
Dirty Statistical Models

Eunho Yang

Department of Computer Science
University of Texas at Austin
eunho@cs.utexas.edu

Pradeep Ravikumar

Department of Computer Science
University of Texas at Austin
pradeepr@cs.utexas.edu

Abstract

We provide a unified framework for the high-dimensional analysis of “superposition-structured” or “dirty” statistical models: where the model parameters are a *superposition* of structurally constrained parameters. We allow for any number and types of structures, and any statistical model. We consider the general class of M -estimators that minimize the sum of any loss function, and an instance of what we call a “hybrid” regularization, that is the infimal convolution of weighted regularization functions, one for each structural component. We provide corollaries showcasing our unified framework for varied statistical models such as linear regression, multiple regression and principal component analysis, over varied superposition structures.

1 Introduction

High-dimensional statistical models have been the subject of considerable focus over the past decade, both theoretically as well as in practice. In these high-dimensional models, the ambient dimension of the problem p may be of the same order as, or even substantially larger than the sample size n . It has now become well understood that even in this type of high-dimensional $p \gg n$ scaling, it is possible to obtain statistically consistent estimators provided one imposes structural constraints on the statistical models. Examples of such structural constraints include sparsity constraints (e.g. compressed sensing), graph-structure (for graphical model estimation), low-rank structure (for matrix-structured problems), and sparse additive structure (for non-parametric models), among others. For each of these structural constraints, a large body of work have proposed and analyzed statistically consistent estimators. For instance, a key subclass leverage such structural constraints via specific regularization functions. Examples include ℓ_1 -regularization for sparse models, nuclear norm regularization for low-rank matrix-structured models, and so on.

A caveat to this strong line of work is that imposing such “clean” structural constraints such as sparsity or low-rank structure, is typically too stringent for real-world messy data. What if the parameters are not exactly sparse, or not exactly low rank? Indeed, over the last couple of years, there has been an emerging line of work that address this caveat by “mixing and matching” different structures. Chandrasekaran et al. [5] consider the problem of recovering an unknown low-rank and an unknown sparse matrix, given the sum of the two matrices; for which they point to applications in system identification in linear time-invariant systems, and optical imaging systems among others. Chandrasekaran et al. [6] also apply this matrix decomposition estimation to the learning of latent-variable Gaussian graphical models, where they estimate an inverse covariance matrix that is the sum of sparse and low-rank matrices. A number of papers have applied such decomposition estimation to robust principal component analysis: Candès et al. [3] learn a covariance matrix that is the sum of a low-rank factored matrix and a sparse “error/outlier” matrix, while [9, 15] learn a covariance matrix that is the sum of a low-rank matrix and a column-sparse error matrix. Hsu et al. [7] analyze this estimation of a sum of a low-rank and elementwise sparse matrix in the noisy setting; while Agarwal et al. [1] extend this to the sum of a low-rank matrix and a matrix with general structure. Another application is multi-task learning, where [8] learn a multiple-linear-regression coefficient

matrix that is the sum of a sparse and a block-sparse matrix. This strong line of work can be seen to follow the resume of estimating a superposition of two structures; and indeed their results show this simple extension provides a vast increase in the practical applicability of structurally constrained models. The statistical guarantees in these papers for the corresponding M -estimators typically require fairly extensive technical arguments that extend the analyses of specific *single-structured* regularized estimators in highly non-trivial ways.

This long-line of work above on M -estimators and analyses for specific pairs of super-position structures for specific statistical models, lead to the question: is there a unified framework for studying *any general* tuple (i.e. not just a pair) of structures, for any *general statistical model*? This is precisely the focus of this paper: we provide a unified framework of “superposition-structured” or “dirty” statistical models, with any number and any types of structures, for any statistical model. By such “superposition-structure,” we mean the constraint that the parameter be a *superposition* of “clean” structurally constrained parameters. In addition to the motivation above, of unifying the burgeoning list of works above, as well as to provide guarantees for many novel superpositions (of for instance more than two structures) not yet considered in the literature; another key motivation is to provide insights on the key ingredients characterizing the statistical guarantees for such dirty statistical models. Our unified analysis allows the following very general class of M -estimators, which are the sum of *any loss function*, and an instance of what we call a “hybrid” regularization function, that is the infimal convolution of any weighted regularization functions, one for each structural component. As we show, this is equivalent to an M -estimator that is the sum of (a) a loss function applied to the sum of the multiple parameter vectors, one corresponding to each structural component; and (b) a weighted sum of regularization functions, one for each of the parameter vectors. We stress that our analysis allows for *general loss functions*, and *general component regularization functions*. We provide corollaries showcasing our unified framework for varied statistical models such as linear regression, multiple regression and principal component analysis, over varied superposition structures.

2 Problem Setup

We consider the following general statistical modeling setting. Consider a random variable Z with distribution \mathbb{P} , and suppose we are given n observations $Z_1^n := \{Z_1, \dots, Z_n\}$ drawn i.i.d. from \mathbb{P} . We are interested in estimating some parameter $\theta^* \in \mathbb{R}^p$ of the distribution \mathbb{P} . We assume that the statistical model parameter θ^* is “superposition-structured,” so that it is the sum of parameter components, each of which is constrained by a specific structure. For a formalization of the notion of structure, we first review some terminology from [11]. There, they use subspace pairs $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, where $\mathcal{M} \subseteq \overline{\mathcal{M}}$, to capture any structured parameter. \mathcal{M} is the *model subspace* that captures the constraints imposed on the model parameter, and is typically *low-dimensional*. $\overline{\mathcal{M}}^\perp$ is the *perturbation subspace* of parameters that represents perturbations away from the model subspace. They also define the property of *decomposability* of a regularization function, which captures the suitability of a regularization function \mathcal{R} to particular structure. Specifically, a regularization function \mathcal{R} is said to be *decomposable* with respect to a subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, if

$$\mathcal{R}(u + v) = \mathcal{R}(u) + \mathcal{R}(v), \quad \text{for all } u \in \mathcal{M}, v \in \overline{\mathcal{M}}^\perp.$$

For any structure such as sparsity, low-rank, etc., we can define the corresponding low-dimensional model subspaces, as well as regularization functions that are decomposable with respect to the corresponding subspace pairs.

I. Sparse vectors. Given any subset $S \subseteq \{1, \dots, p\}$ of the coordinates, let $\mathcal{M}(S)$ be the *subspace* of vectors in \mathbb{R}^p that have support contained in S . It can be seen that any parameter $\theta \in \mathcal{M}(S)$ would be atmost $|S|$ -sparse. For this case, we use $\overline{\mathcal{M}}(S) = \mathcal{M}(S)$, so that $\overline{\mathcal{M}}^\perp(S) = \mathcal{M}^\perp(S)$. As shown in [11], the ℓ_1 norm $\mathcal{R}(\theta) = \|\theta\|_1$, commonly used as a sparsity-encouraging regularization function, is decomposable with respect to subspace pairs $(\mathcal{M}(S), \overline{\mathcal{M}}^\perp(S))$.

II. Low-rank matrices. Consider the class of matrices $\Theta \in \mathbb{R}^{k \times m}$ that have rank $r \leq \min\{k, m\}$. For any given matrix Θ , we let $\text{row}(\Theta) \subseteq \mathbb{R}^m$ and $\text{col}(\Theta) \subseteq \mathbb{R}^k$ denote its row space and column space respectively. For a given pair of r -dimensional subspaces $U \subseteq \mathbb{R}^k$ and $V \subseteq \mathbb{R}^m$, we define the subspace pairs as follows: $\mathcal{M}(U, V) := \{\Theta \in \mathbb{R}^{k \times m} \mid \text{row}(\Theta) \subseteq V, \text{col}(\Theta) \subseteq U\}$ and

$\overline{\mathcal{M}}^\perp(U, V) := \{\Theta \in \mathbb{R}^{k \times m} \mid \text{row}(\Theta) \subseteq V^\perp, \text{col}(\Theta) \subseteq U^\perp\}$. As [11] show, the nuclear norm $\mathcal{R}(\theta) = \|\theta\|_1$ is decomposable with respect to the subspace pairs $(\mathcal{M}(U, V), \overline{\mathcal{M}}^\perp(U, V))$.

In our dirty statistical model setting, we do not just have one, but a set of structures; suppose we index them by the set I . Our key structural constraint can then be stated as: $\theta^* = \sum_{\alpha \in I} \theta_\alpha^*$, where θ_α^* is a “clean” structured parameter with respect to a subspace pair $(\mathcal{M}_\alpha, \overline{\mathcal{M}}_\alpha^\perp)$, for $\mathcal{M}_\alpha \subseteq \overline{\mathcal{M}}_\alpha$. We also assume we are given a set of regularization functions $\mathcal{R}_\alpha(\cdot)$, for $\alpha \in I$ that are suited to the respective structures, in the sense that they are decomposable with respect to the subspace pairs $(\mathcal{M}_\alpha, \overline{\mathcal{M}}_\alpha^\perp)$.

