Output:** Optimal solution $\hat{\theta}$ and the objective value f^* to $\text{OPT}_{0 \rightarrow T}(\alpha)$
 - 2: Find the maximal zero-feasible sequences;
 - 3: Construct the DAG \mathcal{G} with weights defined as (9);
 - 4: Find the shortest path $p = (v_1, v_2, \dots, v_r)$ between the vertices 0 and $Z+1$ in \mathcal{G} ;
 - 5: Set $\hat{\theta}_{j_{v_i}+1:i_{v_{i+1}}-1} = \theta_{j_{v_i}+1 \rightarrow i_{v_{i+1}}-1}^{\text{Greedy}}$ and $\hat{\theta}_{i_{v_l}:j_{v_{l+1}}} = 0$ for every $l = 1, 2, \dots, r$;
 - 6: **Return** $\{\hat{\theta}\}_{t=0}^T$ and f^*
-

Theorem 5. The shortest path from 0 to $Z+1$ on \mathcal{G} has value f^* .

The above theorem implies that the optimal solution to (2) can be obtained via Algorithm 2.

Theorem 6. Problem (8) can be solved in $\mathcal{O}(ZT)$ time and memory.

Since $Z = \mathcal{O}(T)$ and solving (2) requires solving $\mathcal{O}(p)$ instances of (8), we find the total complexity stated in Theorem 2. Note however that if $Z = \mathcal{O}(1)$ for every instance of (2), then the overall complexity reduces to $\mathcal{O}(pT)$, which is linear in the total number of variables. In the next section, we will show that the practical runtime of the proposed algorithm is near-linear with respect to the number of variables.

7 Numerical Analysis

In this section, we provide detailed information about the performance of the proposed estimator in different case studies. In the first case study, our goal is to compare the statistical performance of our proposed

method with two other state-of-the-art methods, namely time-varying Graphical Lasso [15], and a modified version of the elementary ℓ_1 estimator [39, 41]. We will show that the proposed estimator outperforms the other two estimators, in terms of both sparsity recovery and estimation error. In the second case study, we showcase the statistical and computational performance of the proposed method on massive-scale datasets. In particular, we will show that our proposed estimation method can solve instances of the problem with more than 500 million variables in less than one hour, with almost perfect sparsity recovery. Moreover, we demonstrate the improvements in the runtime of our algorithm with parallelization. Finally, we conduct a case study on the correlation network inference in stock markets. In particular, we show that the inferred time-varying graphical model can correctly identify the stock market spikes based on the historical data.

All simulations are run on a desktop computer with an Intel Core i9 3.50 GHz CPU and 128GB RAM. The reported results are for an implementation in MATLAB R2020b.

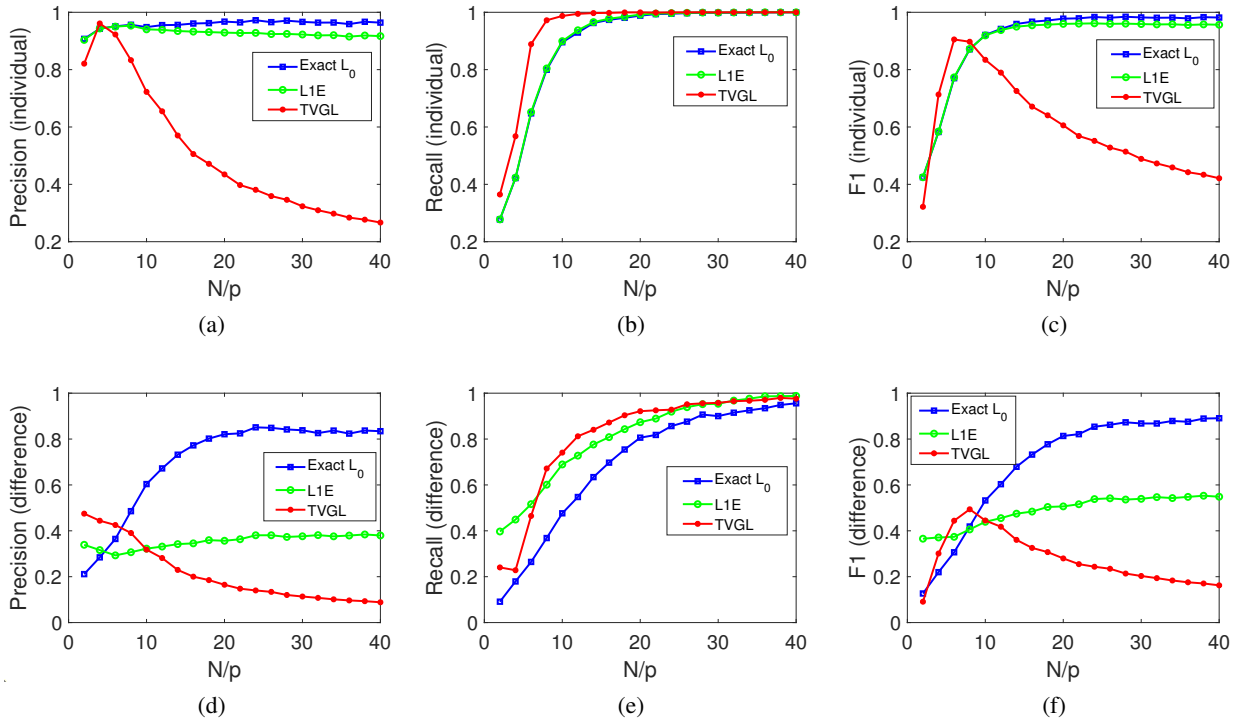


Figure 3: Precision, Recall, and F1-score for the estimated precision matrices and their differences using the proposed method (denoted as Exact L_0), L1E, and TVGL (averaged over 10 independent trials).

7.1 Case Study on Small Datasets

In this case study, we evaluate the statistical performance of the proposed estimator, compared to two other methods, namely time-varying Graphical Lasso (TVGL) [15, 11], and a modified version of the elementary ℓ_1 estimator (L1E) introduced in [39, 41]. As mentioned in the introduction, TVGL is a well-known regularized MLE approach for estimating the sparsely-changing GMRFs. On the other hand, different variants of L1E have been used to estimate static MRFs [41], and differential networks with sparsity imposed only on the parameter differences [39]. Consider an ℓ_1 relaxation of the proposed estimator (2), where the ℓ_0 penalties in the objective function are replaced with ℓ_1 penalties. The resulted estimator reduces to that of [41] for $T = 0$, and [39] for $T = 1$ and $\gamma = 1$.




Figure 4: The normalized ℓ_∞ -norm and induced 2-norm of the estimation error for the estimated precision matrices and their differences using the proposed method, L1E, and TVGL (averaged over 10 independent trials).

