Sparse Trace Norm Regularization

Jianhui Chen*
GE Global Research, Niskayuna, NY

Jieping Ye
Arizona State University, Tempe, AZ

June 5, 2012

Abstract

We study the problem of estimating multiple predictive functions from a dictionary of basis functions in the nonparametric regression setting. Our estimation scheme assumes that each predictive function can be estimated in the form of a linear combination of the basis functions. By assuming that the coefficient matrix admits a sparse low-rank structure, we formulate the function estimation problem as a convex program regularized by the trace norm and the ℓ_1-norm simultaneously. We propose to solve the convex program using the accelerated gradient (AG) method and the alternating direction method of multipliers (ADMM) respectively; we also develop efficient algorithms to solve the key components in both AG and ADMM. In addition, we conduct theoretical analysis on the proposed function estimation scheme: we derive a key property of the optimal solution to the convex program; based on an assumption on the basis functions, we establish a performance bound of the proposed function estimation scheme (via the composite regularization). Simulation studies demonstrate the effectiveness and efficiency of the proposed algorithms.

1 Introduction

We study the problem of estimating multiple predictive functions from noisy observations. Such a problem has received broad attention in many areas of statistics and machine learning [6, 15, 16, 18]. This line of work can be roughly divided into two categories: parametric estimation and non-parametric estimation; a common and important theme for both categories is the appropriate assumption of the structure in the model parameters (parametric setting) or the coefficients of the dictionary (nonparametric setting).

There has been an enormous amount of literature on effective function estimation based on different sparsity constraints, including the estimation of the sparse linear regression via ℓ_1-norm penalty [3, 6, 27, 32], and the estimation of the linear regression functions using group lasso estimator [15, 16]. More recently, trace norm regularization has become a popular tool for approximating a set of linear models and the associated low-rank matrices in the high-dimensional setting [18, 24]; the trace norm is the tightest convex surrogate [12] for the (non-convex) rank function under certain conditions, encouraging the sparsity in the singular values of the matrix of interest. One limitation of the use of trace norm regularization is that the resulting model is dense in general. However, in many real-world applications [21], the underlying structure of multiple predictive functions may be sparse as well as low-rank; the sparsity leads to explicitly interpretable prediction models and the low-rank implies essential subspace structure information. Similarly, the ℓ_1-norm is the tightest convex surrogate for the non-convex cardinality function [5], encouraging the sparsity in the entries of the matrix. This motivates us to explore the use of the combination of the trace norm and the ℓ_1-norm as a composite regularization (called sparse trace norm regularization) to induce the desirable sparse low-rank structure.

---

*This work was done when the first author was a Ph.D. student at Arizona State University.

Trace norm regularization (minimization) has been investigated extensively in recent years. Efficient algorithms have been developed for solving convex programs with trace norm regularization [29, 12]; sufficient conditions for exact recovery from trace norm minimization have been established in [22]; consistency of trace norm minimization has been studied in [1]; trace norm minimization has been applied for matrix completion [7] and collaborative filtering [25, 23]. Similarly, ℓ_1-norm regularization has been well studied in the literature, just to mention a few, from the efficient algorithms for convex optimization [11, 13, 29], theoretical guarantee of the performance [9, 32], and model selection consistency [33].

In this paper, we focus on estimating multiple predictive functions simultaneously from a finite dictionary of basis functions in the nonparametric regression setting. Our function estimation scheme assumes that each predictive function can be approximated using a linear combination of those basis functions. By assuming that the coefficient matrix of the basis functions admits a sparse low-rank structure, we formulate the function estimation problem as a convex formulation, in which the combination of the trace norm and the ℓ_1-norm is employed as a composite regularization to induce a sparse low-rank structure in the coefficient matrix. The simultaneous sparse and low-rank structure is different from the incoherent sparse and low-rank structures studied in [8, 10]. We propose to solve the function estimation problem using the accelerated gradient method and the alternating direction method of multipliers; we also develop efficient algorithms to solve the key components involved in both methods. We conduct theoretical analysis on the proposed convex formulation: we first present some basic properties of the optimal solution to the convex formulation (Lemma 4.1); we then present an assumption associated with the geometric nature of the basis functions over the prescribed observations; based on such an assumption, we derive a performance bound for the combined regularization for function estimation (Theorem 4.1). We conduct simulations on benchmark data to demonstrate the effectiveness and efficiency of the proposed algorithms.

Notation Denote $\mathbb{N}_n = \{1, \cdots, n\}$. For any matrix Θ, denote its trace norm by $\|\Theta\|_*$, i.e., the sum of the singular values; denote its operator norm by $\|\Theta\|_2$, i.e., the largest singular value; denote its ℓ_1-norm by $\|\Theta\|_1$, i.e., the sum of absolute value of all entries.

2 Problem Formulation

Let $\{(x_1, y_1), \cdots, (x_n, y_n)\} \subset \mathbb{R}^d \times \mathbb{R}^k$ be a set of prescribed sample pairs (fixed design) associated with k unknown functions $\{f_1, \cdots, f_k\}$ as

$$y_{ij} = f_j(x_i) + w_{ij}, \quad i \in \mathbb{N}_n, \ j \in \mathbb{N}_k, \tag{1}$$

where $f_j : \mathbb{R}^d \to \mathbb{R}$ is an unknown regression function, y_{ij} denotes the j-th entry of the response vector $y_i \in \mathbb{R}^k$, and $w_{ij} \sim \mathcal{N}(0, \sigma_w^2)$ is a stochastic noise variable. Let $X = [x_1, \cdots, x_n]^T \in \mathbb{R}^{n \times d}$, $Y = [y_1, \cdots, y_n]^T \in \mathbb{R}^{n \times k}$, and $W = (w_{ij})_{i,j} \in \mathbb{R}^{n \times k}$. Denoting

$$\mathcal{F} = (f_j(x_i))_{i,j} \in \mathbb{R}^{n \times k}, \quad i \in \mathbb{N}_n, \ j \in \mathbb{N}_k, \tag{2}$$

we can rewrite Eq. (1) in a compact form as $Y = \mathcal{F} + W$. Let $\{g_1, \cdots, g_h\}$ be a set of h pre-specified basis functions as $g_i : \mathbb{R}^d \to \mathbb{R}$, and let $\Theta = [\theta_1, \cdots, \theta_k] \in \mathbb{R}^{h \times k}$ be the coefficient matrix. We define

$$\hat{g}_j(x) = \sum_{i=1}^{h} \theta_{ij} g_i(x), \quad j \in \mathbb{N}_k, \tag{3}$$

where θ_{ij} denotes the i-th entry in the vector θ_j. Note that in practice the basis functions $\{g_i\}$ can be estimators from different methods, or different values of the tuning parameters of the same method.

We consider the problem of estimating the unknown functions $\{f_1, \cdots, f_k\}$ using the composite functions $\{\hat{g}_1, \cdots, \hat{g}_k\}$ defined in Eq. (3), respectively. Denote

$$\mathcal{G}_X = (g_j(x_i))_{i,j} \in \mathbb{R}^{n \times h}, \quad i \in \mathbb{N}_n, \ j \in \mathbb{N}_h, \tag{4}$$

and define the empirical error as

$$\widehat{S}(\Theta) = \frac{1}{nk} \sum_{i=1}^{n} \sum_{j=1}^{k} (\hat{g}_j(x_i) - y_{ij})^2 = \frac{1}{N} \|\mathcal{G}_X \Theta - Y\|_F^2, \tag{5}$$

where $N = n \times k$. Our goal is to estimate the model parameter Θ of a sparse low-rank structure from the given n sample pairs $\{(x_i, y_i)\}_{i=1}^{n}$. Such a structure induces the sparsity and the low rank simultaneously in a single matrix of interest.