Let $\mathcal{L} : \Omega \times \mathcal{Z}^n \mapsto \mathbb{R}$ be some loss function that assigns a cost to any parameter $\theta \in \Omega \subseteq \mathbb{R}^p$, for a given set of observations Z_1^n . For ease of notation, in the sequel, we adopt the shorthand $\mathcal{L}(\theta)$ for $\mathcal{L}(\theta; Z_1^n)$. We are interested in the following “super-position” estimator:

$$\min_{(\theta_\alpha)_{\alpha \in I}} \mathcal{L}\left(\sum_{\alpha \in I} \theta_\alpha\right) + \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha), \quad (1)$$

where $(\lambda_\alpha)_{\alpha \in I}$ are the regularization penalties. This optimization problem involves not just one parameter vector, but multiple parameter vectors, one for each structural component: while the loss function applies only to the sum of these, separate regularization functions are applied to the corresponding parameter vectors. We will now see that this can be re-written to a standard M -estimation problem which minimizes, over a *single parameter vector*, the sum of a loss function and a special “dirty” regularization function.

Given a vector $\mathbf{c} := (c_\alpha)_{\alpha \in I}$ of convex-combination weights, suppose we define the following “dirty” regularization function, that is the infimal convolution of a set of regularization functions:

$$\mathcal{R}(\theta; \mathbf{c}) = \inf \left\{ \sum_{\alpha \in I} c_\alpha \mathcal{R}_\alpha(\theta_\alpha) : \sum_{\alpha \in I} \theta_\alpha = \theta \right\}. \quad (2)$$

It can be shown that provided the individual regularization functions $\mathcal{R}_\alpha(\cdot)$, for $\alpha \in I$, are *norms*, $\mathcal{R}(\cdot; \mathbf{c})$ is a norm as well. We discuss this and other properties of this hybrid regularization function $\mathcal{R}(\cdot; \mathbf{c})$ in Appendix A.

Proposition 1. *Suppose $(\hat{\theta}_\alpha)_{\alpha \in I}$ is the solution to the M -estimation problem in (1). Then $\hat{\theta} := \sum_{\alpha \in I} \hat{\theta}_\alpha$ is the solution to the following problem:*

$$\min_{\theta \in \Omega} \mathcal{L}(\theta) + \lambda \mathcal{R}(\theta; \mathbf{c}), \quad (3)$$

where $c_\alpha = \lambda_\alpha / \lambda$. Similarly, if $\hat{\theta}$ is the solution to (3), then there is a solution $(\hat{\theta}_\alpha)_{\alpha \in I}$ to the M -estimation problem (1), such that $\hat{\theta} := \sum_{\alpha \in I} \hat{\theta}_\alpha$.

Proposition 1 shows that the optimization problems (1) and (3) are equivalent. While the tuning parameters in (1) correspond to the regularization penalties $(\lambda_\alpha)_{\alpha \in I}$, the tuning parameters in (3) correspond to the weights $(c_\alpha)_{\alpha \in I}$ specifying the “dirty” regularization function. In our unified analysis theorem, we will provide guidance on setting these tuning parameters as a function of various model-parameters.

3 Error Bounds for Convex M -estimators

Our goal is to provide error bounds $\|\hat{\theta} - \theta^*\|$, between the target parameter θ^* , the minimizer of the population risk, and our M -estimate $\hat{\theta}$ from (1), for *any error norm* $\|\cdot\|$. A common example of an error norm for instance is the ℓ_2 norm $\|\cdot\|_2$. We now turn to the properties of the loss function and regularization function that underlie our analysis. We first restate some natural assumptions on the loss and regularization functions.

(C1) The loss function \mathcal{L} is convex and differentiable.

(C2) The regularizers \mathcal{R}_α are norms, and are decomposable with respect to the subspace pairs $(\mathcal{M}_\alpha, \overline{\mathcal{M}}_\alpha^\perp)$, where $\mathcal{M}_\alpha \subseteq \overline{\mathcal{M}}_\alpha$.

Our next assumption is a *restricted strong convexity* assumption [11]. Specifically, we will require the loss function \mathcal{L} to satisfy:

(C3) (Restricted Strong Convexity) For all $\Delta_\alpha \in \Omega_\alpha$, where Ω_α is the parameter space for the parameter component α ,

$$\delta\mathcal{L}(\Delta_\alpha; \theta^*) := \mathcal{L}(\theta^* + \Delta_\alpha) - \mathcal{L}(\theta^*) - \langle \nabla_{\theta} \mathcal{L}(\theta^*), \Delta_\alpha \rangle \geq \kappa_{\mathcal{L}} \|\Delta_\alpha\|^2 - g_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha),$$

where $\kappa_{\mathcal{L}}$ is a ‘‘curvature’’ parameter, and $g_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha)$ is a ‘‘tolerance’’ parameter.

Note that these conditions (C1)-(C3) are imposed even when the model has a single clean structural constraint; see [11]. Note that g_α is usually a function on the problem size decreasing in the sample size; in the standard Lasso with $|I| = 1$ for instance, $g_\alpha = \frac{\log p}{n}$.

Our next assumption is on the interaction between the different structured components.

(C4) (Structural Incoherence) For all $\Delta_\alpha \in \Omega_\alpha$,

$$\left| \mathcal{L}(\theta^* + \sum_{\alpha \in I} \Delta_\alpha) + (|I| - 1)\mathcal{L}(\theta^*) - \sum_{\alpha \in I} \mathcal{L}(\theta^* + \Delta_\alpha) \right| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 + \sum_{\alpha \in I} h_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha).$$

Note that for a model with a single clean structural constraint, with $|I| = 1$, the condition (C4) is trivially satisfied since the LHS becomes 0. We will see in the sequel that for a large collection of loss functions including all linear loss functions, the condition (C4) simplifies considerably, and moreover holds with high probability, typically with $h_\alpha = 0$. We note that this condition is much weaker than ‘‘incoherence’’ conditions typically imposed when analyzing specific instances of such superposition-structured models (see e.g. references in the introduction), where the assumptions typically include (a) assuming that the structured subspaces $(\mathcal{M}_\alpha)_{\alpha \in I}$ intersect only at $\{0\}$, and (b) that the sizes of these subspaces are extremely small.

Finally, we will use the notion of *subspace compatibility constant* defined in [11], that captures the relationship between the regularization function $\mathcal{R}(\cdot)$ and the error norm $\|\cdot\|$, over vectors in the subspace \mathcal{M} : $\Psi(\mathcal{M}, \|\cdot\|) := \sup_{u \in \mathcal{M} \setminus \{0\}} \frac{\mathcal{R}(u)}{\|u\|}$.

Theorem 1. *Suppose we solve the M -estimation problem in (3), with hybrid regularization $\mathcal{R}(\cdot; \mathbf{c})$, where the convex-combination weights \mathbf{c} are set as $c_\alpha = \lambda_\alpha / \sum_{\alpha \in I} \lambda_\alpha$, with $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha} \mathcal{L}(\theta^*; Z_1^n))$. Further, suppose conditions (C1) - (C4) are satisfied. Then, the parameter error bounds are given as:*

$$\|\widehat{\theta} - \theta^*\| \leq \left(\frac{3|I|}{2\bar{\kappa}} \right) \max_{\alpha \in I} \lambda_\alpha \Psi_\alpha(\overline{\mathcal{M}}_\alpha) + (|I|\sqrt{\tau_{\mathcal{L}}}/\sqrt{\bar{\kappa}}),$$

where

$$\begin{aligned} \bar{\kappa} &:= \frac{\kappa_{\mathcal{L}}}{2} - 32\bar{g}^2|I| \left(\max_{\alpha \in I} \lambda_\alpha \Psi_\alpha(\overline{\mathcal{M}}_\alpha) \right)^2, & \bar{g} &:= \max_{\alpha} \frac{1}{\lambda_\alpha} \sqrt{g_\alpha + h_\alpha}, \\ \tau_{\mathcal{L}} &:= \sum_{\alpha \in I} \left[32\bar{g}^2 \lambda_\alpha^2 \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta^*)) + \frac{2\lambda_\alpha}{|I|} \mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta^*)) \right]. \end{aligned}$$

Remarks: (R1) It is instructive to compare Theorem 1 to the main Theorem in [11], where they derive parameter error bounds for any M -estimator with a decomposable regularizer, for any ‘‘clean’’ structure. Our theorem can be viewed as a generalization: we recover their theorem when we have a single structure with $|I| = 1$. We cannot derive our result in turn from their theorem applied to the M -estimator (3) with the hybrid regularization function $\mathcal{R}(\cdot; \mathbf{c})$: the ‘‘superposition’’ structure is not captured by a pair of subspaces, nor is the hybrid regularization function decomposable, as is required by their theorem. Our setting as well as analysis is strictly more general, because of which we needed the additional structural incoherence assumption (C4) (which is trivially satisfied when $|I| = 1$).

(R2) Agarwal et al. [1] provide Frobenius norm error bounds for the matrix-decomposition problem of recovering the sum of low-rank and a general structured matrix. In addition to the greater generality of our theorem and framework, Theorem 1 addresses two key drawbacks of their theorem even in their specific setting. First, the proof for their theorem requires the regularization

penalty λ for the second structure to be strongly bounded away from zero: their convergence rate does not approach zero even with infinite number of samples n . Theorem 1, in contrast, imposes the weaker condition $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha}\mathcal{L}(\theta^*; Z_1^n))$, which as we show in the corollaries, allows for the convergence rates to go to zero as a function of the samples. Second, they assumed much stronger conditions for their theorem to hold; in Theorem 1 in contrast, we pose much milder “local” RSC conditions (C3), and a structural incoherence condition (C4).