We consider randomly generated instances of sparsely-changing GMRFs, where the true inverse covariance matrix is constructed as follows: at time $t = 0$, we set $\Theta_0 = I_{d \times d} + \sum_{(i,j) \in \mathcal{S}} A^{(i,j)}$, where $d = 50$ and $A^{(i,j)}$ is a sparse positive semidefinite matrix with exactly two nonzero off-diagonal elements. In particular, we randomly select 100 edges in the graph (corresponding to 200 off-diagonal entries in Θ_0) and collect their indices in \mathcal{S} . For every $(i, j) \in \mathcal{S}$, we set $A_{ij}^{(i,j)} = A_{ji}^{(i,j)} = -0.4$ and $A_{ii}^{(i,j)} = A_{jj}^{(i,j)} = 0.4$. Clearly, $A^{(i,j)} \succeq 0$, and hence, $\Theta_0 \succ 0$. Moreover, at every time $t = 1, \dots, 9$, exactly 20 nonzero off-diagonal entries are added to Θ_t according to the aforementioned rules, and 20 nonzero off-diagonal entries are deleted by reversing the above procedure. Our goal is to estimate the true sparsely-changing precision matrices $\{\Theta_t\}_{t=0}^9$ based on a varying number of samples N_t . We evaluate the accuracy of the different methods in terms of Recall, Precision, and F1-score values, defined as

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad \text{F1-score} = \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}}, \quad (10)$$

where TP, FP, and FN respectively denote the number of true positives, false positives, and false negatives in the estimated sequence of precision matrices. In all of our experiments, we set $\alpha = 0.7$. Moreover, according to Theorem 3, we set $\lambda_t = C_1 \frac{\log d}{N_t}$ and $\nu_t = C_2 \frac{\log d}{N_t}$ for the proposed method and L1E, where C_1 and C_2 are constants that are fine-tuned for an instance of the problem, and are kept unchanged throughout the simulations. In particular, we perform an exhaustive search over the constants C_1 and C_2 , and select the ones that achieve the smallest estimation error for a random instance of the problem. Similarly, we set the regularization coefficients $\gamma_1 = C_3 \frac{\log d}{N_t}$ and $\gamma_2 = C_4 \frac{\log d}{N_t}$ for TVGL (1), where the constants C_3 and C_4 are selected in a similar fashion.

Figure 3 illustrates the accuracy of the estimated precision matrices for different number of samples. It can be seen that the proposed estimator outperforms L1E and TVGL in terms of Precision value, but has a slightly worse Recall value. In particular, both L1E and TVGL tend to *overestimate* the number of nonzero elements in the precision matrices. This overestimation naturally reduces the number of false negatives (leading to better Precision values), while significantly increasing the number of false positives (leading to worse Recall values). Moreover, F1-score shows the overall performance of the estimates in terms of the sparsity recovery. It can be seen that the proposed estimator outperforms the other two methods. In particular, both L1E and TVGL perform poorly on the sparsity recovery of the parameter

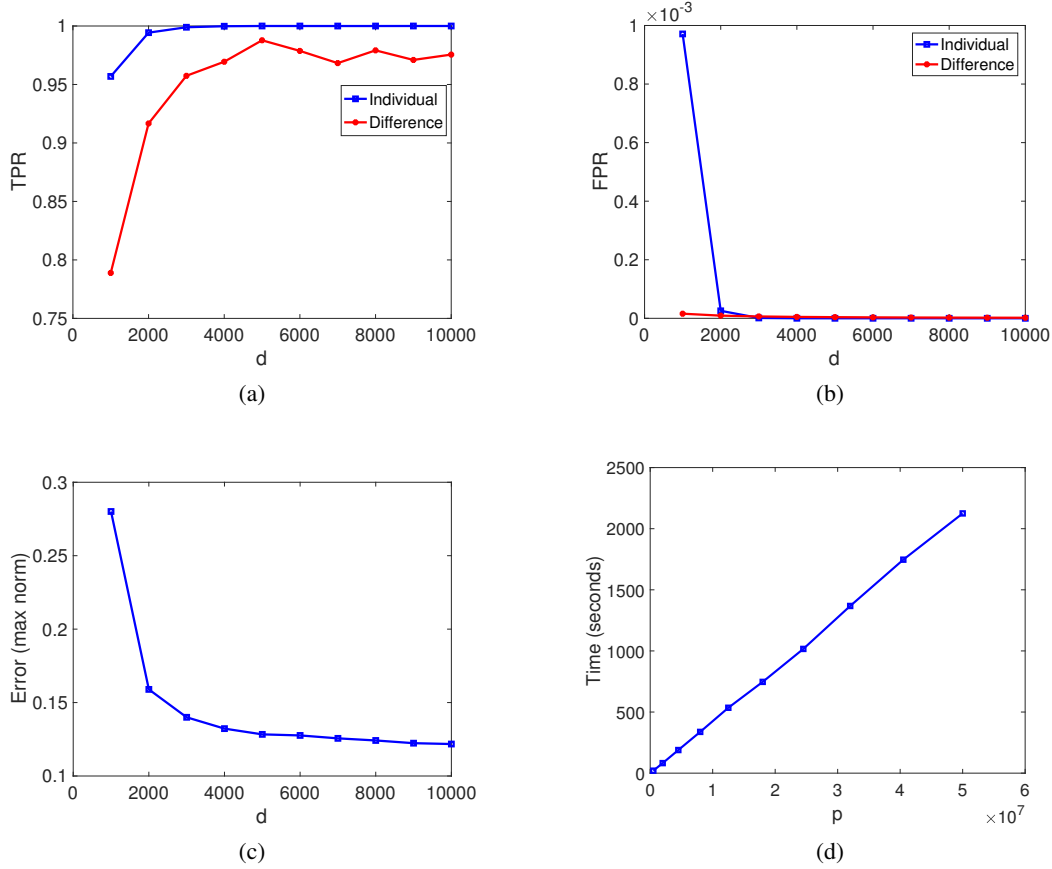


Figure 5: TPR, FPR, ℓ_1 -norm estimation error, and the runtime of the proposed method for fixed T and different values of d . The number of samples N_t is set to $d/2$ for every t . The runtime is shown with respect to $p = d(d + 1)/2$.

differences. Finally, Figure 4 depicts the normalized ℓ_∞ -norm and induced 2-norm estimation errors. It can be seen that TVGL incurs a relatively large ℓ_∞ -norm error due to the shrinking effect of its regularization.

7.2 Case Study on Large Datasets

In this case study, we analyze the performance of the proposed estimator on large datasets, with different values of d and T . In particular, we will analyze the runtime of the proposed algorithm and its statistical performance in high dimensional settings, where $N_t < d$ for every $t = 1, 2, \dots, T$. Moreover, we will report the improvements in the runtime with parallelization, and analyze the robustness of the estimator for different choices of the regularization coefficient α .

Consider the class of synthetically generated sparsely-changing GMRFs with random precision matrices, as explained in Subsection 7.1. In the first experiment, we fix $T = 10$ and change the values of d . The number of nonzero elements in the individual precision matrices and their differences are set to $3d$ and $0.04d$, respectively. We evaluate the performance of the proposed method in the high dimensional settings, where $N_t = d/2$ for every $t = 0, \dots, T$. The parameters λ_t and ν_t are fine-tuned similar to the previous case study and $\alpha = 0.7$ in all instances. Moreover, define TPR and FPR for the individual parameters and their differences as the TP and FP values, normalized by the total number of nonzero and zero elements in the true precision matrices and their differences, respectively. Clearly, both TPR and FPR are between 0




Figure 6: (a) The runtime of the parallelized algorithm with respect to the number of variables Tp , for different number of cores. (b) The normalized mismatch error with respect to the regularization coefficient α , for the choices of parameters $d = 4000$, $T = 10$, and $N_t = 2000$ for every t .




Figure 7: TPR, FPR, ℓ_1 -norm estimation error, and the runtime of the proposed method for fixed d and different values of T . The number of samples N_t is set to $2d$ for every t .

and 1, with $\text{TPR} = 1$ and $\text{FPR} = 0$ corresponding to the perfect recovery of the sparsity patterns. Figure 5 depicts TPR, FPR, and the ℓ_1 -norm error of the estimated parameters, as well as the runtime of our algorithm for different values of d . It can be seen that both TPR and FPR values improve with the dimension for the estimated parameters and their differences. Moreover, the runtime of our algorithm scales almost linearly with $p = d(d + 1)/2$, which is in line with the result of Theorem 2. Using our algorithm, we reliably infer instances of sparsely-changing GMRFs with more than 500 million variables in less than one hour.