Given that the functions $\{f_1, \cdots, f_k\}$ are coupled via Θ in some coherent sparse and low-rank structure, we propose to estimate Θ as

$$\widehat{\Theta} = \arg\min_{\Theta} \left(\widehat{S}(\Theta) + \alpha \|\Theta\|_* + \beta \|\Theta\|_1 \right), \tag{6}$$

where α and β are regularization parameters (estimated via cross-validation), and the linear combination of $\|\Theta\|_*$ and $\|\Theta\|_1$ is used to induce the sparse low-rank structure in Θ. The optimization problem in Eq. (6) is non-smooth convex and hence admits a globally optimal solution; it can be solved using many sophisticated optimization techniques [28, 12]; in Section 3, we propose to apply the accelerated gradient method [19] and the alternating direction method of multipliers [4] to solve the optimization problem in Eq. (6).

3 Optimization Algorithms

In this section, we consider to apply the accelerated gradient (AG) algorithm [2, 19, 20] and the alternating direction method of multipliers (ADMM) [4], respectively, to solve the (non-smooth and convex) optimization problem in Eq. (6). We also develop efficient algorithms to solve the key components involved in both AG and ADMM.

3.1 Accelerated Gradient Algorithm

The AG algorithm has attracted extensive attention in the machine learning community due to its optimal convergence rate among all first order techniques and its ability of dealing with large scale data. The general scheme in AG for solving Eq. (6) can be described as below: at the k-th iteration, the intermediate (feasible) solution Θ_k can be obtained via

$$\Theta_k = \arg\min_{\Theta} \left(\frac{\gamma_k}{2} \left\| \Theta - \left(\Phi_k - \frac{1}{\gamma_k} \nabla \widehat{S}(\Phi_k) \right) \right\|_F^2 + \alpha \|\Theta\|_* + \beta \|\Theta\|_1 \right), \tag{7}$$

where Φ_k denotes a searching point constructed on the intermediate solutions from previous iterations, $\nabla \widehat{S}(\Phi_k)$ denotes the derivative of the loss function in Eq. (5) at Φ_k, and γ_k specifies the step size which can be determined by iterative increment until the condition

$$\widehat{S}(\Theta_k) \leq \widehat{S}(\Phi_k) + \langle \nabla f(\Phi_k), \Theta_k - \Phi_k \rangle + \frac{\gamma_k}{2} \|\Theta_k - \Phi_k\|_F^2$$

is satisfied. The operation in Eq. (7) is commonly referred to as proximal operator [17], and its efficient computation is critical for the practical convergence of the AG-type algorithm. Next we present an efficient alternating optimization procedure to solve Eq. (7) with a given γ_k.

3.1.1 Dual Formulation

The problem in Eq. (7) is not easy to solve directly; next we show that this problem can be efficiently solved in its dual form. By reformulating $\|\Theta\|_*$ and $\|\Theta\|_1$ into the equivalent dual forms, we convert Eq. (7) into a max-min formulation as

$$\max_{L,S} \min_{\Theta} \ \|\Theta - \widehat{\Phi}\|_F^2 + \widehat{\alpha}\langle L, \Theta \rangle + \widehat{\beta}\langle S, \Theta \rangle, \quad \text{subject to } \|L\|_2 \leq 1, \ \|S\|_\infty \leq 1, \tag{8}$$

where $\widehat{\Phi} = \Phi_k - \nabla \widehat{S}(\Phi_k)/\gamma_k$, $\widehat{\alpha} = 2\alpha/\gamma_k$, and $\widehat{\beta} = 2\beta/\gamma_k$. It can be verified that in Eq. (8) the Slater condition is satisfied and strong duality holds [5]. Also the optimal Θ can be expressed as a function of L and S given by

$$\Theta = \widehat{\Phi} - \frac{1}{2}(\widehat{\alpha}L + \widehat{\beta}S). \tag{9}$$

By substituting Eq. (9) into Eq. (8), we obtain the dual form of Eq. (7) as

$$\min_{L,S} \ \|\widehat{\alpha}L + \widehat{\beta}S - 2\widehat{\Phi}\|_F^2, \quad \text{subject to } \|L\|_2 \leq 1, \ \|S\|_\infty \leq 1. \tag{10}$$

3.1.2 Alternating Optimization

The optimization problem in Eq. (10) is smooth convex and it has two optimization variables. For such type of problems, coordinate descent (CD) method is routinely used to compute its globally optimal solution [14]. To solve Eq. (10), the CD method alternatively optimizes one of the two variables with the other variable fixed. Our analysis below shows that the variables L and S in Eq. (10) can be optimized efficiently. Note that the convergence rate of the CD method is not known, however, it converges very fast in practice (less than 10 iterations in our experiments).

Optimization of L For a given S, the variable L can be optimized via solving the following problem:

$$\min_{L} \ \|L - \widehat{L}\|_F^2, \quad \text{subject to } \|L\|_2 \leq 1, \tag{11}$$

where $\widehat{L} = (2\widehat{\Phi} - \widehat{\beta}S)/\widehat{\alpha}$. The optimization on L above can be interpreted as computing an optimal projection of a given matrix over a unit spectral norm ball. Our analysis shows that the optimal solution to Eq. (11) can be expressed in an analytic form as summarized in the following theorem.

Theorem 3.1. *For arbitrary $\widehat{L} \in \mathbb{R}^{h \times k}$ in Eq. (11), denote its SVD by $\widehat{L} = U\Sigma V^T$, where $r = rank(\widehat{L})$, $U \in \mathbb{R}^{h \times r}$, $V \in \mathbb{R}^{k \times r}$, and $\Sigma = diag(\sigma_1, \cdots, \sigma_r) \in \mathbb{R}^{r \times r}$. Let $\hat{\sigma}_i^* = \min(\sigma_i, 1)$, $i = 1, \cdots, r$. Then the optimal solution to Eq. (11) is given by*

$$L^* = U\hat{\Sigma}V^T, \ \hat{\Sigma} = diag(\hat{\sigma}_1^*, \cdots, \hat{\sigma}_r^*). \tag{12}$$

Proof. Assume the existence of a set of left and right singular vector pairs shared by the optimal L^* to Eq. (11) and the given $\widehat{L}$ for their non-zero singular values. Under such an assumption, it can be verified that the singular values of L^* can be obtained via

$$\min_{\{\hat{\sigma}_i\}} \ (\hat{\sigma}_i - \sigma_i)^2, \quad \text{subject to } 0 \leq \hat{\sigma}_i \leq 1, \ i = 1, \cdots, r,$$

to which the optimal solution is given by $\hat{\sigma}_i^* = \min(\sigma_i, 1)$ $(\forall i)$; hence the expression of L^* coincides with Eq. (12). Therefore, all that remains is to show that our assumption (on the left and right singular vector pairs of L^* and $\widehat{L}$) holds.