(R3) The statement in the theorem is *deterministic* for fixed choices of (λ_α) . We also note that the theorem holds for any set of subspace pairs $(\mathcal{M}_\alpha, \overline{\mathcal{M}}_\alpha^\perp)_{\alpha \in I}$ with respect to which the corresponding regularizers are decomposable. As noted earlier, the \mathcal{M}_α should ideally be set to the structured subspace in which the true parameter at least approximately lies, and which we want to be as small as possible (note that the bound includes a term that depends on the size of this subspace via the subspace compatibility constant). In particular, if we assume that the subspaces are chosen so that $\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*) = 0$ i.e. $\theta_\alpha^* \in \mathcal{M}_\alpha$, then we obtain the simpler bound in the following corollary.

Corollary 1. *Suppose we solve the M-estimation problem in (1), with hybrid regularization $\mathcal{R}(\cdot; \mathbf{c})$, where the convex-combination weights \mathbf{c} are set as $c_\alpha = \lambda_\alpha / \sum_{\alpha \in I} \lambda_\alpha$, with $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha}\mathcal{L}(\theta^*; Z_1^n))$, and suppose conditions (C1) - (C4) are satisfied. Further, suppose that the subspace-pairs are chosen so that $\theta_\alpha^* \in \mathcal{M}_\alpha$. Then, the parameter error bounds are given as:*

$$\|\hat{\theta} - \theta^*\| \leq \left(\frac{3|I|}{2\kappa} \right) \max_{\alpha \in I} \lambda_\alpha \Psi_\alpha(\overline{\mathcal{M}}_\alpha).$$

It is now instructive to compare the bounds of Theorem 1, and Corollary 1. Theorem 1 has two terms: the first of which is the sole term in the bound in Corollary 1. This first term can be thought of as the “estimation error” component of the error bound, when the parameter has exactly the structure being modeled by the regularizers. The second term can be thought of as the “approximation error” component of the error bound, which is the penalty for the parameter not exactly lying in the structured subspaces modeled by the regularizers. The key term in the “estimation error” component, in Theorem 1, and Corollary 1, is:

$$\Phi = \max_{\alpha \in I} \lambda_\alpha \Psi_\alpha(\overline{\mathcal{M}}_\alpha).$$

Note that each λ_α is larger than a particular norm of the sample score function (gradient of the loss at the true parameter): since the expected value of the score function is zero, the magnitude of the sample score function captures the amount of “noise” in the data. This is in turn scaled by $\Psi_\alpha(\overline{\mathcal{M}}_\alpha)$, which captures the size of the structured subspace corresponding to the parameter component θ_α^* . Φ can thus be thought of as capturing the amount of noise in the data *relative* to the particular structure at hand.

We now provide corollaries showcasing our unified framework for varied statistical models such as linear regression, multiple regression and principal component analysis, over varied superposition structures.

4 Convergence Rates for Linear Regression

In this section, we consider the linear regression model:

$$Y = X\theta^* + w, \tag{4}$$

where $Y \in \mathbb{R}^n$ is the observation vector, and $\theta^* \in \mathbb{R}^p$ is the true parameter. $X \in \mathbb{R}^{n \times p}$ is the “observation” matrix; while $w \in \mathbb{R}^n$ is the observation noise. For this class of statistical models, we will consider the instantiation of (1) with the loss function \mathcal{L} consisting of the squared loss:

$$\min_{(\theta_\alpha)_{\alpha \in I}} \left\{ \frac{1}{n} \|Y - X(\sum_{\alpha \in I} \theta_\alpha)\|_2^2 + \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha) \right\}. \tag{5}$$

For this regularized least squared estimator (5), conditions (C1-C2) in Theorem 1 trivially hold. The *restricted strong convexity* condition (C3) reduces to the following. Noting that $\mathcal{L}(\theta^* + \Delta_\alpha) - \mathcal{L}(\theta^*) - \langle \nabla_{\theta} \mathcal{L}(\theta^*), \Delta_\alpha \rangle = \frac{1}{n} \|X\Delta_\alpha\|_2^2$, we obtain the following *restricted eigenvalue* condition:

(D3) $\frac{1}{n} \|X \Delta_\alpha\|_2^2 \geq \kappa_{\mathcal{L}} \|\Delta_\alpha\|^2 - g_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha)$ for all $\Delta_\alpha \in \Omega_\alpha$.

Finally, our *structural incoherence* condition reduces to the following: Noting that $|\mathcal{L}(\theta^* + \sum_{\alpha \in I} \Delta_\alpha) + (|I| - 1)\mathcal{L}(\theta^*) - \sum_{\alpha \in I} \mathcal{L}(\theta^* + \Delta_\alpha)| = \frac{2}{n} |\sum_{\alpha < \beta} \langle X \Delta_\alpha, X \Delta_\beta \rangle|$ in this specific case,

(D4) $\frac{2}{n} |\sum_{\alpha < \beta} \langle X \Delta_\alpha, X \Delta_\beta \rangle| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 + \sum_{\alpha \in I} h_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha)$.

4.1 Structural Incoherence with Gaussian Design

We now show that the condition (D4) required for Theorem 1, holds with high probability when the observation matrix is drawn from a so-called Σ -Gaussian ensemble: where each row X_i is independently sampled from $N(0, \Sigma)$. Before doing so, we first state some assumption on the population covariance matrix Σ . Let \mathcal{P}_M denote the matrix corresponding to the projection operator for the subspace M . We will then require the following assumption:

(C-Linear) Let $\Lambda := \max_{\gamma_1, \gamma_2} \left\{ 2 + \frac{3\lambda_{\gamma_1} \Psi_{\gamma_1}(\bar{\mathcal{M}}_{\gamma_1})}{\lambda_{\gamma_2} \Psi_{\gamma_2}(\bar{\mathcal{M}}_{\gamma_2})} \right\}$. For any $\alpha, \beta \in I$,

$$\max \left\{ \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_\alpha} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_\beta} \right), \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_\alpha} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_\beta^\perp} \right), \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_\alpha^\perp} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_\beta} \right) \right\} \leq \frac{\kappa_{\mathcal{L}}}{8 \binom{|I|}{2} \Lambda^2 |I|}. \quad (6)$$

Proposition 2. *Suppose each row X_i of the observation matrix X is independently sampled from $N(0, \Sigma)$, and the condition (C-Linear) (6) holds. Further, suppose that $\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\theta_\alpha^*) = 0$, for all $\alpha \in I$. Then, it holds that with probability at least $1 - \frac{4}{\max\{n, p\}}$,*

$$\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X \Delta_\alpha, X \Delta_\beta \rangle \right| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha} \|\Delta_\alpha\|_2^2,$$

when the number of samples scales as $n \geq c \left(\frac{\binom{|I|}{2} \Lambda^2 |I|}{\kappa_{\mathcal{L}}} \right)^2 (\max_{\alpha} \Psi_{\alpha}(\bar{\mathcal{M}}_{\alpha})^2 + \max\{\log p, \log n\})$, for some constant c that depends only on the distribution of X .

Condition (D3) is the usual restricted eigenvalue condition which has been analyzed previously in “clean-structured” model estimation, so that we can directly appeal to previous results [10, 12] to show that it holds with high probability when the observation matrix is drawn from the Σ -Gaussian ensemble.

We are now ready to derive the consequences of the deterministic bound in Theorem 1 for the case of the linear regression model above.

4.2 Linear Regression with Sparse and Group-sparse structures

We now consider the following superposition structure, comprised of both sparse and group-sparse structures. Suppose that a set of groups $\mathcal{G} = \{G_1, G_2, \dots, G_q\}$ are disjoint subsets of the index-set $\{1, \dots, p\}$, each of size at most $|G_i| \leq m$. Suppose that the linear regression parameter θ^* is a superposition of a group-sparse component θ_g^* with respect to this set of groups \mathcal{G} , as well as a sparse component θ_s^* with respect to the remaining indices $\{1, \dots, p\} \setminus \cup_{i=1}^q G_i$, so that $\theta^* = \theta_g^* + \theta_s^*$. Then, we use the hybrid regularization function $\sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha) = \lambda_s \|\theta_s\|_1 + \lambda_g \|\theta_g\|_{1,a}$ where $\|\theta\|_{1,a} := \sum_{t=1}^q \|\theta_{G_t}\|_a$ for $a \geq 2$.

Corollary 2. *Consider the linear model (4) where θ^* is the sum of exact s -sparse θ_s^* and exact s_g group-sparse θ_g^* . Suppose that each row X_i of the observation matrix X is independently sampled from $N(0, \Sigma)$. Further, suppose that (6) holds and w is sub-Gaussian with parameter σ . Then, if we solve (5) with*

$$\lambda_s = 8\sigma \sqrt{\frac{\log p}{n}} \quad \text{and} \quad \lambda_g = 8\sigma \left\{ \frac{m^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log q}{n}} \right\},$$

then, with probability at least $1 - c_1 \exp(-c_2 n \lambda_s^2) - c_3/q^2$, we have the error bound:

$$\|\hat{\theta} - \theta^*\|_2 \leq \frac{24\sigma}{\bar{\kappa}} \max \left\{ \sqrt{\frac{s \log p}{n}}, \frac{\sqrt{s_g} m^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{s_g \log q}{n}} \right\}.$$

Let us briefly compare the result from Corollary 2 with those from single-structured regularized estimators. Since the total sparsity of θ^* is bounded by $\|\theta\|_0 \leq ms_g + s$, “clean” ℓ_1 regularized least squares, with high probability, gives the bound [11]: $\|\widehat{\theta}_{\ell_1} - \theta^*\|_2 = O\left(\sqrt{\frac{(ms_g + s)\log p}{n}}\right)$. On the other hand, the support of θ^* also can be interpreted as comprising $s_g + s$ disjoint groups in the worst case, so that “clean” ℓ_1/ℓ_2 group regularization entails, with high probability, the bound [11]: $\|\widehat{\theta}_{\ell_1/\ell_2} - \theta^*\|_2 = O\left(\sqrt{\frac{(s_g + s)m}{n}} + \sqrt{\frac{(s_g + s)\log q}{n}}\right)$. We can easily verify that Corollary 2 achieves better bounds, considering the fact $p \leq mq$.