As mentioned before, our proposed optimization framework is amenable to parallelization due to its elementwise decomposable nature. Figure 6a illustrates the runtime of our parallelized algorithm with respect to the total number of variables (fixed T and varying p), for different number of cores. Using 5 cores, the runtime of our algorithm is improved by 40% on average. On the other hand, using 10 cores deteriorates the performance due to the shared memory limitations. Finally, we evaluate the accuracy of the estimated parameters for different choices of the regularization coefficient α . In particular, we fix $d = 4000$, $T = 10$, and $N_t = 2000$ for every t , and depict the normalized mismatch error in the sparsity pattern of the estimated parameters and their differences for $\alpha \in \{0.1, 0.2, \dots, 1\}$. Based on this figure, it can be concluded that overall performance of the proposed method is not too sensitive to specific choice of the regularization parameter α . In particular, it can be seen that the normalized mismatch error remains

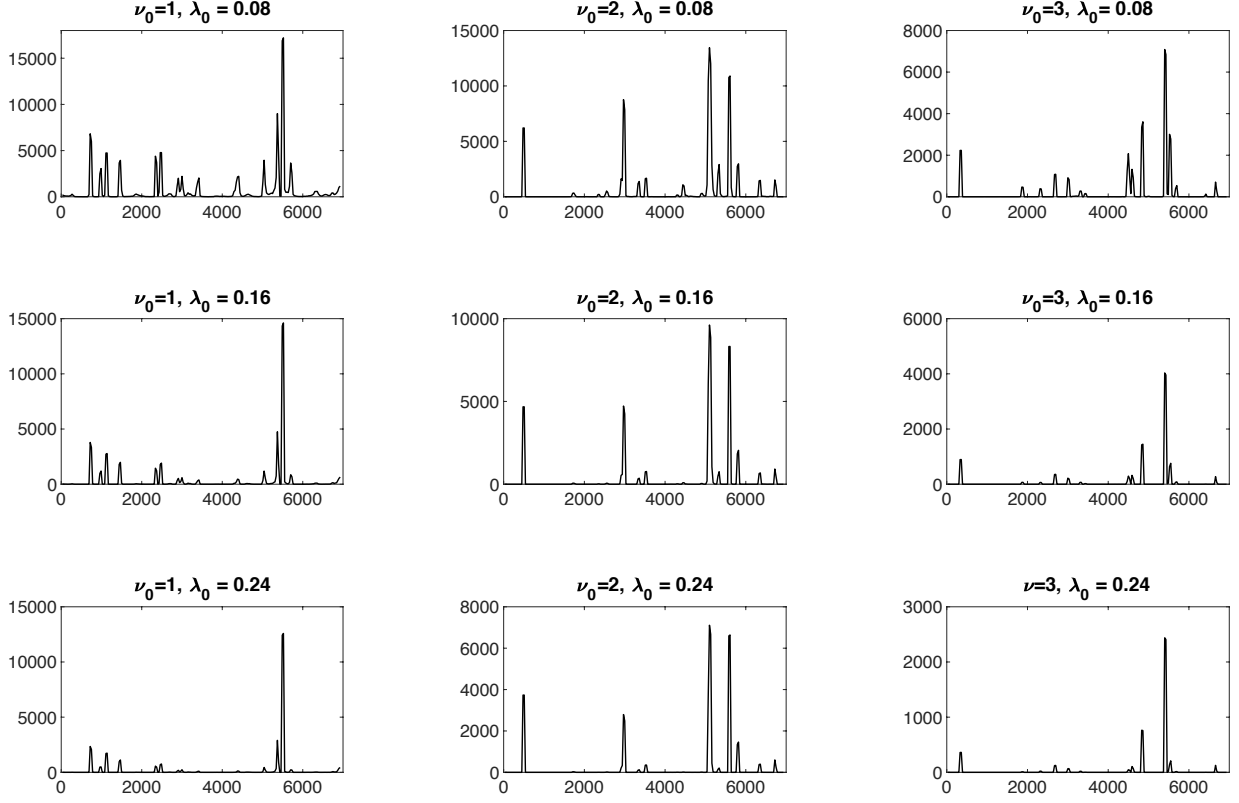


Figure 8: The number of changes in the estimated stock correlation network, for different choices of ν_0 and λ_0 . The x -axis represent the day indexes.

approximately the same for $\alpha \in [0.5, 0.8]$.

In the next experiment, we set $d = 1000$ and $N_t = 2d$, and evaluate the performance of the proposed method for different values of $T \in \{10, 20, 30, \dots, 200\}$. Figure 7a shows TPR for the estimated precision matrices and their differences. It can be seen that TPR for the estimated precision matrices is close to 1 for all values of T . Moreover, the TPR for the differences of the estimated precision matrices is at least 0.966. On the other hand, Figure 7b shows that the FPR for the estimated precision matrices is close to zero. Finally, Figure 7c shows that the runtime of the proposed algorithm scales almost linearly with T . Together with Figure 5d, this implies that the empirical complexity of the algorithm is linear in both p and T .

7.3 Case Study on Stock Market

Finally, we illustrate the performance of our algorithm for the inference of stock correlation network. We consider the daily stock prices for 214 securities from 1990/01/04 to 2017/08/10, with the total number of 6990 days ($d = 214$ and $T = 6990$). Due to the continuously changing nature of the stock correlation network, we will use the kernel averaging approach that was introduced in Subsection 5.1 to estimate the underlying time-varying network. In particular, we consider a Gaussian kernel with bandwidth $h = 0.3T^{-1/3}$ to obtain the sequence of weighted sample covariance matrices. Using the constructed sample covariance matrices, we estimate the sparsely-changing precision matrix $\Theta(t/T)$ at discrete times $t \in \{30, 60, 90, \dots, 6990\}$. Moreover, we set $\alpha = 0.9$, $\lambda_t = \lambda_0 \sqrt{\frac{\log(d)}{T_h}}$, and $\nu_t = \nu_0 \sqrt{\frac{\log(d)}{T_h}}$, for some constants λ_0 and ν_0 to be defined later. Note that these choices of the parameters are consistent with the

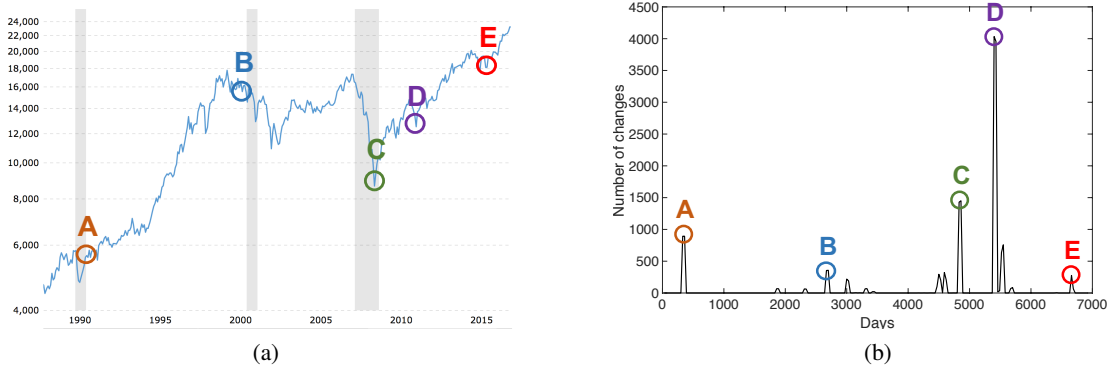


Figure 9: (a) NASDAQ historical chart from 1988 to 2017 [2]. (b) The number of changes in the estimated correlation network for $\nu_0 = 3$ and $\lambda_0 = 0.16$.

assumptions of Theorem 4.