Denote the Lagrangian associated with the problem in Eq. (11) as $h(L, \lambda) = \|L - \widehat{L}\|_F^2 + \lambda(\|L\|_2 - 1)$, where λ denotes the dual variable. Since $\mathbf{0}$ is strictly feasible in Eq. (11), namely, $\|\mathbf{0}\|_2 < 1$, strong duality holds for Eq. (11). Let λ^* be the optimal dual variable to Eq. (11). Therefore we have $L^* = \arg\min_L h(L, \lambda^*)$. It is well known that L^* minimizes $h(L, \lambda^*)$ if and only if $\mathbf{0}$ is a subgradient of $h(L, \lambda^*)$ at L^*, i.e.,

$$\mathbf{0} \in 2(L^* - \widehat{L}) + \lambda^* \partial \|L^*\|_2. \tag{13}$$

For any matrix Z, the subdifferential of $\|Z\|_2$ is given by [31] $\partial \|Z\|_2 = \text{conv}\left\{u_z v_z^T : \|u_z\| = \|v_z\| = 1, Zv_z = \|Z\|_2 u_z\right\}$, where $\text{conv}\{c\}$ denotes the convex hull of the set c. Specifically, any element of $\partial\|Z\|_2$ has the form

$$\sum_i \alpha_i u_{zi} v_{zi}^T, \ \alpha_i \geq 0, \ \sum_i \alpha_i = 1,$$

4

where u_{zi} and v_{zi} are any left and right singular vectors of Z corresponding to its largest singular value (the top singular values may share a common value). From Eq. (13) and the definition of $\partial\|Z\|_2$, there exist $\{\hat{\alpha}_i\}$ such that $\hat{\alpha}_i > 0$, $\sum_i \hat{\alpha}_i = 1$, $\sum_i \hat{\alpha}_i u_{li} v_{li}^T \in \partial\|L^*\|_2$, and

$$\widehat{L} = L^* + \frac{\lambda^*}{2} \sum_i \hat{\alpha}_i u_{li} v_{li}^T, \tag{14}$$

where u_{li} and v_{li}^T correspond to any left and right singular vectors of L^* corresponding to its largest singular value. Since $\lambda^*, \hat{\alpha}_i > 0$, Eq. (14) verifies the existence of a set of left and right singular vector pairs shared by L^* and $\widehat{L}$. This completes the proof. $\square$

Optimization of S For a given L, the variable S can be optimized via solving the following problem:

$$\min_S \ \|S - \widehat{S}\|_F^2, \quad \text{subject to } \|S\|_\infty \le 1, \tag{15}$$

where $\widehat{S} = (2\widehat{\Phi} - \widehat{\alpha}L)/\widehat{\beta}$. Similarly, the optimization on S can be interpreted as computing a projection of a given matrix over an infinity norm ball. It also admits an analytic solution as summarized in the following theorem.

Lemma 3.1. *For any matrix $\widehat{S}$, the optimal solution to Eq. (15) is given by*

$$S^* = sgn(\widehat{S}) \circ \min(|\widehat{S}|, \mathbf{1}), \tag{16}$$

where $\circ$ denotes the component-wise multiplication operator, and $\mathbf{1}$ denotes the matrix with entries 1 of appropriate size.

3.2   Alternating Direction Method of Multipliers

The ADMM algorithm [4] is suitable for dealing with non-smooth (convex) optimizations problems, as it blends the decomposability of dual ascent with the superior convergence of the method of multipliers. We present two implementations of the ADMM algorithm for solving Eq. (6). Due to the space constraint, we move the detailed discussion of two ADMM implementations to the supplemental material.

4   Theoretical Analysis

In this section, we present a performance bound for the function estimation scheme in Eq. (3). Such a performance bound measures how well the estimation scheme can approximate the regression functions $\{f_j\}$ in Eq. (2) via the sparse low-rank coefficient Θ.

4.1   Basic Properties of the Optimal Solution

We first present some basic properties of the optimal solution defined in Eq. (6); these properties are important building blocks of our following theoretical analysis.

Lemma 4.1. *Consider the optimization problem in Eq. (6) for $h, k \ge 2$ and $n \ge 1$. Given n sample pairs as $X = [x_1, \cdots, x_n]^T \in \mathbb{R}^{n \times d}$ and $Y = [y_1, \cdots, y_n]^T \in \mathbb{R}^{n \times k}$. Let $\mathcal{F}$ and $\mathcal{G}_X$ be defined in Eq. (2) and Eq. (4), respectively; let $\sigma_{X(l)}$ be the largest singular values of $\mathcal{G}_X$. Assume that $W \in \mathbb{R}^{n \times k}$ has independent and identically distributed (i.i.d.) entries as $w_{ij} \sim \mathcal{N}(0, \sigma_w^2)$. Take*

$$\alpha + \beta = \frac{2\sigma_{X(l)}\sigma_w\sqrt{n}}{N}\left(1 + \sqrt{\frac{k}{n}} + t\right), \tag{17}$$

where $N = n \times k$ and t is a universal constant. Then with probability of at least $1 - \exp\left(-nt^2/2\right)$, for the minimizer $\widehat{\Theta}$ in Eq. (6) and any $\Theta \in \mathbb{R}^{h \times k}$, we have

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq \frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2 + 2\alpha\|\mathcal{S}_0(\widehat{\Theta} - \Theta)\|_* + 2\beta\|(\widehat{\Theta} - \Theta)_{J(\Theta)}\|_1, \tag{18}$$

where $\mathcal{S}_0$ is an operator defined in Lemma 1 of the supplemental material.

Proof. From the definition of $\widehat{\Theta}$ in Eq. (6), we have $\widehat{S}(\widehat{\Theta}) + \alpha\|\widehat{\Theta}\|_* + \beta\|\widehat{\Theta}\|_1 \leq \widehat{S}(\Theta) + \alpha\|\Theta\|_* + \beta\|\Theta\|_1$. By substituting $Y = \mathcal{F} + W$ and Eq. (5) into the previous inequality, we have

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq \frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2 + \frac{2}{N}\langle W, \mathcal{G}_X(\widehat{\Theta} - \Theta)\rangle + \alpha\left(\|\Theta\|_* - \|\widehat{\Theta}\|_*\right) + \beta\left(\|\Theta\|_1 - \|\widehat{\Theta}\|_1\right).$$

Define the random event

$$\mathcal{A} = \left\{\frac{1}{N}\|\mathcal{G}_X^T W\|_2 \leq \frac{\alpha + \beta}{2}\right\}. \tag{19}$$

Taking $\alpha + \beta$ as the value in Eq. (17), it follows from Lemma 3 of the supplemental materia that $\mathcal{A}$ holds with probability of at least $1 - \exp\left(-nt^2/2\right)$. Therefore, we have

$$\langle W, \mathcal{G}_X(\widehat{\Theta} - \Theta)\rangle = \frac{\alpha + \beta}{\alpha + \beta}\langle W, \mathcal{G}_X(\widehat{\Theta} - \Theta)\rangle$$

$$\leq \frac{\alpha}{\alpha + \beta}\|\mathcal{G}_X^T W\|_2\|\widehat{\Theta} - \Theta\|_* + \frac{\beta}{\alpha + \beta}\|\mathcal{G}_X^T W\|_\infty\|\widehat{\Theta} - \Theta\|_1 \leq \frac{N}{2}\left(\alpha\|\widehat{\Theta} - \Theta\|_* + \beta\|\widehat{\Theta} - \Theta\|_1\right),$$

where the second inequality follows from $\|\mathcal{G}_X^T W\|_2 \geq \|\mathcal{G}_X^T W\|_\infty$. Therefore, under $\mathcal{A}$, we have

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2$$

$$\leq \frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2 + \alpha\|\widehat{\Theta} - \Theta\|_* + \beta\|\widehat{\Theta} - \Theta\|_1 + \alpha\left(\|\Theta\|_* - \|\widehat{\Theta}\|_*\right) + \beta\left(\|\Theta\|_1 - \|\widehat{\Theta}\|_1\right).$$

From Corollary 1 and Lemma 2 of the supplemental material, we complete the proof. $\qquad\square$

4.2   Main Assumption

We introduce a key assumption on the dictionary of basis functions $\mathcal{G}_X$. Based on such an assumption, we derive a performance bound for the sparse trace norm regularization formulation in Eq. (6).