5 Convergence Rates for Multiple Regression

In this section, we consider the multiple linear regression model, with m linear regressions written jointly as

$$Y = X\Theta^* + W, \quad (7)$$

where $Y \in \mathbb{R}^{n \times m}$ is the observation matrix: with each column corresponding to a separate linear regression task, and $\Theta^* \in \mathbb{R}^{p \times m}$ is the collated set of parameters. $X \in \mathbb{R}^{n \times p}$ is the “observation” matrix; while $W \in \mathbb{R}^{n \times m}$ is collated set of observation noise vectors. For this class of statistical models, we will consider the instantiation of (1) with the loss function \mathcal{L} consisting of the squared loss:

$$\min_{\{\Theta_\alpha\}_{\alpha \in I}} \left\{ \frac{1}{n} \|Y - X \left(\sum_{\alpha \in I} \Theta_\alpha \right)\|_F^2 + \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Theta_\alpha) \right\}. \quad (8)$$

In contrast to the linear regression model in the previous section, the model (7) has a *matrix-structured* parameter; nonetheless conditions (C3-C4) in Theorem 1 reduce to the following conditions that are very similar to those in the previous section, with the Frobenius norm replacing the ℓ_2 norm:

- (D3) $\frac{1}{n} \|X\Delta_\alpha\|_F^2 \geq \kappa_{\mathcal{L}} \|\Delta_\alpha\|^2 - g_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha)$ for all $\Delta_\alpha \in \Omega_\alpha$.
- (D4) $\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X\Delta_\alpha, X\Delta_\beta \rangle \right| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 + \sum_{\alpha \in I} h_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha)$.

where the notation $\langle A, B \rangle$ denotes the trace inner product, $\text{trace}(A^\top B) = \sum_i \sum_j A_{ij} B_{ij}$.

As in the previous linear regression example, we again impose the assumption (C-Linear) on the population covariance matrix of a Σ -Gaussian ensemble, but in this case with the notational change of $\mathcal{P}_{\bar{\mathcal{M}}_\alpha}$ denoting the matrix corresponding to projection operator onto the *row-spaces* of matrices in $\bar{\mathcal{M}}_\alpha$. Thus, with the low-rank matrix structure discussed in Section 2, we would have $\mathcal{P}_{\bar{\mathcal{M}}_\alpha} = UU^\top$. Under the (C-Linear) assumption, the following proposition then extends Proposition 2:

Proposition 3. *Consider the problem (8) with the matrix parameter Θ . Under the same assumptions as in Proposition 2, we have with probability at least $1 - \frac{4}{\max\{n, p\}}$,*

$$\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X\Delta_\alpha, X\Delta_\beta \rangle \right| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha} \|\Delta_\alpha\|_F^2.$$

Consider an instance of this multiple linear regression model with the superposition structure consisting of row-sparse, column-sparse and elementwise sparse matrices: $\Theta^* = \Theta_r^* + \Theta_c^* + \Theta_s^*$. In order to obtain estimators for this model, we use the hybrid regularization function $\sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha) = \lambda_r \|\Theta_r\|_{r,a} + \lambda_c \|\Theta_c\|_{c,a} + \lambda_s \|\Theta_s\|_1$ where $\|\cdot\|_{r,a}$ denotes the sum of ℓ_a norm of rows for $a \geq 2$, and similarly $\|\cdot\|_{c,a}$ is the sum of ℓ_a norm of columns, and $\|\cdot\|_1$ is entrywise ℓ_1 norm for matrix.

Corollary 3. *Consider the multiple linear regression model (7) where Θ^* is the sum of Θ_r^* with s_r nonzero rows, Θ_c^* with s_c nonzero columns, and Θ_s^* with s nonzero elements. Suppose that the design matrix X is Σ -Gaussian ensemble with the properties of column normalization and $\sigma_{\max}(X) \leq \sqrt{n}$. Further, suppose that (6) holds and W is elementwise sub-Gaussian with parameter σ . Then, if we solve (8) with*

$$\lambda_s = 8\sigma \sqrt{\frac{\log p + \log m}{n}}, \quad \lambda_r = 8\sigma \left\{ \frac{m^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log p}{n}} \right\}, \quad \text{and} \quad \lambda_c = 8\sigma \left\{ \frac{p^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log m}{n}} \right\},$$

with probability at least $1 - c_1 \exp(-c_2 n \lambda_s^2) - \frac{c_3}{p^2} - \frac{c_3}{m^2}$, the error of the estimate $\widehat{\Theta}$ is bounded as:

$$\|\widehat{\Theta} - \Theta^*\|_2 \leq \frac{36\sigma}{\bar{\kappa}} \max \left\{ \sqrt{\frac{s(\log p + \log m)}{n}}, \frac{\sqrt{s_r m^{1-1/a}}}{\sqrt{n}} + \sqrt{\frac{s_r \log p}{n}}, \frac{\sqrt{s_c p^{1-1/a}}}{\sqrt{n}} + \sqrt{\frac{s_c \log m}{n}} \right\}.$$

6 Convergence Rates for Principal Component Analysis

In this section, we consider the robust/noisy principal component analysis problem, where we are given n i.i.d. random vectors $Z_i \in \mathbb{R}^p$ where $Z_i = U_i + v_i$, $U_i \sim N(0, \Theta^*)$ is the ‘‘uncorrupted’’ set of observations, with a low-rank covariance matrix $\Theta^* = LL^T$, for some loading matrix $L \in \mathbb{R}^{p \times r}$. $v_i \in \mathbb{R}^p$ is a noise/error vector; in standard factor analysis, v_i is a spherical Gaussian noise vector: $v_i \sim N(0, \sigma^2 I_{p \times p})$ (or $v_i = 0$); and the goal is to recover the loading matrix L from samples.

In PCA with sparse noise, $v_i \sim N(0, \Gamma^*)$, where Γ^* is elementwise sparse. In this case, the covariance matrix of Z_i has the form $\Sigma = \Theta^* + \Gamma^*$, where Θ^* is low-rank, and Γ^* is sparse. We can thus write the sample covariance model as: $Y := \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T = \Theta^* + \Gamma^* + W$, where $W \in \mathbb{R}^{p \times p}$ is a Wishart distributed random matrix. For this class of statistical models, we will consider the following instantiation of (1):

$$\min_{(\Theta, \Gamma)} \left\{ \|Y - \Theta - \Gamma\|_F^2 + \lambda_\Theta \|\Theta\|_1 + \lambda_\Gamma \|\Gamma\|_1 \right\}. \quad (9)$$

where $\|\cdot\|_1$ denotes the nuclear norm while $\|\cdot\|_1$ does the element-wise ℓ_1 norm (we will use $\|\cdot\|_2$ for the spectral norm.).

In contrast to the previous two examples, (9) includes a trivial design matrix, $X = I_{p \times p}$, which allows (D4) to hold under the simpler (C-linear) condition: where Λ is $\max_{\gamma_1, \gamma_2} \left\{ 2 + \frac{3\lambda_{\gamma_1} \Psi_{\gamma_1}(\mathcal{M}_{\gamma_1})}{\lambda_{\gamma_2} \Psi_{\gamma_2}(\mathcal{M}_{\gamma_2})} \right\}$, $\max \left\{ \sigma_{\max}(\mathcal{P}_{\mathcal{M}_\Theta} \mathcal{P}_{\mathcal{M}_\Gamma}), \sigma_{\max}(\mathcal{P}_{\mathcal{M}_\Theta} \mathcal{P}_{\mathcal{M}_\Gamma^\perp}), \sigma_{\max}(\mathcal{P}_{\mathcal{M}_\Theta^\perp} \mathcal{P}_{\mathcal{M}_\Gamma}), \sigma_{\max}(\mathcal{P}_{\mathcal{M}_\Theta^\perp} \mathcal{P}_{\mathcal{M}_\Gamma^\perp}) \right\} \leq \frac{1}{16\Lambda^2}$. (10)

Corollary 4. Consider the principal component analysis model where Θ^* has the rank r at most and Γ^* has s nonzero entries. Suppose that (10) holds. Then, given the choice of

$$\lambda_\Theta = 16\sqrt{\|\Sigma\|_2} \sqrt{\frac{p}{n}}, \quad \lambda_\Gamma = 32\rho(\Sigma) \sqrt{\frac{\log p}{n}},$$

where $\rho(\Sigma) = \max_j \Sigma_{jj}$, the optimal error of (9) is bounded by

$$\|\widehat{\Theta} - \Theta^*\|_2 \leq \frac{48}{\bar{\kappa}} \max \left\{ \sqrt{\|\Sigma\|_2} \sqrt{\frac{rp}{n}}, 2\rho(\Sigma) \sqrt{\frac{s \log p}{n}} \right\},$$

with probability at least $1 - c_1 \exp(-c_2 \log p)$.