Figure 8 shows the number of changes in the sparsity pattern of the estimated correlation network, for different choices of the parameters ν_0 and λ_0 . A drastic change in the correlation network signals a *spike* in the stock market, which may reflect the market’s response to unexpected global events. It can be seen that, for small values of ν_0 and λ_0 , the estimated network can detect both small and large spikes. As the values of ν_0 and λ_0 increase, the small spikes gradually diminish, and the estimated network only “picks up” major changes in the network. Nonetheless, there is a recurring pattern of spikes in these plots that is almost insensitive to different values of ν_0 and λ_0 . A closer look at this recurring pattern sheds light on the behavior of the market. Figure 9 shows the number of changes in the estimated network, for the choices of $\nu_0 = 3$ and $\lambda_0 = 0.16$, together with the historical chart of National Association of Securities Dealers Automated Quotations (NASDAQ) [1]. It can be seen that the major spikes in the estimated network can be attributed to the historical stock market *crashes*. For instance, the spikes A, B, and C respectively correspond to the “early 1990s recession”, “dot-com bubble”, and “global financial crisis”; see [5] for more details. Interestingly, the estimated network can also detect other historical (but less severe) downturns in 2011 (point D) and 2016 (point E).

8 Conclusion

In this paper, we study the inference of sparsely-changing Markov random fields (MRF), where the goal is to estimate a sequence of time-varying Markov graphs from a limited number of samples, while promoting sparsity of the individual Markov graphs and their differences. We introduce a new class of constrained optimization problems which, unlike existing alternatives, is based on the exact ℓ_0 regularization. We provide an efficient algorithm that can solve the proposed optimization problem to optimality in polynomial time. The proposed estimator benefits from strong statistical guarantees. As a special case, we show that, using our method, the sparsely-changing Gaussian MRFs can be reliably estimated with as few as one sample per time. Finally, we illustrate the performance of the proposed method in different case studies on synthetic data and financial markets: problems with more than 500 million unknown parameters can be solved in less than one hour.

References

- [1] About nasdaq. <https://www.nasdaq.com/about>, 2020.
- [2] Nasdaq composite - 45 year historical chart. <https://www.macrotrends.net/1320/nasdaq-historical-chart>, 2020.
- [3] R. K. Ahuja, D. S. Hochbaum, and J. B. Orlin. A cut-based algorithm for the nonlinear dual of the minimum cost network flow problem. *Algorithmica*, 39(3):189–208, 2004.
- [4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows*. 1988.
- [5] R. Z. Aliber and C. P. Kindleberger. *Manias, panics, and crashes: A history of financial crises*. Springer, 2017.
- [6] A. Atamtürk and A. Gómez. Strong formulations for quadratic optimization with m-matrices and indicator variables. *Mathematical Programming*, 170(1):141–176, 2018.
- [7] A. Atamtürk, A. Gómez, and S. Han. Sparse and smooth signal estimation: Convexification of ℓ_0 -formulations. *arXiv preprint arXiv:1811.02655*, 2018.
- [8] F. Bach. Submodular functions: from discrete to continuous domains. *Mathematical Programming*, 175(1-2):419–459, 2019.
- [9] M. Benzi and N. Razouk. Decay bounds and $\mathcal{O}(n)$ algorithms for approximating functions of sparse matrices. *Electron. Trans. Numer. Anal*, 28:16–39, 2007.
- [10] M. Benzi and V. Simoncini. Decay bounds for functions of hermitian matrices with banded or kronecker structure. *SIAM Journal on Matrix Analysis and Applications*, 36(3):1263–1282, 2015.
- [11] B. Cai, G. Zhang, A. Zhang, J. M. Stephen, T. W. Wilson, V. D. Calhoun, and Y.-P. Wang. Capturing dynamic connectivity from resting state fmri using time-varying graphical lasso. *IEEE Transactions on Biomedical Engineering*, 66(7):1852–1862, 2018.
- [12] S. Demko, W. F. Moss, and P. W. Smith. Decay rates for inverses of band matrices. *Mathematics of computation*, 43(168):491–499, 1984.
- [13] S. Fattahi and S. Sojoudi. Graphical lasso and thresholding: Equivalence and closed-form solutions. *The Journal of Machine Learning Research*, 20(1):364–407, 2019.
- [14] K. Greenewald, S. Park, S. Zhou, and A. Giessing. Time-dependent spatially varying graphical models, with application to brain fmri data analysis. In *Advances in Neural Information Processing Systems*, pages 5832–5840, 2017.
- [15] D. Hallac, Y. Park, S. Boyd, and J. Leskovec. Network inference via the time-varying graphical lasso. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 205–213, 2017.
- [16] D. S. Hochbaum. An efficient algorithm for image segmentation, markov random fields and related problems. *Journal of the ACM (JACM)*, 48(4):686–701, 2001.
- [17] S. Huang, J. Li, L. Sun, J. Ye, A. Fleisher, T. Wu, K. Chen, and E. Reiman. Learning brain connectivity of alzheimer’s disease by sparse inverse covariance estimation. *NeuroImage*, 50(3):935–949, 2010.

- [18] S. Jewell and D. Witten. Exact spike train inference via ℓ_0 optimization. *arXiv preprint arXiv:1703.08644*, 2017.
- [19] R. E. Kalman. A new approach to linear filtering and prediction problems. 1960.
- [20] D. Kershaw. Inequalities on the elements of the inverse of a certain tridiagonal matrix. *Mathematics of computation*, pages 155–158, 1970.
- [21] J. Kim and W. Pan. Highly adaptive tests for group differences in brain functional connectivity. *NeuroImage: Clinical*, 9:625–639, 2015.
- [22] C. Lam and J. Fan. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of statistics*, 37(6B):4254, 2009.
- [23] B. Le Bars, P. Humbert, A. Kalogeratos, and N. Vayatis. Learning the piece-wise constant graph structure of a varying ising model. *International Conference on Machine Learning*, 2020.
- [24] S. Liu, K. Fukumizu, and T. Suzuki. Learning sparse structural changes in high-dimensional markov networks. *Behaviormetrika*, 44(1):265–286, 2017.
- [25] X. Liu and J. H. Duyn. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, 110(11):4392–4397, 2013.
- [26] A. Miller. *Subset Selection in Regression*. CRC Press, 2002.
- [27] M. Narayan, G. I. Allen, and S. Tomson. Two sample inference for populations of graphical models with applications to functional connectivity. *arXiv preprint arXiv:1502.03853*, 2015.
- [28] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty. High-dimensional ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- [29] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- [30] A. Rinaldo et al. Properties and refinements of the fused lasso. *The Annals of Statistics*, 37:2922–2952, 2009.
- [31] A. J. Rothman, P. J. Bickel, E. Levina, J. Zhu, et al. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- [32] M. Rubinov and O. Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.
- [33] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992.
- [34] J. C. Shellenbarger. Estimation of covariance parameters for an adaptive kalman filter. 1966.
- [35] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67:91–108, 2005.
- [36] C. R. Vogel and M. E. Oman. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17:227–238, 1996.