Assumption 4.1. *For a matrix pair Θ and Δ of size $h \times k$, let $s \leq \min(h, k)$ and $q \leq h \times k$. We assume that there exist constants $\kappa_1(s)$ and $\kappa_2(q)$ such that*

$$\kappa_1(s) \triangleq \min_{\Delta \in \mathcal{R}(s,q)} \frac{\|\mathcal{G}_X\Delta\|_F}{\sqrt{N}\|\mathcal{S}_0(\Delta)\|_*} > 0, \; \kappa_2(q) \triangleq \min_{\Delta \in \mathcal{R}(s,q)} \frac{\|\mathcal{G}_X\Delta\|_F}{\sqrt{N}\|\Delta_{J(\Theta)}\|_1} > 0, \tag{20}$$

where the restricted set $\mathcal{R}(s, q)$ is defined as

$$\mathcal{R}(s,q) = \left\{\Delta \in \mathbb{R}^{h \times k}, \Theta \in \mathbb{R}^{h \times k} \,|\, \Delta \neq 0, \; rank(\mathcal{S}_0(\Delta)) \leq s, \; |J(\Theta)| \leq q\right\},$$

and $|J(\Theta)|$ denotes the number of nonzero entries in the matrix Θ.

Our assumption on $\kappa_1(s)$ in Eq. (20) is closely related to but less restrictive than the RSC condition used in [18]; its denominator is only a part of the one in RSC and in a different matrix norm as well. Our assumption on $\kappa_2(q)$ is similar to the RE condition used in [3] except that its denominator is in a different matrix norm; our assumption can also be implied by sufficient conditions similar to the ones in [3].

4.3 Performance Bound

We derive a performance bound for the sparse trace norm structure obtained by solving Eq. (6). This bound measures how well the optimal $\widehat{\Theta}$ can be used to approximate $\mathcal{F}$ by evaluating the averaged estimation error, i.e., $\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2/N$.

Theorem 4.1. *Consider the optimization problem in Eq. (6) for $h, k \geq 2$ and $n \geq 1$. Given n sample pairs as $X = [x_1, \cdots, x_n]^T \in \mathbb{R}^{n \times d}$ and $Y = [y_1, \cdots, y_n]^T \in \mathbb{R}^{n \times k}$, let $\mathcal{F}$ and $\mathcal{G}_X$ be defined in Eqs. (2) and (4), respectively; let $\sigma_{X(l)}$ be the largest singular value of $\mathcal{G}_X$. Assume that $W \in \mathbb{R}^{n \times k}$ has i.i.d. entries as $w_{ij} \sim \mathcal{N}(0, \sigma_w^2)$. Take $\alpha + \beta$ as the value in Eq. (17). Then with probability of at least $1 - \exp\left(-nt^2/2\right)$, for the minimizer $\widehat{\Theta}$ in Eq. (6), we have*

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq (1+\epsilon)\inf_{\Theta}\left\{\frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2\right\} + \mathcal{E}(\epsilon)\left(\frac{\alpha^2}{\kappa_1^2(2r)} + \frac{\beta^2}{\kappa_2^2(c)}\right), \tag{21}$$

where $\inf$ is taken over all $\Theta \in \mathbb{R}^{h \times k}$ with $\text{rank}(\Theta) \leq r$ and $|J(\Theta)| \leq c$, and $\mathcal{E}(\epsilon) > 0$ is a constant depending only on ϵ.

Proof. Denote $\Delta = \widehat{\Theta} - \Theta$ in Eq. (18). We have

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq \frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2 + 2\alpha\|\mathcal{S}_0(\Delta)\|_* + 2\beta\|\Delta_{J(\Theta)}\|_1. \tag{22}$$

Given $\mathcal{S}_0(\Delta) \leq 2r$ (from Lemma 1 of the supplemental material) and $|J(\Theta)| \leq c$, we derive upper bounds on the components $2\alpha\|\mathcal{S}_0(\Delta)\|_*$ and $2\beta\|\Delta_{J(\Theta)}\|_1$ over the restrict set $\mathcal{R}(2r, c)$ based on Assumptions 4.1, respectively. It follows that

$$\begin{aligned}
2\alpha\|\mathcal{S}_0(\Delta)\|_* &\leq \frac{2\alpha}{\kappa_1(2r)\sqrt{N}}\|\mathcal{G}_X(\widehat{\Theta} - \Theta)\|_F \leq \frac{2\alpha}{\kappa_1(2r)\sqrt{N}}\left(\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F + \|\mathcal{G}_X\Theta - \mathcal{F}\|_F\right) \\
&\leq \frac{\alpha^2\tau}{\kappa_1^2(2r)} + \frac{1}{N\tau}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 + \frac{\alpha^2\tau}{\kappa_1^2(2r)} + \frac{1}{N\tau}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2,
\end{aligned} \tag{23}$$

where the last inequality above follows from $2ab \leq a^2\tau + b^2/\tau$ for $\tau > 0$. Similarly, we have

$$2\beta\|\Delta_{J(\Theta)}\|_1 \leq \frac{\beta^2\tau}{\kappa_2^2(c)} + \frac{1}{N\tau}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 + \frac{\beta^2\tau}{\kappa_2^2(c)} + \frac{1}{N\tau}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2. \tag{24}$$

Substituting Eqs. (23) and (24) into Eq. (22), we have

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq \frac{\tau+2}{(\tau-2)N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2 + \frac{2\tau^2}{\tau-2}\left(\frac{\alpha^2}{\kappa_1^2(2r)} + \frac{\beta^2}{\kappa_2^2(c)}\right).$$

Setting $\tau = 2 + 4/\epsilon$ and $\mathcal{E}(\epsilon) = 2(\epsilon+2)^2/\epsilon$ in the inequality above, we complete the proof. □

By choosing specific values for α and β, we can refine the performance bound described in Eq. (21). It follows from Eq. (17) we have

$$\min_{\alpha,\beta,\alpha+\beta=\gamma}\left(\frac{\alpha^2}{\kappa_1^2(2r)} + \frac{\beta^2}{\kappa_2^2(c)}\right) = \frac{\gamma^2}{\kappa_1^2(2r) + \kappa_2^2(c)}, \quad \gamma = \frac{2\sigma_{X(l)}\sigma_w\sqrt{n}}{N}\left(1 + \sqrt{\frac{k}{n}} + t\right), \tag{25}$$

where the equality of the first equation is achieved by setting α and β proportional to $\kappa_1^2(2r)$ and $\kappa_2^2(q)$, i.e., $\alpha = \gamma\kappa_1^2(2r)/\left(\kappa_1^2(2r) + \kappa_2^2(c)\right)$ and $\beta = \gamma\kappa_2^2(c)/\left(\kappa_1^2(2r) + \kappa_2^2(c)\right)$. Thus the performance bound in Eq. (21) can be refined as

$$\frac{1}{N}\|\mathcal{G}_X\widehat{\Theta} - \mathcal{F}\|_F^2 \leq (1+\epsilon)\inf_{\Theta}\left\{\frac{1}{N}\|\mathcal{G}_X\Theta - \mathcal{F}\|_F^2\right\} + \frac{4\mathcal{E}(\epsilon)\sigma_{X(l)}^2\sigma_w^2 n}{N^2\left(\kappa_1^2(2r) + \kappa_2^2(c)\right)}\left(1 + \sqrt{\frac{k}{n}} + t\right)^2.$$

Note that the performance bound above is independent of the value of α and β, and it is tighter than the one described in Eq. (21).