Remarks. Agarwal et al. [1] also analyze this model, and propose to use the M -estimator in (9), with the additional constraint of $\|\Theta\|_\infty \leq \frac{\alpha}{p}$. Under a stricter ‘‘global’’ RSC condition, they

compute the error bound $\|\widehat{\Theta} - \Theta^*\|_2 \asymp \max \left\{ \sqrt{\|\Sigma\|_2} \sqrt{\frac{rp}{n}}, \rho(\Sigma) \sqrt{\frac{s \log p}{n}} + \frac{\alpha}{p} \right\}$ where α is a parameter between 1 and p . This bound is similar to that in Corollary 4, but with an additional term $\frac{\alpha}{p}$, so that it does not go to zero as a function of n . It also faces a trade-off: a smaller value of α to reduce error bound would make the assumption on the maximum element of Θ^* stronger as well. Our corollaries do not suffer these lacunae; see also our remarks in (R2) in Theorem 1. [14] extended the result of [1] to the special case where $\Theta^* = \Theta_r^* + \Theta_s^*$ using the notation of the previous section; the remarks above also apply here. Note that our work and [1] derive Frobenius error bounds under restricted strong convexity conditions; other recent works such as [7] also derive such Frobenius error bounds but under stronger conditions (see [1] for details).

Acknowledgments

We acknowledge the support of ARO via W911NF-12-1-0390 and NSF via IIS-1149803, DMS-1264033.

References

- [1] A. Agarwal, S. Negahban, and M. J. Wainwright. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *Annals of Statistics*, 40(2):1171–1197, 2012.
- [2] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [3] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3), May 2011.
- [4] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. In *48th Annual Allerton Conference on Communication, Control and Computing*, 2010.
- [5] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2), 2011.
- [6] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky. Latent variable graphical model selection via convex optimization. *Annals of Statistics (with discussion)*, 40(4), 2012.
- [7] D. Hsu, S. M. Kakade, and T. Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Trans. Inform. Theory*, 57:7221–7234, 2011.
- [8] A. Jalali, P. Ravikumar, S. Sanghavi, and C. Ruan. A dirty model for multi-task learning. In *Neur. Info. Proc. Sys. (NIPS)*, 23, 2010.
- [9] M. McCoy and J. A. Tropp. Two proposals for robust pca using semidefinite programming. *Electron. J. Statist.*, 5:1123–1160, 2011.
- [10] S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Annals of Statistics*, 39(2):1069–1097, 2011.
- [11] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [12] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research (JMLR)*, 99:2241–2259, 2010.
- [13] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [14] H. Xu and C. Leng. Robust multi-task regression with grossly corrupted observations. *Inter. Conf. on AI and Statistics (AISTATS)*, 2012.
- [15] H. Xu, C. Caramanis, and S. Sanghavi. Robust pca via outlier pursuit. *IEEE Transactions on Information Theory*, 58(5):3047–3064, 2012.

Appendix

A Properties of Hybrid Regularization Function $\mathcal{R}(\cdot; \mathbf{c})$

Let us briefly study some properties of this hybrid regularization function $\mathcal{R}(\cdot; \mathbf{c})$. In the remainder of this section, we suppress the dependence of the notation on the weights \mathbf{c} . From (2), it can be seen that $\mathcal{R}(\cdot)$ satisfies the triangle inequality:

$$\mathcal{R}(\theta + \theta') \leq \mathcal{R}(\theta) + \mathcal{R}(\theta'). \quad (11)$$

Under the additional assumption that each $\mathcal{R}_\alpha(\cdot)$ is 1-sub-homogeneous, so that $\mathcal{R}_\alpha(\beta\theta) \leq \beta\mathcal{R}_\alpha(\theta)$, we get that:

$$\mathcal{R}(\beta\theta + (1 - \beta)\theta') \leq \mathcal{R}(\beta\theta) + \mathcal{R}((1 - \beta)\theta') \leq \beta\mathcal{R}(\theta) + (1 - \beta)\mathcal{R}(\theta'), \quad (12)$$

so that \mathcal{R} is convex. Moreover, provided the individual regularization functions $\mathcal{R}_\alpha(\cdot)$, for $\alpha \in I$, are non-negative, it can be seen that $\mathcal{R}(\cdot)$ is non-negative as well.

From the above properties, it follows that provided the the individual regularization functions $\mathcal{R}_\alpha(\cdot)$, for $\alpha \in I$, are *norms*, $\mathcal{R}(\cdot)$ is a norm as well.

We can also derive the dual of the hybrid regularization function, assuming it is a norm:

$$\begin{aligned} \mathcal{R}^*(u) &= \sup_{\theta} \frac{\langle \theta, u \rangle}{\mathcal{R}(\theta)} = \sup_{(\theta_\alpha)} \frac{\sum_{\alpha} \langle u, \theta_\alpha \rangle}{\sum_{\alpha} c_\alpha \mathcal{R}_\alpha(\theta_\alpha)} = \sup_{(\theta_\alpha)} \frac{\sum_{\alpha} \langle u/c_\alpha, \theta_\alpha \rangle}{\sum_{\alpha} \mathcal{R}_\alpha(\theta_\alpha)} \\ &\leq \sup_{(\theta_\alpha)} \frac{\sum_{\alpha} \mathcal{R}_\alpha^*(u/c_\alpha) \mathcal{R}_\alpha(\theta_\alpha)}{\sum_{\alpha} \mathcal{R}_\alpha(\theta_\alpha)} \leq \max_{\alpha \in I} \mathcal{R}_\alpha^*(u)/c_\alpha. \end{aligned}$$

Noting that the inequalities become equalities by setting $\theta'_\alpha = 0$, for all α' not attaining the maximum in the last expression, we obtain the required expression for the dual norm:

$$\mathcal{R}^*(u) = \max_{\alpha \in I} \mathcal{R}_\alpha^*(u)/c_\alpha. \quad (13)$$

On the other hand, its Fenchel conjugate (again, assuming it is a norm) is given by:

$$\begin{aligned} \mathcal{R}_f(u) &= \sup_{\theta} \left\{ \langle \theta, u \rangle - \mathcal{R}(\theta) \right\} = \sup_{\theta} \left\{ \langle \theta, u \rangle - \inf_{(\theta_\alpha)_{\alpha \in I}: \sum_{\alpha \in I} \theta_\alpha = \theta} \sum_{\alpha \in I} c_\alpha \mathcal{R}_\alpha(\theta_\alpha) \right\} \\ &= \sum_{\alpha \in I} \sup_{\theta_\alpha} \left\{ \langle \theta, u/c_\alpha \rangle - \mathcal{R}_\alpha(\theta_\alpha) \right\} = \sum_{\alpha \in I} \mathcal{R}_{\alpha;f}(u)/c_\alpha. \end{aligned}$$

B Proof of Theorem 1

Throughout the proof, we make use of the fact that the optimal error vector $\hat{\theta} - \theta^*$ is guaranteed to belong to a specific set, as a consequence of the decomposability of the respective regularization functions:

Proposition 4. *Suppose that conditions (C1) and (C2) are satisfied. Then for any optimal solution $\hat{\theta}$ of (1), with the regularization penalties satisfying $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha} \mathcal{L}(\theta^*; Z_1^n))$, the error $\hat{\Delta}$ lies in the set*

$$\begin{aligned} \mathbb{C}(\mathcal{M}_1, \bar{\mathcal{M}}_1^\perp, \dots, \mathcal{M}_{|I|}, \bar{\mathcal{M}}_{|I|}^\perp; \theta^*) &:= \left\{ (\Delta_1, \dots, \Delta_{|I|}) \in \Omega_1 \times \dots \times \Omega_{|I|} \mid \right. \\ &\left. \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \leq \sum_{\alpha \in I} \lambda_\alpha \left[3\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \right\} \quad (14) \end{aligned}$$

See Appendix B.1 for the proof of this claim.

To set up the first crucial ingredient of the proof, we define the function $F : \Omega_1 \times \dots \times \Omega_{|I|} \mapsto \mathbb{R}$:

$$F(\Delta_1, \dots, \Delta_{|I|}) := \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) + \sum_{\alpha \in I} \left[\lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha^* + \Delta_\alpha) - \mathcal{R}_\alpha(\theta_\alpha^*) \right]. \quad (15)$$

By virtue of its optimality, it can be seen that the error $(\widehat{\Delta}_1, \dots, \widehat{\Delta}_{|I|})$ of the optimal solution $(\widehat{\theta}_1, \dots, \widehat{\theta}_{|I|})$ satisfies $F(\widehat{\Delta}_1, \dots, \widehat{\Delta}_{|I|}) \leq 0$. Also note that $F(\vec{0}, \dots, \vec{0}) = 0$.