- [37] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. Tree-reweighted belief propagation algorithms and approximate ml estimation by pseudo-moment matching. In *AISTATS*, volume 3, page 3, 2003.
- [38] M. J. Wainwright and M. I. Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [39] B. Wang, Y. Qi, et al. Fast and scalable learning of sparse changes in high-dimensional gaussian graphical model structure. In *International Conference on Artificial Intelligence and Statistics*, pages 1691–1700. PMLR, 2018.
- [40] Y. Weiss and W. T. Freeman. Correctness of belief propagation in gaussian graphical models of arbitrary topology. In *Advances in Neural Information Processing Systems*, pages 673–679, 2000.
- [41] E. Yang, A. C. Lozano, and P. K. Ravikumar. Elementary estimators for graphical models. In *Advances in neural information processing systems*, pages 2159–2167, 2014.
- [42] S. D. Zhao, T. T. Cai, and H. Li. Direct estimation of differential networks. *Biometrika*, 101(2):253–268, 2014.
- [43] S. Zhou, J. Lafferty, and L. Wasserman. Time varying undirected graphs. *Machine Learning*, 80(2-3):295–319, 2010.

A Proofs

A.1 Proof of Theorem 1

Due to the feasibility of $\{\hat{\theta}_t\}_{t=0}^T$, one can write $\|\hat{\theta}_t - \tilde{F}^*(\hat{\mu}_t)\|_\infty \leq \lambda_t$. Combined with the first assumption of the theorem, this implies that

$$\begin{aligned} \|\hat{\theta}_t - \theta_t^*\|_\infty &= \|\hat{\theta}_t - \tilde{F}^*(\hat{\mu}_t) + \tilde{F}^*(\hat{\mu}_t) - \theta_t^*\|_\infty \\ &\leq \|\hat{\theta}_t - \tilde{F}^*(\hat{\mu}_t)\|_\infty + \|\theta_t^* - \tilde{F}^*(\hat{\mu}_t)\|_\infty \\ &< 2\lambda_t, \end{aligned} \tag{11}$$

thereby establishing the element-wise estimation error bound. We proceed to show the sparsistency of the estimated parameters. First, suppose that $\theta_{t,i}^* \neq 0$ for some time t and index i . One can write

$$\begin{aligned} |\hat{\theta}_{t,i}| &= |\hat{\theta}_{t,i} - \theta_{t,i}^* + \theta_{t,i}^*| \\ &\geq |\theta_{t,i}^*| - |\hat{\theta}_{t,i} - \theta_{t,i}^*| \\ &> 0 \end{aligned} \tag{12}$$

where the last inequality is due to the second assumption of the theorem and (11). This implies that $\text{supp}(\theta_t^*) \subseteq \text{supp}(\hat{\theta}_t)$. Similarly, suppose that $\theta_{t,i}^* - \theta_{t-1,i}^* \neq 0$ for some time $t > 0$ and index i . One can write

$$\begin{aligned} |\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i}| &= |\hat{\theta}_{t,i} - \theta_{t,i}^* + \theta_{t,i}^* - \theta_{t-1,i}^* + \theta_{t-1,i}^* - \hat{\theta}_{t-1,i}| \\ &\geq |\theta_{t,i}^* - \theta_{t-1,i}^*| - |\hat{\theta}_{t,i} - \theta_{t,i}^*| - |\hat{\theta}_{t-1,i} - \theta_{t-1,i}^*| \\ &> 0 \end{aligned} \tag{13}$$

where the last inequality is due to the third assumption of the theorem and (11). This implies that $\text{supp}(\theta_t^* - \theta_{t-1}^*) \subseteq \text{supp}(\hat{\theta}_t - \hat{\theta}_{t-1})$. Finally, due to the optimality of $\{\hat{\theta}_t\}_{t=0}^T$ and feasibility of $\{\theta_t^*\}_{t=0}^T$, one can write

$$\begin{aligned} (1 - \alpha) \sum_{t=0}^T \|\hat{\theta}_t\|_0 + \alpha \sum_{t=1}^T \|\hat{\theta}_t - \hat{\theta}_{t-1}\|_0 &\leq (1 - \alpha) \sum_{t=0}^T \|\theta_t^*\|_0 + \alpha \sum_{t=1}^T \|\theta_t^* - \theta_{t-1}^*\|_0 \\ \implies (1 - \alpha) \sum_{t=0}^T \left(\sum_{i \in [p] \setminus \mathcal{S}_t} |\hat{\theta}_{t,i}|_0 + \sum_{i \in \mathcal{S}_t} |\hat{\theta}_{t,i}|_0 \right) &+ \alpha \sum_{t=1}^T \left(\sum_{i \in [p] \setminus \mathcal{D}_t} |\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i}|_0 + \sum_{i \in \mathcal{D}_t} |\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i}|_0 \right) \\ &\leq (1 - \alpha) \sum_{t=0}^T \sum_{i \in \mathcal{S}_t} |\theta_{t,i}^*|_0 + \alpha \sum_{t=1}^T \sum_{i \in \mathcal{D}_t} |\theta_{t,i}^* - \theta_{t-1,i}^*|_0 \\ \implies (1 - \alpha) \sum_{t=0}^T \sum_{i \in [p] \setminus \mathcal{S}_t} |\hat{\theta}_{t,i}|_0 &+ \alpha \sum_{t=1}^T \sum_{i \in [p] \setminus \mathcal{D}_t} |\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i}|_0 \leq 0 \end{aligned} \tag{14}$$

where the last inequality follows from $\text{supp}(\theta_t^*) \subseteq \text{supp}(\hat{\theta}_t)$ and $\text{supp}(\theta_t^* - \theta_{t-1}^*) \subseteq \text{supp}(\hat{\theta}_t - \hat{\theta}_{t-1})$, which implies $\sum_{i \in \mathcal{S}_t} |\hat{\theta}_{t,i}|_0 - |\theta_{t,i}^*|_0 \geq 0$ and $\sum_{i \in \mathcal{D}_t} |\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i}|_0 - |\theta_{t,i}^* - \theta_{t-1,i}^*|_0 \geq 0$ for every t . Due to $0 < \alpha < 1$, the above inequality implies that $\hat{\theta}_{t,i} = 0$ for every t and $i \in [p] \setminus \mathcal{S}_t$, and $\hat{\theta}_{t,i} - \hat{\theta}_{t-1,i} = 0$ for every $t > 0$ and $i \in [p] \setminus \mathcal{D}_t$. This implies that $\text{supp}(\hat{\theta}_t) \subseteq \text{supp}(\theta_t^*)$ and $\text{supp}(\hat{\theta}_t - \hat{\theta}_{t-1}) \subseteq \text{supp}(\theta_t^* - \theta_{t-1}^*)$. Finally, since $\text{supp}(\hat{\theta}_t) \subseteq \text{supp}(\theta_t^*)$, we have $|\text{supp}(\hat{\theta}_t - \theta_t^*)| = |\mathcal{S}_t|$. This, together with (11) implies that $\|\hat{\theta}_t - \theta_t^*\|_2 \leq \sqrt{|\mathcal{S}_t|} \|\hat{\theta}_t - \theta_t^*\|_\infty \leq 2\sqrt{|\mathcal{S}_t|} \lambda_t$, thereby completing the proof. \square

A.2 Proof of Theorem 3

The proof is inspired by Corollary 1 in [41]. First, we present the following key lemmas.

Lemma 1 (Lemma 2 of [41] and Lemma 1 of [29]). *Suppose that $X^{(i)} \sim \mathcal{N}(0, \Sigma)$ for $i = 1, \dots, N$, and $\hat{\Sigma} = \frac{1}{N} \sum_{i=1}^N X^{(i)} X^{(i)\top}$. Then, we have*

$$\left\| \hat{\Sigma} - \Sigma \right\|_{\infty/\infty} \leq 8 \left(\max_i \Sigma_{ii} \right) \sqrt{\frac{\tau \log d}{N}} \quad (15)$$

with probability of at least $1 - 4d^{-\tau+2}$ for any $\tau > 2$, provided that $N \geq 40 (\max_i \Sigma_{ii})$.