Table 1: Averaged performance (with standard derivation) comparison in terms of AUC, Macro F1, and Micro F1. Note that n, d, and m denote the sample size, dimensionality, and label number, respectively.

	Data Set (n, d, m)	Business (9968, 16621, 17)	Arts (7441, 17973, 19)	Health (9109, 18430, 14)	Scene (2407, 294, 6)
AUC	S.TraceNorm	85.42 ± 0.31	76.31 ± 0.15	86.18 ± 0.56	91.54 ± 0.18
	TraceNorm	83.43 ± 0.41	75.90 ± 0.27	85.24 ± 0.42	90.33 ± 0.24
	OneNorm	81.95 ± 0.26	70.47 ± 0.18	83.60 ± 0.32	88.42 ± 0.31
Macro F1	S.TraceNorm	48.83 ± 0.13	32.83 ± 0.25	60.05 ± 0.36	51.65 ± 0.33
	TraceNorm	47.24 ± 0.15	31.90 ± 0.31	58.91 ± 0.24	50.59 ± 0.08
	OneNorm	46.28 ± 0.25	31.03 ± 0.46	58.01 ± 0.18	46.57 ± 1.10
Micro F1	S.TraceNorm	78.26 ± 0.71	42.91 ± 0.27	67.22 ± 0.47	52.83 ± 0.35
	TraceNorm	78.84 ± 0.11	42.08 ± 0.11	66.92 ± 0.42	52.06 ± 0.49
	OneNorm	78.16 ± 0.17	40.64 ± 0.52	66.37 ± 0.19	47.32 ± 0.13

5 Experiments

In this section, we evaluate the effectiveness of the sparse trace norm regularization formulation in Eq. (6) on benchmark data sets; we also conduct numerical studies on the convergence of AG and two ADMM implementations including ADMM1 and ADMM2 (see details in Section E of the supplemental material) for solving Eq. (6) and the convergence of the alternating optimization algorithm for solve Eq. (10). Note that we use the least square loss for the following experiments.

Performance Evaluation We apply the sparse trace norm regularization formulation (S.TraceNorm) on multi-label classification problems, in comparison with the trace norm regularization formulation (TraceNorm) and the ℓ_1-norm regularization formulation (OneNorm). AUC, Macro F1, and Micro F1 are used as the classification performance measures. Four benchmark data sets, including Business, Arts, and Health from Yahoo webpage data sets [30] and Scene from LIBSVM multi-label data sets[1], are employed in this experiment. The reported experimental results are averaged over 10 random repetitions of the data sets into training and test sets of the ratio $1:9$. We use the AG method to solve the S.TraceNorm formulation, and stop the iterative procedure of AG if the change of the objective values in two successive iterations is smaller than 10^{-8} or the iteration numbers larger than 10^5. The regularization parameters α and β are determined via double cross-validation from the set $\{10^{-2} \times i\}_{i=1}^{10} \cup \{10^{-1} \times i\}_{i=2}^{10} \cup \{2 \times i\}_{i=1}^{10}$.

We present the averaged performance of the competing algorithms in Table 1. The main observations are summarized as follows: (1) S.TraceNorm achieves the best performance on all benchmark data sets (except on Business data) in this experiment; this result demonstrates the effectiveness of the induced sparse low-rank structure for multi-label classification tasks; (2) TraceNorm outperforms OneNorm on all benchmark data sets; this result demonstrates the effectiveness of modeling a shared low-rank structure for high-dimensional text and image data analysis.

Numerical Study We study the practical convergence of AG and ADMM2 by solving Eq. (6) on Scene data. In our experiments, we observe that ADMM1 is much slower than ADMM2 and we thus only focus on ADMM2. Note that in AG, we set $\alpha = 1, \beta = 1$; in ADMM2, we set $\alpha = 1$, $\beta = 1$, $\rho_1 = \rho_2 = 10$. For other parameter settings, we observe similar trends.

In the first experiment, we compare AG and ADMM2 in term of the practical convergence. We stop ADMM2 when the change of the objective values in two successive iterations smaller than 10^{-4}; the attained objective value in ADMM2 is used as the stopping criterion for AG, that is, we stop AG if the attained objective value in AG is equal to or smaller than that objective value attained in ADMM2. The convergence curves of ADMM2 and AG are presented in the left plot of Figure 1. Clearly, we can observe that AG converges much faster than ADMM2. In the second experiment, we study the convergence of AG. We stop AG when the change of the objective values in two successive iterations smaller than 10^{-8}. The convergence curves is presented in the middle plot of Figure 1. We observe that AG converges very fast, and its convergence speed is consistent with the theoretical convergence analysis in [19].

We also conduct numerical study on the alternating optimization algorithm (in Section 3.1.2) for solving the dual formulation of the proximal operator in Eq. (10). Similarly, the alternating optimization algorithm

---

[1]http://www.csie.ntu.edu.tw/~cjlin

Figure 1: Convergence comparison of AG and ADMM2 for solving Eq. (6) (left plot); convergence plot of AG for solving Eq. (6) (middle plot); and the alternating optimization algorithm for solving the dual formulation of the proximal operator in Eq. (10) (right plot).

is stopped when the change of the objective values in two successive iterations smaller than 10^{-8}. For illustration, in Eq. (10) we randomly generate the matrix $\widehat{\Phi}$ of size 10000 by 5000 from $\mathcal{N}(0,1)$; we then apply the alternating optimization algorithm to solve Eq. (10) and plot its convergence curve in the right plot of Figure 1. Our experimental results show that the alternating optimization algorithm generally converges within 10 iterations and our results demonstrate the practical efficiency of this algorithm.

6   Conclusion

We study the problem of estimating multiple predictive functions simultaneously in the nonparametric regression setting. In our estimation scheme, each predictive function is estimated using a linear combination of a dictionary of pre-specified basis functions. By assuming that the coefficient matrix admits a sparse low-rank structure, we formulate the function estimation problem as a convex program with the trace norm and the ℓ_1-norm regularization. We propose to employ AG and ADMM algorithms to solve the function estimation problem and also develop efficient algorithms for the key components involved in AG and ADMM. We derive a key property of the optimal solution to the convex program; moreover, based on an assumption associated with the basis functions, we establish a performance bound of the proposed function estimation scheme using the composite regularization. Our simulation studies demonstrate the effectiveness and the efficiency of the proposed formulation. In the future, we plan to derive a formal sparse oracle inequality for the convex problem in Eq. (6) as in [3]; we also plan to apply the proposed function estimation formulation to other real world applications.