As the following lemma shows, in order to compute an upper bound on the optimal error, $\sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\|$, it suffices to control the sign of function F over the following set:

$$\mathbb{K}(\delta) := \mathbb{C} \cap \left\{ \sum_{\alpha \in I} \|\Delta_\alpha\| = \delta \right\}$$

Lemma 1. *Suppose that conditions (C1) and (C2) are satisfied. If $F(\Delta_1, \dots, \Delta_{|I|}) > 0$ for all possible vectors $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{K}(\delta)$, then the optimal error $(\widehat{\Delta}_1, \dots, \widehat{\Delta}_{|I|})$ satisfies $\sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\| \leq \delta$.*

See Appendix B.2 for the proof of this claim.

The other crucial ingredient of the proof, is that we show the following ‘‘global’’ restricted strong convexity condition given the ‘‘local’’ restricted strong convexity conditions in (C3) for just individual structures:

Lemma 2. *Suppose that conditions (C3) and (C4) are satisfied for the M -estimation problem (1). Then, for $(\theta_1, \dots, \theta_{|I|}) \in (\Omega_1 \times \dots \times \Omega_{|I|})$ the following ‘‘global’’ restricted strong convexity (RSC) condition holds:*

$$\begin{aligned} \delta \mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*) &:= \mathcal{L}\left(\sum_{\alpha \in I} (\theta_\alpha^* + \Delta_\alpha)\right) - \mathcal{L}\left(\sum_{\alpha \in I} \theta_\alpha^*\right) - \sum_{\alpha \in I} \langle \nabla_\theta \mathcal{L}\left(\sum_{\alpha \in I} \theta_\alpha^*\right), \Delta_\alpha \rangle \\ &\geq \left[\frac{\kappa_{\mathcal{L}}}{2} - 32\bar{g}^2 |I| \Phi^2 \right] \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - 32\bar{g}^2 |I| \sum_{\alpha \in I} \left[\lambda_\alpha^2 \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right], \end{aligned} \quad (16)$$

where \bar{g} is defined as $\max_\alpha \frac{1}{\lambda_\alpha} \sqrt{g_\alpha + h_\alpha}$.

See Appendix B.3 for the proof of this claim.

We will next prove the following bound on $\sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\|$:

$$\sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\| \leq \frac{|I|}{\bar{\kappa}} \left(\frac{3}{2} \Phi + \sqrt{\bar{\kappa} \tau_{\mathcal{L}}} \right), \quad (17)$$

where $\bar{\kappa} := \frac{\kappa_{\mathcal{L}}}{2} - 32\bar{g}^2 |I| \Phi^2$ and $\tau_{\mathcal{L}} := \sum_{\alpha \in I} \left[32\bar{g}^2 \lambda_\alpha^2 \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) + \frac{2\lambda_\alpha}{|I|} \mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right]$.

Given this bound, we can use the following inequality and complete the proof:

$$\begin{aligned} \|\widehat{\theta} - \theta^*\| &= \left\| \sum_{\alpha} \widehat{\theta}_\alpha - \sum_{\alpha} \theta_\alpha^* \right\| = \left\| \sum_{\alpha} (\widehat{\theta}_\alpha - \theta_\alpha^*) \right\| \\ &\leq \sum_{\alpha} \|\widehat{\theta}_\alpha - \theta_\alpha^*\| \leq \sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\| \leq \frac{|I|}{\bar{\kappa}} \left(\frac{3}{2} \Phi + \sqrt{\bar{\kappa} \tau_{\mathcal{L}}} \right). \end{aligned}$$

It thus remains to show the bound (17).

As shown in Lemma 2, given conditions (C3) and (C4), the loss function satisfies the following ‘‘global’’ RSC condition with curvature $\bar{\kappa}$ and tolerance $\bar{\tau}$:

$$\delta \mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*) \geq \bar{\kappa} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - \bar{\tau}(\theta^*).$$

Then, by the construction of F , for an arbitrary error vector $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{C}$, we have

$$F(\Delta_1, \dots, \Delta_{|I|}) \geq \sum_{\alpha \in I} \langle \nabla_\theta \mathcal{L}(\theta^*), \Delta_\alpha \rangle + \bar{\kappa} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - \bar{\tau}(\theta^*) + \sum_{\alpha \in I} \left[\lambda_\alpha \mathcal{R}_\alpha(\theta_\alpha^* + \Delta_\alpha) - \mathcal{R}_\alpha(\theta_\alpha^*) \right].$$

Using the bounds (18) and (19), we obtain

$$\begin{aligned}
F(\Delta_1, \dots, \Delta_{|I|}) &\geq - \sum_{\alpha \in I} \frac{\lambda_\alpha}{2} \left[\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \right] + \bar{\kappa} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - \bar{\tau}(\theta^*) \\
&\quad + \sum_{\alpha \in I} \lambda_\alpha \left[\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) - \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) - 2\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \\
&\geq \sum_{\alpha \in I} \left[\bar{\kappa} \|\Delta_\alpha\|^2 - \frac{\lambda_\alpha}{2} \left[3\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \right] - \bar{\tau}(\theta^*),
\end{aligned}$$

where in the last inequality we dropped the term $\frac{\lambda_\alpha}{2} \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha))$, since it is always positive.

The following proposition allows us to reduce the task of bounding a multivariate function to that of bounding a univariate function.

Proposition 5. *Consider the K -variate quadratic function: $F(x_1, \dots, x_K) = \sum_{k=1}^K (ax_k^2 + bx_k) + c$ for some constants a, b and c . Suppose further that: $a > 0$, $x_k \geq 0$ for all k and $\sum_k x_k = \delta > 0$. Then $F(x_1, \dots, x_K)$ attains its minimum value at $x_1 = \dots = x_K = \delta/K$.*

Denote $\Phi := \max_{\alpha \in I} \lambda_\alpha \Psi_\alpha(\bar{\mathcal{M}}_\alpha)$ and $D := \frac{1}{|I|} \bar{\tau}(\theta^*) + \sum_{\alpha \in I} \frac{2\lambda_\alpha}{|I|} \mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*))$. Armed with these notations, and using the definition of subspace compatibility constant,

$$\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) \leq \Psi_\alpha(\bar{\mathcal{M}}_\alpha) \|\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)\| \leq \Psi_\alpha(\bar{\mathcal{M}}_\alpha) \|\Delta_\alpha\|,$$

we obtain:

$$F(\Delta_1, \dots, \Delta_{|I|}) \geq \sum_{\alpha \in I} (\bar{\kappa} \|\Delta_\alpha\|^2 - \frac{3}{2} \Phi \|\Delta_\alpha\| - D).$$

Now consider an error vector $(\Delta_1, \dots, \Delta_{|I|})$ s.t. $\sum_{\alpha \in I} \|\Delta_\alpha\| = \delta$. From Proposition 5,

$$F(\Delta_1, \dots, \Delta_{|I|}) \geq \bar{\kappa} \delta^2 - \frac{3}{2} \Phi \delta - D|I|.$$

It follows that $F(\Delta_1, \dots, \Delta_{|I|}) > 0$ so long as $\delta > \frac{|I|}{\bar{\kappa}} (\frac{3}{2} \Phi + \sqrt{\bar{\kappa} D})$. Therefore, for all error vectors $(\Delta_1, \dots, \Delta_{|I|})$ such that $\sum_{\alpha \in I} \|\Delta_\alpha\| \geq \frac{|I|}{\bar{\kappa}} (\frac{3}{2} \Phi + \sqrt{\bar{\kappa} D})$, and in particular, if $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{K}(\frac{|I|}{\bar{\kappa}} (\frac{3}{2} \Phi + \sqrt{\bar{\kappa} D}))$, we can guarantee $F(\Delta_1, \dots, \Delta_{|I|}) > 0$. This satisfies the conditions for Lemma 1, whose statement then completes the proof.

B.1 Proof of Proposition 4

In the proof, we use the fact that $F(\hat{\Delta}_1, \dots, \hat{\Delta}_{|I|}) \leq 0$.

For any decomposable regularizer \mathcal{R}_α on $(\mathcal{M}_\alpha, \bar{\mathcal{M}}_\alpha^\perp)$, it is known that the following inequality holds (See [11] for the proof):

$$\mathcal{R}_\alpha(\theta_\alpha^* + \Delta_\alpha) - \mathcal{R}_\alpha(\theta_\alpha^*) \geq \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) - \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) - 2\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)). \quad (18)$$

At the same time, by the convexity of \mathcal{L} , we have

$$\begin{aligned}
\mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) &= \mathcal{L}\left(\sum_{\alpha \in I} \theta_\alpha^* + \sum_{\alpha \in I} \Delta_\alpha\right) - \mathcal{L}\left(\sum_{\alpha \in I} \theta_\alpha^*\right) \\
&\geq \sum_{\alpha \in I} \langle \nabla_{\theta_\alpha} \mathcal{L}(\theta^*), \Delta_\alpha \rangle \\
&\stackrel{(i)}{\geq} - \sum_{\alpha \in I} \mathcal{R}_\alpha^*(\nabla_{\theta_\alpha} \mathcal{L}(\theta^*)) \mathcal{R}_\alpha(\Delta_\alpha) \\
&\stackrel{(ii)}{\geq} - \sum_{\alpha \in I} \frac{\lambda_\alpha}{2} \left[\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \right] \quad (19)
\end{aligned}$$

where the inequality (i) comes from the generalized Cauchy-Schwarz inequality, and the inequality (ii) from the triangular inequality and the assumption in the statement.