Lemma 2 (Lemma 1 of [41]; modified). *Under the conditions of Lemma 1, we have*

$$\left\| S_{T_{\nu}}(\hat{\Sigma}) - \Sigma \right\|_{\infty} \leq 5\nu^{1-q} s(q, d) + 24\nu^{-q} s(q, d) \left(\max_i \Sigma_{ii} \right) \sqrt{\frac{\tau \log d}{N}} \quad (16)$$

with probability of at least $1 - 4d^{-\tau+2}$ for any $\tau > 2$, provided that $N \geq 40 (\max_i \Sigma_{ii})$.

Based on the above lemmas, we proceed to present the proof of Corollary 3.

Proof of Corollary 3. It suffices to show that the conditions of Theorem 1 are satisfied with the proposed choices of λ_t and ν_t . It is easy to see that

$$\begin{aligned} \left\| \Theta_t - [S_{T_{\nu_t}}(\hat{\Sigma}_t)]^{-1} \right\|_{\infty/\infty} &= \left\| [S_{T_{\nu_t}}(\hat{\Sigma}_t)]^{-1} (S_{T_{\nu_t}}(\hat{\Sigma}_t) \Theta_t - I) \right\|_{\infty/\infty} \\ &\leq \left\| [S_{T_{\nu_t}}(\hat{\Sigma}_t)]^{-1} \right\|_{\infty} \left\| \Theta_t \right\|_{\infty} \left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) - \Sigma_t \right\|_{\infty/\infty} \end{aligned} \quad (17)$$

We provide separate bounds for different terms of the above inequality. Due to Assumption 1, one can write $\left\| \Theta_t \right\|_{\infty} \leq \kappa_1$. Moreover, due to Lemma 1, the following inequality holds with probability of at least $1 - 4d^{-\tau+2}$ for any $\tau > 2$

$$\begin{aligned} \left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) - \Sigma_t \right\|_{\infty/\infty} &\leq \left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) - \hat{\Sigma}_t \right\|_{\infty/\infty} + \left\| \hat{\Sigma}_t - \Sigma_t \right\|_{\infty/\infty} \\ &\leq \nu_t + 8\kappa_3 \sqrt{\frac{\tau \log d}{N_t}} \\ &= 16\kappa_3 \sqrt{\frac{\tau \log d}{N_t}} \end{aligned} \quad (18)$$

provided that $N_t \geq 40\kappa_3$ and $\nu_t = 8\kappa_3 \sqrt{\frac{\tau \log d}{N_t}}$. Finally, for any vector w , one can write

$$\begin{aligned} \left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) w \right\|_{\infty} &\geq \left\| \Sigma_t w \right\|_{\infty} - \left\| (S_{T_{\nu_t}}(\hat{\Sigma}_t) - \Sigma_t) w \right\|_{\infty} \\ &\geq \left(\kappa_2 - \left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) - \Sigma_t \right\|_{\infty} \right) \left\| w \right\|_{\infty} \end{aligned} \quad (19)$$

On the other hand, the aforementioned choice of ν_t and Lemma 2 implies that

$$\left\| S_{T_{\nu_t}}(\hat{\Sigma}_t) - \Sigma_t \right\|_{\infty} \leq 64\kappa_3^{1-q} s(q, d) \left(\frac{\tau \log d}{N_t} \right)^{\frac{1-q}{2}} \quad (20)$$

Combining this inequality with (19) leads to

$$\left\| \text{ST}_{\nu_t}(\widehat{\Sigma}_t) - \Sigma_t \right\|_{\infty} \leq \frac{\kappa_2}{2} \quad (21)$$

provided that

$$N_t \geq \left(\frac{128s(q, d)}{\kappa_2} \right)^{\frac{2}{1-q}} \kappa_3^2 \tau \log d \quad (22)$$

This implies that $\left\| \text{ST}_{\nu_t}(\widehat{\Sigma}_t)w \right\|_{\infty} \geq \frac{\kappa_2}{2} \|w\|_{\infty}$, and hence, $\left\| [\text{ST}_{\nu_t}(\widehat{\Sigma}_t)]^{-1} \right\|_{\infty} \leq \frac{2}{\kappa_2}$. Combining these bounds with (23) yields

$$\left\| \Theta_t - [\text{ST}_{\nu_t}(\widehat{\Sigma}_t)]^{-1} \right\|_{\infty/\infty} \leq \frac{32\kappa_1\kappa_3}{\kappa_2} \sqrt{\frac{\tau \log d}{N_t}} = \lambda_t \quad (23)$$

with probability of at least $1 - 4d^{-\tau+2}$. Finally, we need to verify that the conditions $\lambda_t \leq \Theta_t^{\min}/2$ and $\lambda_t + \lambda_{t-1} \leq \Delta\Theta_t^{\min}/2$ hold. Based on the above definition of λ_t , it is easy to see that both of these conditions are satisfied if

$$\begin{aligned} N_t &\geq \left(\frac{128\kappa_1\kappa_3}{\kappa_2} \right)^2 \max \left\{ (\Theta_t^{\min})^{-2}, (\Delta\Theta_t^{\min})^{-2}, (\Delta\Theta_{t-1}^{\min})^{-2} \right\} \tau \log d \\ \implies N_t &\gtrsim \tau \log d \end{aligned} \quad (24)$$

Based on our assumption, we have $T + 1 \leq Cd^{\zeta}$ for some universal constant $C > 0$. Therefore, a simple union bound over $t = 0, \dots, T$ implies that the statements of the corollary holds for every $t = 0, \dots, T$ with the probability of at least

$$1 - 4 \sum_{t=0}^T d^{-\tau+2} \geq 1 - 4(T+1)d^{-\tau+2} \geq 1 - 4d^{\zeta-\tau+2} \quad (25)$$

Selecting $\tau > \zeta + 2$ completes the proof. \square

A.3 Proof of Theorem 4

First, we delineate the imposed assumptions on the selected kernel function.

Assumption 5 ([14]). *The kernel $K(x)$ satisfies the following conditions:*

- $\int_{-1}^1 K(x)dx = 1$,
- $\int_{-1}^1 x^2 K(x)dx \leq \infty$,
- $K(x)$ is uniformly bounded on its support,
- $\sup_{-1 \leq x \leq 1} K''(x/h) = \mathcal{O}(h^{-4})$.

The following key lemmas are borrowed from [14].

Lemma 3 (Lemma 5 of [14]). *For any fixed t , we have*

$$\left\| \mathbb{E}[\Sigma_t^w] - \Sigma(t/T) \right\|_{\infty/\infty} \lesssim C \left(h + \frac{1}{T^2 h^5} \right) \quad (26)$$

for some constant $C > 0$.

Lemma 4 (Lemma 2 of [14]). *There exists a constant $c > 0$ such that*

$$\mathbb{P}(|[\Sigma_t^w]_{ij} - \mathbb{E}[\Sigma_t^w]_{ij}| \geq \epsilon) \leq 2 \exp(-cTh\epsilon^2) \quad (27)$$

for every $\epsilon > 0$ and any fixed t .

Combining the above lemmas gives rise to the following result.

Lemma 5. *Assume that $h = T^{-1/3}$. Then, the following inequality holds for any t and $\tau > 2$*

$$\left\| \hat{\Sigma}_t^w - \Sigma(t/T) \right\|_{\infty/\infty} \lesssim \frac{\sqrt{\tau \log d}}{T^{1/3}} \quad (28)$$

with probability of at least $1 - d^{-(\tau-2)}$.