References

[1] F. Bach. Consistency of trace norm minimization. *Journal of Machine Learning Research*, 9:1019–1048, 2008.

[2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Science*, 2:183–202, 2009.

[3] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *Annals of Statistics*, 37:1705–1732, 2009.

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 2010.

[5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[6] F. Bunea, A. B. Tsybakov, and M. H. Wegkamp. Aggregation and sparsity via ℓ_1 penalized least squares. In *COLT*, 2006.

[7] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *CoRR*, abs/0805.4471, 2008.

[8] E.J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of ACM*, 2011.

[9] E.J. Candès and T Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51:4203–4215, 2005.

[10] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Sparse and low-rank matrix decompositions. In *SYSID*, 2009.

[11] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–451, 2004.

[12] M. Fazel, H. Hindi, and S. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *ACL*, 2001.

[13] J. Friedman, T. Hastie, H. Hofling, and R. Tibshirani. Pathwise coordinate optimization. *Annals of Statistics*, 1:302–332, 2007.

[14] L. Grippoa and M. Sciandrone. On the convergence of the block nonlinear gauss-seidel method under convex constraints. *Operation Research Letters*, 26:127–136, 2000.

[15] J. Huang, T. Zhang, and D. N. Metaxas. Learning with structured sparsity. In *ICML*, 2009.

[16] K. Lounici, M. Pontil, A. B. Tsybakov, and S. van de Geer. Taking advantage of sparsity in multi-task learning. In *COLT*, 2008.

[17] J.-J. Moreau. Proximité et dualité dans un espace hilbertien. *Bull. Soc. Math. France*, 93:273–299, 1965.

[18] S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. In *ICML*, 2010.

[19] Y. Nesterov. Introductory lectures on convex programming. 1998. Lecture Notes.

[20] Y. Nesterov. Gradient methods for minimizing composite objective function. *CORE Discussion Paper*, 2007.

[21] Y. C. Pati and T. Kailath. Phase-shifting masks for microlithography: automated design and mask requirements. *Journal of the Optical Society of America A*, 11(9):2438–2452, 1994.

[22] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization. *SIAM Review*, (3):471–501, 2010.

[23] J. D. M. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *ICML*, 2005.

[24] A. Rohde and A. B. Tsybakov. Estimation of high-dimensional low rank matrices. *Preprint available at 0912.5338v2*, 2010.

[25] N. Srebro, J. D. M. Rennie, and T. Jaakkola. Maximum-margin matrix factorization. In *NIPS*, 2004.

[26] S. J. Szarek. Condition numbers of random matrices. *Journal of Complexity*, 7(2):131–149, 1991.

[27] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.

[28] K. C. Toh, M. J. Todd, and R.H. Tutuncu. SDPT3: a MATLAB software package for semidefinite programming. *Optimization Methods and Software*, 11:545–581, 1999.

[29] K. C. Toh and S. W. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *Pacific Journal of Optimization*, 2009.

[30] N. Ueda and K. Saito. Single-shot detection of multiple categories of text using parametric mixture models. In *KDD*, 2002.

[31] G. A. Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra and its Applications*, (170):33–45, 1992.

[32] T. Zhang. Some sharp performance bounds for least squares regression with l_1 regularization. *Annals of Statistics*, 37:2109–2144, 2009.

[33] P. Zhao and B. Yu. On model selection consistency of lasso. *Journal of Machine Learning Research*, 7:2541–2563, 2006.

Sparse Trace Norm Regularization: Supplemental Material

A. Operators $\mathcal{S}_0$ and $\mathcal{S}_1$

We define two operators, namely $\mathcal{S}_0$ and $\mathcal{S}_1$, on an arbitrary matrix pair (of the same size) based on Lemma 3.4 in [22], as summarized in the following lemma.

Lemma 1. *Given any Θ and Δ of size $h \times k$, let rank$(\Theta) = r$ and denote the SVD of Θ as*

$$\Theta = U \begin{bmatrix} \Sigma & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} V^T,$$

where $U \in \mathbb{R}^{h \times h}$ and $V \in \mathbb{R}^{k \times k}$ are orthogonal, and $\Sigma \in \mathbb{R}^{r \times r}$ is diagonal consisting of the non-zero singular values on its main diagonal. Let

$$\widehat{\Delta} = U^T \Delta V = \begin{bmatrix} \widehat{\Delta}_{11} & \widehat{\Delta}_{12} \\ \widehat{\Delta}_{21} & \widehat{\Delta}_{22} \end{bmatrix},$$

where $\widehat{\Delta}_{11} \in \mathbb{R}^{r \times r}$, $\widehat{\Delta}_{12} \in \mathbb{R}^{r \times (k-r)}$, $\widehat{\Delta}_{21} \in \mathbb{R}^{(h-r) \times r}$, and $\widehat{\Delta}_{22} \in \mathbb{R}^{(h-r) \times (k-r)}$. Define $\mathcal{S}_0$ and $\mathcal{S}_1$ as

$$\mathcal{S}_0(\Theta, \Delta) = U \begin{bmatrix} \widehat{\Delta}_{11} & \widehat{\Delta}_{12} \\ \widehat{\Delta}_{21} & \mathbf{0} \end{bmatrix} V^T, \ \mathcal{S}_1(\Theta, \Delta) = U \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \widehat{\Delta}_{22} \end{bmatrix} V^T.$$

Then the following conditions hold: rank$\left(\mathcal{S}_0(\Theta, \Delta)\right) \leq 2r$, $\Theta \mathcal{S}_1(\Theta, \Delta)^T = 0$, $\Theta^T \mathcal{S}_1(\Theta, \Delta) = 0$.

The result presented in Lemma 1 implies a condition under which the trace norm on a matrix pair is additive. From Lemma 1 we can easily verify that

$$\|\Theta + \mathcal{S}_1(\Theta, \Delta)\|_* = \|\Theta\|_* + \|\mathcal{S}_1(\Theta, \Delta)\|_*, \tag{26}$$

for arbitrary Θ and Δ of the same size. To avoid clutter notation, we denote $\mathcal{S}_0(\Theta, \Delta)$ by $\mathcal{S}_0(\Delta)$, and $\mathcal{S}_1(\Theta, \Delta)$ by $\mathcal{S}_1(\Delta)$ throughout this paper, as the appropriate Θ can be easily determined from the context.

B. Bound on Trace Norm

As a consequence of Lemma 1, we derive a bound on the trace norm of the matrices of interest as summarized below.