Combining (18) and (19) yields

$$0 \geq F(\widehat{\Delta}_1, \dots, \widehat{\Delta}_{|I|}) \geq \sum_{\alpha \in I} \frac{\lambda_\alpha}{2} \left[\mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(\widehat{\Delta}_\alpha)) - 3\mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha}(\widehat{\Delta}_\alpha)) - 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right],$$

as claimed.

B.2 Proof of Lemma 1

We first show that the set \mathbb{C} has some special structure, and then show that that structure guarantees the error upper bounded as in the statement.

Let $(\Delta_1, \dots, \Delta_{|I|})$ be an arbitrary error vectors in the set \mathbb{C} . Then, for any $t \in (0, 1)$, we have

$$\begin{aligned} & \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(t\Delta_\alpha)) \\ \stackrel{(i)}{=} & \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(t\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \\ \stackrel{(ii)}{=} & t \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \\ \stackrel{(iii)}{\leq} & t \sum_{\alpha \in I} \lambda_\alpha \left[3\mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha}(\Delta_\alpha)) + 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \\ \stackrel{(iv)}{=} & \sum_{\alpha \in I} \lambda_\alpha \left[3\mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha}(t\Delta_\alpha)) + 4t\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \\ \stackrel{(v)}{\leq} & \sum_{\alpha \in I} \lambda_\alpha \left[3\mathcal{R}_\alpha(\Pi_{\overline{\mathcal{M}}_\alpha}(t\Delta_\alpha)) + 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \end{aligned} \quad (20)$$

where step (i) uses the fact that

$$\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(t\Delta_\alpha) = \operatorname{argmin}_{\gamma \in \overline{\mathcal{M}}_\alpha^\perp} \|t\Delta_\alpha - \gamma\| = t \operatorname{argmin}_{\gamma \in \overline{\mathcal{M}}_\alpha^\perp} \|\Delta_\alpha - \frac{\gamma}{t}\| = t\Pi_{\overline{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha),$$

step (ii) uses the positive homogeneity of norms, and step (iii) holds since $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{C}$. Moreover, step (iv) holds similarly as in equalities (i) and (ii), and finally step (v) trivially holds for any $t \leq 1$. Therefore, if $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{C}$, then the line segment $\{(t\Delta_1, \dots, t\Delta_{|I|}) \mid t \in (0, 1)\}$ between $(\Delta_1, \dots, \Delta_{|I|})$ and the origin $(\vec{0}, \dots, \vec{0})$ also lies in \mathbb{C} .

Now, we show the statement in the lemma by its contrapositive; suppose $\sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\| > \delta$. Since $\sum_{\alpha \in I} \|t\widehat{\Delta}_\alpha\| = t \sum_{\alpha \in I} \|\widehat{\Delta}_\alpha\|$, there exists some constant $t^* \in (0, 1)$ s.t. $(t^*\widehat{\Delta}_1, \dots, t^*\widehat{\Delta}_{|I|}) \in \mathbb{K}(\delta)$. At the same time, by the convexity of \mathcal{L} and the regularizers,

$$F(t^*\widehat{\Delta}_1, \dots, t^*\widehat{\Delta}_{|I|}) \leq t^*F(\widehat{\Delta}_1, \dots, \widehat{\Delta}_{|I|}) + (1-t^*)F(\vec{0}, \dots, \vec{0}) \leq 0.$$

Therefore, $(t^*\widehat{\Delta}_1, \dots, t^*\widehat{\Delta}_{|I|})$ is in $\mathbb{K}(\delta)$ such that $F(t^*\widehat{\Delta}_1, \dots, t^*\widehat{\Delta}_{|I|}) \leq 0$ by construction. Hence, by its contrapositive, the claim follows.

B.3 Proof of Lemma 2

The definition of $\delta\mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*)$ can be rewritten as

$$\begin{aligned} \delta\mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*) &= \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) - \sum_{\alpha \in I} \langle \nabla_\theta \mathcal{L}(\theta^*), \Delta_\alpha \rangle \\ &= \mathcal{L}(\theta^* + \Delta) + (|I| - 1)\mathcal{L}(\theta^*) - \sum_{\alpha \in I} \mathcal{L}(\theta^* + \Delta_\alpha) \\ &\quad + \sum_{\alpha \in I} \left[\mathcal{L}(\theta^* + \Delta_\alpha) - \mathcal{L}(\theta^*) - \langle \nabla_\theta \mathcal{L}(\theta^*), \Delta_\alpha \rangle \right]. \end{aligned}$$

Then, simply by the inequalities in (C3) and (C4), $\delta\mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*)$ can be lower-bounded by

$$\begin{aligned} & \sum_{\alpha \in I} \left[\kappa_{\mathcal{L}} \|\Delta_\alpha\|^2 - g_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha) \right] - \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - \sum_{\alpha \in I} h_\alpha \mathcal{R}_\alpha^2(\Delta_\alpha) \\ & \geq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - \sum_{\alpha \in I} (g_\alpha + h_\alpha) \mathcal{R}_\alpha^2(\Delta_\alpha). \end{aligned} \quad (21)$$

Now, let us focus on the term $\sum_{\alpha \in I} (g_\alpha + h_\alpha) \mathcal{R}_\alpha^2(\Delta_\alpha)$ in the RHS of (21). By the basic property of the square, we have the following inequalities:

$$\sum_{\alpha \in I} (g_\alpha + h_\alpha) \mathcal{R}_\alpha^2(\Delta_\alpha) \leq \left(\sum_{\alpha \in I} \sqrt{g_\alpha + h_\alpha} \mathcal{R}_\alpha(\Delta_\alpha) \right)^2 \leq \left(\bar{g} \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Delta_\alpha) \right)^2 \quad (22)$$

where in the second inequality we use $\langle x, y \rangle \leq \|x\|_\infty \|y\|_1$ and $\bar{g} := \max_\alpha \frac{1}{\lambda_\alpha} \sqrt{g_\alpha + h_\alpha}$. Since by Lemma 4, for any $(\Delta_1, \dots, \Delta_{|I|}) \in \mathbb{C}$,

$$\begin{aligned} & \sum_{\alpha \in I} \lambda_\alpha \mathcal{R}_\alpha(\Delta_\alpha) \leq \sum_{\alpha \in I} \lambda_\alpha \left[\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + \mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\Delta_\alpha)) \right] \\ & \leq \sum_{\alpha \in I} \lambda_\alpha \left[4\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + 4\mathcal{R}_\alpha(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right], \end{aligned}$$

the LHS of (22), $\sum_{\alpha \in I} (g_\alpha + h_\alpha) \mathcal{R}_\alpha^2(\Delta_\alpha)$, can be upper-bounded by

$$\begin{aligned} & 32\bar{g}^2 |I| \sum_{\alpha \in I} \lambda_\alpha^2 \left[\mathcal{R}_\alpha^2(\Pi_{\bar{\mathcal{M}}_\alpha}(\Delta_\alpha)) + \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right] \\ & \leq 32\bar{g}^2 |I| \sum_{\alpha \in I} \lambda_\alpha^2 \left[\Psi_\alpha^2 \|\Delta_\alpha\|^2 + \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)) \right]. \end{aligned} \quad (23)$$

By plugging (23) back into (21), we can construct the RSC condition as stated in Lemma 2 and complete the proof:

$$\delta\mathcal{L}(\Delta_1, \dots, \Delta_{|I|}; \theta^*) \geq \left[\frac{\kappa_{\mathcal{L}}}{2} - 32\bar{g}^2 |I| \Phi^2 \right] \sum_{\alpha \in I} \|\Delta_\alpha\|^2 - 32\bar{g}^2 |I| \sum_{\alpha \in I} \lambda_\alpha^2 \mathcal{R}_\alpha^2(\Pi_{\mathcal{M}_\alpha^\perp}(\theta_\alpha^*)).$$

C Proof of Proposition 2

Given the pair of complementary subspaces $(\bar{\mathcal{M}}, \bar{\mathcal{M}}^\perp)$, the error vector Δ_α can be written as $\Delta_\alpha = \Pi_{\bar{\mathcal{M}}}(\Delta_\alpha) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta_\alpha)$, we have

$$X\Delta_\alpha = X(\Pi_{\bar{\mathcal{M}}}(\Delta_\alpha) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta_\alpha)) = (X\mathcal{P}_{\bar{\mathcal{M}}})\Pi_{\bar{\mathcal{M}}}(\Delta_\alpha) + (X\mathcal{P}_{\bar{\mathcal{M}}^\perp})\Pi_{\bar{\mathcal{M}}^\perp}(\Delta_\alpha). \quad (24)$$

We will now assume that the regularization functions \mathcal{R}_α are atomic norms [4] with respect to the orthonormal basis vectors used to describe the subspaces $(\mathcal{M}_\alpha, \bar{\mathcal{M}}_\alpha)$; though the proof generalizes. Specifically, suppose there is a set of orthonormal vectors or ‘‘atoms’’ $\mathcal{A}_\alpha := \{\mathbf{a}_j\}$, such that the regularization function \mathcal{R}_α can be written as

$$\mathcal{R}_\alpha(\theta_\alpha) = \inf_{\mathbf{c}} \left\{ \sum_j c_j : \theta_\alpha = \sum_j c_j \mathbf{a}_j \right\}. \quad (25)$$

We will also assume that the subspace pair $(\mathcal{M}_\alpha, \bar{\mathcal{M}}_\alpha^\perp)$ contain disjoint orthonormal subsets of the basis set; it can be seen that \mathcal{R}_α is decomposable with respect to any such subspace pair.