Proof. Based on Lemma 4, one can write

$$\mathbb{P} \left(\left\| \hat{\Sigma}_t^w - \mathbb{E}[\Sigma_t^w] \right\|_{\infty/\infty} \geq \epsilon \right) \leq 2 \exp(2 \log d - cTh\epsilon^2) \quad (29)$$

Upon choosing $\epsilon = \sqrt{\frac{\tau \log d}{cTh}}$ for some $\tau > 2$, we have

$$\left\| \hat{\Sigma}_t^w - \mathbb{E}[\Sigma_t^w] \right\|_{\infty/\infty} \leq \sqrt{\frac{\tau \log d}{cTh}} \quad (30)$$

with probability of at least $1 - d^{-(\tau-2)}$. Combined with Lemma 3, the following chain of inequalities hold with the same probability

$$\begin{aligned} \left\| \hat{\Sigma}_t^w - \Sigma(t/T) \right\|_{\infty/\infty} &\leq \left\| \hat{\Sigma}_t^w - \mathbb{E}[\Sigma_t^w] \right\|_{\infty/\infty} + \left\| \mathbb{E}[\Sigma_t^w] - \Sigma(t/T) \right\|_{\infty/\infty} \\ &\leq \sqrt{\frac{\tau \log d}{cTh}} + C \left(h + \frac{1}{T^2 h^5} \right) \end{aligned} \quad (31)$$

Replacing $h = T^{-1/3}$ in the above inequality gives rise to

$$\left\| \hat{\Sigma}_t^w - \Sigma(t/T) \right\|_{\infty/\infty} \lesssim \frac{\sqrt{\tau \log d}}{T^{1/3}} \quad (32)$$

which completes the proof. \square

Lemma 6. *Assume that $h = T^{-1/3}$. Then, the following inequality holds for any t and $\tau > 2$*

$$\left\| S_{T\nu}(\hat{\Sigma}_t^w) - \Sigma(t/T) \right\|_{\infty} \lesssim \nu^{1-q} s(q, d) + \nu^{-q} s(q, d) \frac{\sqrt{\tau \log d}}{T^{1/3}} \quad (33)$$

with probability of at least $1 - d^{-\tau+2}$.

Proof. The proof is implied by Lemma 1 of [41] and Lemma 5. \square

Proof of Corollary 4. We only provide a sketch of the proof, due to its similarity to the proof of Corollary 3. One can write

$$\left\| \Theta(t/T) - [S_{T\nu_t}(\hat{\Sigma}_t^w)]^{-1} \right\|_{\infty/\infty} \leq \left\| [S_{T\nu_t}(\hat{\Sigma}_t^w)]^{-1} \right\|_{\infty} \left\| \Theta(t/T) \right\|_{\infty} \left\| S_{T\nu_t}(\hat{\Sigma}_t^w) - \Sigma(t/T) \right\|_{\infty/\infty} \quad (34)$$

Due to Assumption 3, we have $\|\Theta(t/T)\|_\infty \leq \kappa_1$. Furthermore, similar to (19), one can write

$$\begin{aligned} \left\| \text{ST}_{\nu_t}(\hat{\Sigma}_t^w) - \Sigma(t/T) \right\|_{\infty/\infty} &\leq \left\| \text{ST}_{\nu_t}(\hat{\Sigma}_t^w) - \hat{\Sigma}_t^w \right\|_{\infty/\infty} + \left\| \hat{\Sigma}_t^w - \Sigma(t/T) \right\|_{\infty/\infty} \\ &\lesssim \frac{\sqrt{\tau \log d}}{T^{1/3}} \end{aligned} \quad (35)$$

with probability of at least $1 - d^{-\tau+2}$, where the second inequality follows from Lemma 5 and the choice of $\nu_t \asymp \frac{\sqrt{\tau \log d}}{T^{1/3}}$. Finally, Lemma 6 combined with an argument similar to the proof of Corollary 3 leads to

$$\left\| [\text{ST}_{\nu_t}(\hat{\Sigma}_t^w)]^{-1} \right\|_\infty \leq \frac{2}{\kappa_2} \quad (36)$$

provided that

$$T \gtrsim s(q, d)^{\frac{3}{1-q}} (\tau \log d)^{3/2} \quad (37)$$

Combining these inequalities leads to the desired upper bound on (34). The rest of the proof is similar to that of Corollary 3 and omitted for brevity. \square

A.4 Proof of Proposition 1

Let $\delta_1 < \delta_2 < \dots < \delta_m = T$ be the elements of the set Γ from Algorithm 1, and define $\delta_0 = -1$. By construction, $\Delta_{\delta_{i-1}+1 \rightarrow \delta_i+1}^\cap = \emptyset$ for all $i = 1, \dots, m-1$. It follows that for any θ satisfying bound constraints (8b) and $i = 1, \dots, m-1$, we have that

$$\sum_{t=\delta_{i-1}+1}^{\delta_i} \mathbb{1}\{\theta_{t+1} - \theta_t \neq 0\} \geq 1.$$

Given any $j = 1, \dots, T$, let h be the maximum index such that $\delta_h < j$. Therefore, we find that for any feasible θ ,

$$f_{0 \rightarrow j}(\theta) = \sum_{t=0}^{j-1} \mathbb{1}\{\theta_{t+1} - \theta_t \neq 0\} \geq \sum_{t=0}^{\delta_h} \mathbb{1}\{\theta_{t+1} - \theta_t \neq 0\} = \sum_{i=1}^h \sum_{t=\delta_{i-1}+1}^{\delta_i} \mathbb{1}\{\theta_{t+1} - \theta_t \neq 0\} \geq h.$$

Since $f_{0 \rightarrow j}(\theta^{\text{Greedy}}) = h$ meets this lower bound, it follows that $\{\theta_t^{\text{Greedy}}\}_{t=0}^j$ is indeed an optimal solution to $\text{OPT}_{0 \rightarrow j}(1)$. Setting $j = T$ and $h = m-1$, we find that θ^{Greedy} is optimal for $\text{OPT}_{0 \rightarrow T}(1)$. \square

A.5 Proof of Theorem 5

Before proving this theorem, we need the following intermediate lemma:

Lemma 7. *Given any optimal solution $\hat{\theta}$ to (8), exactly one of the following holds for any given zero-feasible sequence $\mathcal{Z}_{i \rightarrow j}$:*

1. $\hat{\theta}_i = \hat{\theta}_{i+1} = \dots = \hat{\theta}_j = 0$
2. $\hat{\theta}_\tau \neq 0$ for all $\tau = i, \dots, j$.

Proof. Let θ be any feasible solution to (2) that does not satisfy the conditions of Proposition 7, i.e., there exists $\tau = i, \dots, j - 1$ such that either $\theta_\tau = 0$ and $\theta_{\tau+1} \neq 0$, or $\theta_\tau \neq 0$ and $\theta_{\tau+1} = 0$. We now show how to construct a solution $\hat{\theta}$ with improved objective value, i.e., $f_{0 \rightarrow T}(\hat{\theta}) < f_{0 \rightarrow T}(\theta)$.