Corollary 1. *Given an arbitrary matrix pair $\widehat{\Theta}$ and Θ, let $\Delta = \widehat{\Theta} - \Theta$. Then*

$$\|\widehat{\Theta} - \Theta\|_* + \|\Theta\|_* - \|\widehat{\Theta}\|_* \leq 2\|\mathcal{S}_0(\Delta)\|_*.$$

Proof. From Lemma 1 we have $\Delta = \mathcal{S}_0(\Delta) + \mathcal{S}_1(\Delta)$ for the matrix pair Θ and Δ. Moreover,

$$
\begin{aligned}
\|\widehat{\Theta}\|_* &= \|\Theta + \mathcal{S}_0(\Delta) + \mathcal{S}_1(\Delta)\|_* \geq \|\Theta + \mathcal{S}_1(\Delta)\|_* - \|\mathcal{S}_0(\Delta)\|_* \\
&= \|\Theta\|_* + \|\mathcal{S}_1(\Delta)\|_* - \|\mathcal{S}_0(\Delta)\|_*,
\end{aligned}
\tag{27}
$$

where the inequality above follows from the triangle inequality and the last equality above follows from Eq. (26). Using the result in Eq. (27), we have

$$
\begin{aligned}
\|\widehat{\Theta} - \Theta\|_* + \|\Theta\|_* - \|\widehat{\Theta}\|_* &\leq \|\Delta\|_* + \|\Theta\|_* - \|\Theta\|_* - \|\mathcal{S}_1(\Delta)\|_* + \|\mathcal{S}_0(\Delta)\|_* \\
&\leq 2\|\mathcal{S}_0(\Delta)\|_*.
\end{aligned}
$$

We complete the proof of this corollary. $\qquad\square$

C. Bound on ℓ_1-norm

Analogous to the bound on the trace norm in Corollary 1, we also derive a bound on the ℓ_1-norm of the matrices of interest in the following lemma. For arbitrary matrices Θ and Δ, we denote by $J(\Theta) = \{(i,j)\}$ the coordinate set (the location set of nonzero entries) of Θ, and by $J(\Theta)_\perp$ the associated complement (the location set of zero entries); we denote by $\Delta_{J(\Theta)}$ the matrix of the same entries as Δ on the set $J(\Theta)$ and of zero entries on the set $J(\Theta)_\perp$. We now present a result associated with $J(\Theta)$ and $J(\Theta)_\perp$ in the following lemma. Note that a similar result for the vector case is presented in [3].

Lemma 2. *Given a matrix pair $\widehat{\Theta}$ and Θ of the same size, the inequality below always holds*

$$\|\widehat{\Theta} - \Theta\|_1 + \|\Theta\|_1 - \|\widehat{\Theta}\|_1 \leq 2\|\widehat{\Theta}_{J(\Theta)} - \Theta_{J(\Theta)}\|_1. \tag{28}$$

Proof. It can be verified that the inequality

$$\|\Theta_{J(\Theta)}\|_1 - \|\widehat{\Theta}_{J(\Theta)}\|_1 \leq \|(\widehat{\Theta} - \Theta)_{J(\Theta)}\|_1$$

and the equalities

$$\Theta_{J(\Theta)_\perp} = \mathbf{0}, \quad \|(\widehat{\Theta} - \Theta)_{J(\Theta)_\perp}\|_1 - \|\widehat{\Theta}_{J(\Theta)}\|_1 = \mathbf{0}$$

hold. Therefore we can derive

$$
\begin{aligned}
& \|\widehat{\Theta} - \Theta\|_1 + \|\Theta\|_1 - \|\widehat{\Theta}\|_1 \\
=\ & \|(\widehat{\Theta} - \Theta)_{J(\Theta)}\|_1 + \|(\widehat{\Theta} - \Theta)_{J(\Theta)_\perp}\|_1 + \|\Theta_{J(\Theta)}\|_1 + \|\Theta_{J(\Theta)_\perp}\|_1 - \|\widehat{\Theta}_{J(\Theta)}\|_1 - \|\widehat{\Theta}_{J(\Theta)_\perp}\|_1 \\
\leq\ & 2\|(\widehat{\Theta} - \Theta)_{J(\Theta)}\|_1.
\end{aligned}
$$

This completes the proof of this lemma. $\qquad\square$

D. Concentration Inequality

Lemma 3. *Let $\sigma_{X(l)}$ be the maximum singular value of the matrix $\mathcal{G}_X \in \mathbb{R}^{n \times h}$; let $W \in \mathbb{R}^{n \times k}$ be the matrix of i.i.d entries as $w_{ij} \sim \mathcal{N}(0, \sigma_w^2)$. Let $\lambda = 2\sigma_{X(l)}\sigma_w\sqrt{n}\left(1 + \sqrt{k/n} + t\right)/N$. Then*

$$\Pr\left(\|W^T\mathcal{G}_X\|_2/N \leq \lambda/2\right) \geq 1 - \exp\left(-nt^2/2\right).$$

Proof. It is known [26] that a Gaussian matrix $\widehat{W} \in \mathbb{R}^{n \times k}$ with $n \geq k$ and $\hat{w}_{ij} \sim \mathcal{N}(0, 1/n)$ satisfies

$$\Pr\left(\|\widehat{W}\|_2 > 1 + \sqrt{k/n} + t\right) \leq \exp\left(-nt^2/2\right), \tag{29}$$

where t is a universal constant. From the definition of the largest singular value, there exist a vector $b \in \mathbb{R}^h$ of length 1, i.e., $\|b\|_2 = 1$, such that $\|W^T\mathcal{G}_X\|_2 = \|W^T\mathcal{G}_X b\|_2 \leq \|W\|_2\|\mathcal{G}_X b\|_2 \leq \sigma_{X(l)}\|W\|_2$. Since $w_{ij}/(\sigma_w\sqrt{n}) \sim \mathcal{N}(0, 1/n)$, we have

$$\Pr\left(\left\|W^T\mathcal{G}_X\right\|_2/N > \lambda/2\right) \leq \Pr\left(\sigma_{X(l)}\|W\|_2/N > \lambda/2\right).$$

Applying the result in Eq. (29) into the inequality above, we complete the proof of this lemma. $\qquad\square$

E. Implementations of the Alternating Direction Method of Multipliers for Solving Eq. (6)

We employ two variants of the Alternating Direction Method of Multipliers (ADMM) to solve the Eq. (6). The key difference lies in the use of different numbers of auxiliary variables to separate the smooth components from the non-smooth components of the objective function in Eq. (6).

E.1 The First Implementation: ADMM1

By adding an auxiliary variable Ψ, we reformulate Eq. (6) as

$$\min_{\Theta,\Psi} \quad \widehat{S}(\Theta) + \alpha\|\Psi\|_* + \beta\|\Theta\|_1$$
$$\text{subject to} \quad \Theta = \Psi. \tag{30}$$

The augmented Lagrangian of Eq. (30) can be expressed as

$$\mathcal{L}_\rho^1(\Theta, \Psi, \Gamma) = \widehat{S}(\Theta) + \alpha\|\Psi\|_* + \beta\|\Theta\|_1 + \langle\Theta - \Psi, \Gamma\rangle + \frac{\rho}{2}\|\Theta - \Psi\|_F^2. \tag{31}$$

To solve Eq. (30), ADMM1 consists of the following iterations:

$$\Theta_{k+1} = \arg\min_\Theta \mathcal{L}_\rho^1(\Theta, \Psi_k, \Gamma_k), \tag{32}$$
$$\Psi_{k+1} = \arg\min_\Psi \mathcal{L}_\rho^1(\Theta_{k+1}, \Psi, \Gamma_k), \tag{33}$$
$$\Gamma_{k+1} = \Gamma_k + \rho(\Theta_{k+1} - \Psi_{k+1}), \tag{34}$$

where Θ_k, Ψ_k, and Γ_k denote the intermediate solutions of ADMM1 at the k-th iteration, and ρ is a pre-specified constant.

Specifically, if we employ the least squares loss, i.e., $\widehat{S}(\Theta) = \|\mathcal{G}_X\Theta - Y\|_F^2/N$, the optimization problems in Eqs. (32) and (34) can be efficiently solved as below.