Consider the case $\theta_\tau = 0$ and $\theta_{\tau+1} \neq 0$. Define $\hat{\theta}_{\tau+1} = 0$ and $\hat{\theta}_t = \theta_t$ for all other coordinates $t \neq \tau + 1$. Clearly, $\hat{\theta}$ satisfies all bound constraints (2). Moreover,

$$f_{0 \rightarrow T}(\hat{\theta}) = f_{0 \rightarrow T}(\theta) - \underbrace{(1 - \alpha)}_{\hat{\theta}_{\tau+1} = 0} - \underbrace{\alpha}_{\hat{\theta}_\tau = \hat{\theta}_{\tau+1}} + \underbrace{\alpha \mathbb{1}\{\hat{\theta}_{\tau+1} \neq \hat{\theta}_{\tau+2}\}}_{\text{this term is 0 if } \tau + 1 = T} \leq f_{0 \rightarrow T}(\theta) - (1 - \alpha) < f_{0 \rightarrow T}(\theta).$$

The case $\theta_\tau \neq 0$ and $\theta_{\tau+1} = 0$ is handled analogously. \square

Since Lemma 7 holds for any zero-feasible sequence, it holds in particular for all maximal zero-feasible sequences. Based on this lemma, we are ready to present the proof of Theorem 5.

Proof of Theorem 5. Let $\hat{\theta}$ be an optimal solution to (8). Due to the optimality of $\hat{\theta}$, the conditions in Lemma 7 are satisfied for all maximal nonzero intervals. We first show that there exists a path in \mathcal{G} with cost $f(\hat{\theta})$, and then we show that this path is indeed a shortest path.

Let $V_0 = \{v_1, v_2, \dots, v_m\} \subseteq \{1, \dots, Z\}$ be the set of indexes of the maximal zero-feasible sequences where $\hat{\theta}$ vanishes, i.e., $\hat{\theta}_{\mathcal{Z}_{i_s \rightarrow j_s}} = 0$ for every $s \in V_0$. It is easy to verify that f^* is the optimal cost for the following constrained optimization:

$$f_{0 \rightarrow T}(\hat{\theta}) = \min_{\{\theta_t\}_{t=0}^T} (1 - \alpha) \left(T + 1 - \left| \bigcup_{h=1}^m \mathcal{Z}_{i_{v_h} \rightarrow j_{v_h}} \right| \right) + \alpha \sum_{t=1}^T \mathbb{1}\{\theta_t - \theta_{t-1} \neq 0\} \quad (38a)$$

$$\text{subject to } l_t \leq \theta_t \leq u_t \quad t = 0, \dots, T \quad (38b)$$

$$\theta_t = 0 \quad t \in \bigcup_{h=1}^m \mathcal{Z}_{i_{v_h} \rightarrow j_{v_h}}. \quad (38c)$$

The constant term in (38a) reduces to

$$(1 - \alpha) \left(T + 1 - \left| \bigcup_{h=1}^m \mathcal{Z}_{i_{v_h} \rightarrow j_{v_h}} \right| \right) = (1 - \alpha) \left(T + 1 - \sum_{h=1}^m (j_{v_h} - i_{v_h} + 1) \right) \quad (39)$$

$$= (1 - \alpha) \left(i_{v_1} + \sum_{h=2}^m (i_{v_h} - j_{v_{h-1}} - 1) + (T - j_{v_m}) \right). \quad (40)$$

Let the feasible region of (38) be denoted as \mathcal{X} . The second term in (38a), under constraints (38c), decomposes as

$$\begin{aligned} & \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \alpha \sum_{t=1}^T \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} \\ &= \alpha \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=1}^{i_{v_1}} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} + \alpha \sum_{h=1}^{m-1} \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_h}+1}^{i_{v_{h+1}}} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} \\ & \quad + \alpha \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_m}+1}^T \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} \end{aligned} \quad (41)$$

Note that each intermediate term in (41) simplifies as follows:

$$\begin{aligned} \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_h}+1}^{i_{v_h}+1} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} &= \underbrace{\min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_h}+2}^{i_{v_h}+1-1} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\}}_{=f_{j_{v_h}+1 \rightarrow i_{v_h}+1-1}^{\text{Greedy}}} \\ &\quad + \underbrace{\mathbb{1}\{\theta_{j_{v_h}+1} \neq \theta_{j_{v_h}}\}}_{=1} + \underbrace{\mathbb{1}\{\theta_{i_{v_h}+1} \neq \theta_{i_{v_h}+1-1}\}}_{=1}. \end{aligned} \quad (42)$$

Similarly, we find that the first and last term in (41) reduces to

$$\min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=1}^{i_{v_1}} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} = \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=1}^{i_{v_1}-1} \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} + \mathbb{1}\{i_{v_1} \geq 1\} = f_{0 \rightarrow i_{v_1}-1}^{\text{Greedy}} + \mathbb{1}\{i_{v_1} \neq 0\} \quad (43)$$

$$\begin{aligned} \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_m}+1}^T \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} &= \min_{\{\theta_t\}_{t=0}^T \in \mathcal{X}} \left\{ \sum_{t=j_{v_m}+2}^T \mathbb{1}\{\theta_t \neq \theta_{t-1}\} \right\} + \mathbb{1}\{j_{v_m}+1 \leq T\} \\ &= f_{j_{v_m}+1 \rightarrow T}^{\text{Greedy}} + \mathbb{1}\{j_{v_m} < T\}. \end{aligned} \quad (44)$$

Combining (44), (43) with (41) and (40), we find that $f_{0 \rightarrow T}(\hat{\theta})$ is precisely the length of the path $(0, v_1, \dots, v_m, Z+1)$ in the constructed graph \mathcal{G} with weights defined as (9).

Now suppose that there exists a path $(0, \bar{v}_1, \bar{v}_2, \dots, \bar{v}_p, Z+1)$ with length $\bar{d} < f_{0 \rightarrow T}(\hat{\theta})$. Consider a solution $\bar{\theta}$ such that: (i) $\bar{\theta}$ is zero at zero-feasible sequences given by $\bar{v}_1, \bar{v}_2, \dots, \bar{v}_p$, and (ii) $\bar{\theta}$ is obtained from $\text{Greedy}(l, u, 0, i_{v_1}-1)$, $\text{Greedy}(l, u, j_{v_p}+1, T)$ and $\text{Greedy}(l, u, j_{v_h}+1, i_{v_{h+1}}-1)$, otherwise. It is easy to verify that $\bar{\theta}$ is feasible and satisfies $f_{0 \rightarrow T}(\bar{\theta}) \leq \bar{d}$ (the inequality could be strict if any solution reported by a call to the Greedy routine has zero values), which contradicts the optimality of $\hat{\theta}$. Thus, we conclude that $f_{0 \rightarrow T}(\hat{\theta})$ is indeed the length of the shortest $(0, Z+1)$ -path in \mathcal{G} . \square

A.6 Proof of Theorem 6

Algorithm 2 involves three main components: construct graph \mathcal{G} (line 3), solve a shortest problem on the constructed graph (line 4), and recover the optimal solution from the obtained shortest path. Since \mathcal{G} is acyclic, the shortest path problem can be solved in time linear in the number of arcs, which is $\mathcal{O}(Z^2)$, via a simple labeling algorithm; see, e.g., Chapter 4.4. in [4]. Constructing graph \mathcal{G} requires computing the costs of all arcs. A naïve implementation, where Algorithm 1 is called for every arc, would require $\mathcal{O}(Z^2T)$ time and memory. However, from the second statement in Proposition 1, we note that a single call to $\text{Greedy}(l, u, i, T)$ allows us to compute $f_{i \rightarrow j}^{\text{Greedy}}$ for all $i \leq j \leq T$. Therefore, Algorithm 1 needs to be invoked only $\mathcal{O}(Z)$ times, and each call require $\mathcal{O}(T)$ leading to a total complexity of $\mathcal{O}(ZT)$. Moreover, given the shortest path, the optimal solution can be constructed by concatenating the solutions obtained from the calls of Greedy. Finally, since $Z \leq T+2$, we find that the overall complexity is dominated by that of constructing the graph. This completes the proof. \square