Update on Θ The optimal Θ_{k+1} to Eq. (32) can be obtained via

$$\Theta_{k+1} = \arg\min_\Theta \left(\frac{1}{N}\|\mathcal{G}_X\Theta - Y\|_F^2 + \beta\|\Theta\|_1 + \langle\Theta, \Gamma_k\rangle + \frac{\rho}{2}\|\Theta - \Psi_k\|_F^2\right), \tag{35}$$

which can be efficiently solved via the gradient-type methods [2, 20].

Update on Ψ The optimal Ψ_{k+1} to Eq. (33) can be obtained via

$$\Psi_{k+1} = \arg\min_\Psi \left(\alpha\|\Psi\|_* - \langle\Psi, \Gamma_k\rangle + \frac{\rho}{2}\|\Theta_{k+1} - \Psi\|_F^2\right).$$

The optimization problem above admits an analytical solution [22]. Assume $\text{rank}(\Theta_{k+1} + \Gamma_k/\rho) = r$. Let $\Theta_{k+1} + \Gamma_k/\rho = U_r\Sigma_r V_r^T$ be the singular value decomposition of $\Theta_{k+1} + \Gamma_k/\rho$, where U_r and V_r consist of respectively r orthonormal columns, and $\Sigma_r = \text{diag}\{(\sigma_1, \sigma_2, \cdots, \sigma_r)\}$. Then the optimal Ψ_{k+1} is given by

$$\Psi_{k+1} = U_r\hat{\Sigma}V_r^T, \quad \hat{\Sigma} = \text{diag}\left\{\left(\sigma_i - \frac{\alpha}{\rho}\right)_+\right\}, \tag{36}$$

where $(x)_+ = x$ if $x > 0$ and $(x)_+ = 0$ otherwise.

E.2 The Second Implementation: ADMM2

By adding two auxiliary variables Ψ^1 and Ψ^2, we reformulate Eq. (6) as

$$\min_{\Theta,\Psi^1,\Psi^2} \quad \widehat{S}(\Theta) + \alpha\|\Psi^1\|_* + \beta\|\Psi^2\|_1$$
$$\text{subject to} \quad \Theta = \Psi^1, \ \Theta = \Psi^2. \tag{37}$$

Similarly, the augmented Lagrangian of Eq. (37) can be expressed as

$$\mathcal{L}_{\rho_1,\rho_2}^2(\Theta, \Psi^1, \Psi^2, \Gamma^1, \Gamma^2)$$
$$= \widehat{S}(\Theta) + \alpha\|\Psi^1\|_* + \beta\|\Psi^2\|_1 + \langle\Theta - \Psi^1, \Gamma^1\rangle + \langle\Theta - \Psi^2, \Gamma_2\rangle + \frac{\rho_1}{2}\|\Theta - \Psi^1\|_F^2 + \frac{\rho_2}{2}\|\Theta - \Psi^2\|_F^2.$$

To solve Eq. (37), ADMM2 consists of the following iterations:

$$\Theta_{k+1} = \arg\min_{\Theta} \mathcal{L}^2_{\rho_1,\rho_2}(\Theta, \Psi^1_k, \Psi^2_k, \Gamma^1_k, \Gamma^2_k), \tag{38}$$

$$\left(\Psi^1_{k+1}, \Psi^2_{k+1}\right) = \arg\min_{\Psi^1,\Psi^2} \mathcal{L}^2_{\rho_1,\rho_2}(\Theta_{k+1}, \Psi^1, \Psi^2, \Gamma^1_k, \Gamma^2_k), \tag{39}$$

$$\Gamma^1_{k+1} = \Gamma^1_k + \rho_1\left(\Theta_{k+1} - \Psi^1_{k+1}\right), \tag{40}$$

$$\Gamma^2_{k+1} = \Gamma^2_k + \rho_2\left(\Theta_{k+1} - \Psi^2_{k+1}\right), \tag{41}$$

where Θ_k, Ψ^1_k, Ψ^2_k, Γ^1_k, and Γ^2_k denote the intermediate solutions at the k-th iteration of the ADMM2 method.

Specifically, if we employ $\widehat{S}(\Theta) = \|\mathcal{G}_X\Theta - Y\|^2_F/N$ as the loss function in Eq. (37), the optimization problems in Eqs. (38), (39), (40), and (41) can be efficiently solved as below.

Update on Θ The optimal Θ_{k+1} to Eq. (38) can be obtained via

$$\Theta_{k+1} = \arg\min_{\Theta} \left(\frac{1}{N}\|\mathcal{G}_X\Theta - Y\|^2_F + \langle\Theta, \Gamma^1_k + \Gamma^2_k\rangle + \frac{\rho_1}{2}\|\Theta - \Psi^1_k\|^2_F + \frac{\rho_2}{2}\|\Theta - \Psi^2_k\|^2_F\right).$$

Note that the optimal Θ_{k+1} can be obtained via solving a systems of linear equations.

Update on Ψ^1 and Ψ^2 The optimal Ψ^1_{k+1} and Ψ^1_{k+1} to Eq. (39) can be obtained via

$$\Psi^1_{k+1} = \arg\min_{\Psi^1} \left(\alpha\|\Psi^1\|_* - \langle\Psi^1, \Gamma^1_k\rangle + \frac{\rho_1}{2}\|\Theta_{k+1} - \Psi^1\|^2_F\right), \tag{42}$$

$$\Psi^2_{k+1} = \arg\min_{\Psi^2} \left(\beta\|\Psi^2\|_1 - \langle\Psi^2, \Gamma^2_k\rangle + \frac{\rho_2}{2}\|\Theta_{k+1} - \Psi^2\|^2_F\right). \tag{43}$$

It can be verified that Eq. (42) admits an analytical solution. Assume rank $\left(\Theta_{k+1} + \Gamma^1_k/\rho_1\right) = r$. Let $\Theta_{k+1} + \Gamma^1_k/\rho_1 = U_r\Sigma_r V_r^T$ be the singular value decomposition of $\Theta_{k+1} + \Gamma^1_k/\rho_1$, where U_r and V_r consist of respectively r orthonormal columns, and $\Sigma_r = \text{diag}\{(\sigma_1, \sigma_2, \cdots, \sigma_r)\}$. Then the optimal Ψ^1_{k+1} is given by

$$\Psi^1_{k+1} = U_r\hat{\Sigma}V_r^T, \ \hat{\Sigma} = \text{diag}\left\{\left(\sigma_i - \frac{\alpha}{\rho_1}\right)_+\right\}, \tag{44}$$

where $(x)_+ = x$ if $x > 0$ and $(x)_+ = 0$ otherwise.

Moreover, it can also be verified that Eq. (43) admits an analytical solution. Let ψ, θ, and γ be the entries of Ψ^2_{k+1}, Θ_{k+1}, and Γ^2_k at the same coordinates. The optimal ψ is given by

$$\psi = \begin{cases} \theta + \frac{1}{\rho_2}(\gamma - \beta) & \theta + \frac{1}{\rho_2}\gamma > \frac{1}{\rho_2}\beta \\ 0 & -\frac{1}{\rho_2}\beta \le \theta + \frac{1}{\rho_2}\gamma \le \frac{1}{\rho_2}\beta \\ \theta + \frac{1}{\rho_2}(\gamma + \beta) & \theta + \frac{1}{\rho_2}\gamma < -\frac{1}{\rho_2}\beta \end{cases}. \tag{45}$$