Let t be the dimensionality of $\bar{\mathcal{M}}_\alpha$. Suppose $\{a_1, \dots, a_J\} \subseteq \mathcal{A}_\alpha$ are a set of atoms, such that the first t atoms characterize $\bar{\mathcal{M}}_\alpha$, and the remaining $J - t$ atoms characterize $\bar{\mathcal{M}}_\alpha^\perp$. Note that we can then write Δ_α as

$$\Delta_\alpha = \sum_{j=1}^J c_j a_j.$$

We now split the index set $\{1, \dots, J\}$ into subsets $S(1), \dots, S(K)$, such that $S(1) = \{1, \dots, t\}$ contains the atoms characterizing \mathcal{M}_α , the second set $S(2)$ contains the largest t entries from the remaining coefficients (c_{t+1}, \dots, c_J) , $S(3)$ contains the largest set of t entries disjoint from $S(1)$ and $S(2)$ and so on. Let $\bar{\mathcal{M}}_i$ denote the subspace characterized by the atoms indexed by set $S(i)$ suppressing the dependency on α . With these notations, we can rewrite (24) as

$$X\Delta_\alpha = X \sum_i \sum_{j \in S(i)} c_j \mathbf{a}_j = \sum_i \mathcal{P}_{\bar{\mathcal{M}}_i} X \times \left(\sum_{j \in S(i)} c_j \mathbf{a}_j \right). \quad (26)$$

Turning to the condition (C4), we have

$$\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X\Delta_\alpha, X\Delta_\beta \rangle \right| \leq \frac{2}{n} \sum_{\alpha < \beta} |\langle X\Delta_\alpha, X\Delta_\beta \rangle|. \quad (27)$$

For any pair $\alpha < \beta$ in the index set I , we use the equality (26) as follows:

$$\begin{aligned} & \frac{2}{n} |\langle X\Delta_\alpha, X\Delta_\beta \rangle| \\ &= \frac{2}{n} \left| \left\langle \sum_i \mathcal{P}_{\bar{\mathcal{M}}_i} X \left(\sum_{j \in S(i)} c_j \mathbf{a}_j \right), \sum_k \mathcal{P}_{\bar{\mathcal{M}}_k} X \left(\sum_{l \in S(k)} d_l \mathbf{b}_l \right) \right\rangle \right| \\ &= \frac{2}{n} \left| \sum_{i,k} \left\langle \mathcal{P}_{\bar{\mathcal{M}}_i} X \left(\sum_{j \in S(i)} c_j \mathbf{a}_j \right), \mathcal{P}_{\bar{\mathcal{M}}_k} X \left(\sum_{l \in S(k)} d_l \mathbf{b}_l \right) \right\rangle \right| \\ &\leq \frac{2}{n} \sum_{i,k} \left| \left\langle \mathcal{P}_{\bar{\mathcal{M}}_i} X \left(\sum_{j \in S(i)} c_j \mathbf{a}_j \right), \mathcal{P}_{\bar{\mathcal{M}}_k} X \left(\sum_{l \in S(k)} d_l \mathbf{b}_l \right) \right\rangle \right| \\ &\stackrel{(i)}{\leq} 2 \left[\max_{i,k} \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \right] \sum_{i,k} \left\| \left(\sum_{j \in S(i)} c_j \mathbf{a}_j \right) \right\|_2 \cdot \left\| \left(\sum_{l \in S(k)} d_l \mathbf{b}_l \right) \right\|_2 \end{aligned}$$

where the inequality (i) holds by the multiple applications of Cauchy-Schwarz inequality inequalities and Rayleigh quotient. For now, we define $C_{\max} := 2 \left[\max_{i,k} \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \right]$. Note that $\mathcal{P}_{\bar{\mathcal{M}}_i}$ denotes the projector matrix corresponding to the structure α while $\mathcal{P}_{\bar{\mathcal{M}}_k}$ does to the structure β .

Using the fact that the basis vectors are orthonormal, so that $\|\mathbf{a}_j\|_2 = 1$ and $\Pi_{\mathbf{a}_j}(\mathbf{a}_{j'}) = 0$ for all $j \neq j'$, we obtain

$$\frac{2}{n} |\langle X\Delta_\alpha, X\Delta_\beta \rangle| \leq C_{\max} \left(\sum_i \left\| \left(\sum_{j \in S(i)} c_j \right) \right\|_2 \right) \left(\sum_k \left\| \left(\sum_{l \in S(k)} d_l \right) \right\|_2 \right). \quad (28)$$

To upper-bound $\sum_i \left\| \left(\sum_{j \in S(i)} c_j \right) \right\|_2$, we can directly appeal to the following standard bound in [2].

$$\sum_{i=3} \left\| \sum_{j \in S(i)} c_j \right\|_2 \leq t^{-1/2} \sum_{i=2} \sum_{j \in S(i)} c_j.$$

Since $\sum_{i=2} \sum_{j \in S(i)} c_j$ is equal to $\mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha))$ by construction, we have

$$\sum_i \left\| \sum_{j \in S(i)} c_j \right\|_2 \leq 2\|\Delta_\alpha\|_2 + t^{-1/2} \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)).$$

In addition, since Δ_α should belong to \mathbb{C} in (14),

$$\lambda_\alpha \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \leq \sum_{\gamma \in I} \lambda_\gamma \left[3\mathcal{R}_\gamma(\Pi_{\bar{\mathcal{M}}_\gamma}(\Delta_\gamma)) + 4\mathcal{R}_\gamma(\Pi_{\bar{\mathcal{M}}_\gamma^\perp}(\theta_\gamma^*)) \right]$$

Under the assumption $\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\theta_\alpha^*) = 0$ for all $\alpha \in I$ for simplicity,

$$t^{-1/2} \mathcal{R}_\alpha(\Pi_{\bar{\mathcal{M}}_\alpha^\perp}(\Delta_\alpha)) \leq \sum_{\gamma \in I} \frac{3\lambda_\gamma}{\lambda_\alpha \sqrt{t}} \mathcal{R}_\gamma(\Pi_{\bar{\mathcal{M}}_\gamma}(\Delta_\gamma)) \leq \sum_{\gamma \in I} \frac{3\lambda_\gamma \Psi_\gamma(\bar{\mathcal{M}}_\gamma)}{\lambda_\alpha \Psi_\alpha(\bar{\mathcal{M}}_\alpha)} \|\Delta_\gamma\|_2.$$

Let Λ be $\max_{\gamma_1, \gamma_2} \left\{ 2 + \frac{3\lambda_{\gamma_1} \Psi_{\gamma_1}(\bar{\mathcal{M}}_{\gamma_1})}{\lambda_{\gamma_2} \Psi_{\gamma_2}(\bar{\mathcal{M}}_{\gamma_2})} \right\}$. Then, we obtain

$$\sum_i \left\| \sum_{j \in S(i)} c_j \right\|_2 \leq \Lambda \sum_{\gamma} \|\Delta_{\gamma}\|_2.$$

Similarly,

$$\sum_k \left\| \sum_{l \in S(k)} d_l \right\|_2 \leq \Lambda \sum_{\gamma} \|\Delta_{\gamma}\|_2.$$

Combining all the pieces together, we obtain

$$\frac{2}{n} |\langle X\Delta_{\alpha}, X\Delta_{\beta} \rangle| \leq C_{\max} \left(\Lambda \sum_{\gamma} \|\Delta_{\gamma}\|_2 \right)^2,$$

and hence,

$$\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X\Delta_{\alpha}, X\Delta_{\beta} \rangle \right| \leq C_{\max} \binom{|I|}{2} \left(\Lambda \sum_{\alpha} \|\Delta_{\alpha}\|_2 \right)^2 \leq C_{\max} \binom{|I|}{2} \Lambda^2 |I| \sum_{\alpha} \|\Delta_{\alpha}\|_2^2. \quad (29)$$

In order to complete the proof, we use the following lemma to control the remaining parameter C_{\max} used in the above inequality.

Lemma 3. *Suppose that X has independent sub-Gaussian rows. Then, for any fixed i, k and every $\delta \geq 0$, with probability at least $1 - 4 \exp(-c_2 \delta^2)$,*

$$\begin{aligned} & \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \\ & \leq \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_k} \right) + c_1 \max(\eta, \eta^2) \end{aligned} \quad (30)$$

where $\eta = \sqrt{\frac{t}{n}} + \frac{\delta}{\sqrt{n}}$, and constants c_1, c_2 only depend on the distribution of the rows in X .

See Appendix C.1 for the proof of this claim.

In order to build a bound (30) for all i, k and additionally for all $\alpha < \beta$, we use the standard union bound with the appropriate choice of δ ; let p' be the $\max\{n, p\}$. By choosing $\delta = \sqrt{c_3 \log p'}$ and the union bound, (30) holds with $\eta = \sqrt{\frac{t}{n}} + \sqrt{\frac{c_3 \log p'}{n}}$ for all i and k with probability at least $1 - 4 \exp(-c_2 c_3 \log p' + 2 \log p + \log \binom{|I|}{2})$ for some large enough constant c_3 so that $1 - 4 \exp(-c_2 c_3 \log p' + 2 \log p) \geq 1 - 4 \exp(-\log p') = 1 - \frac{4}{p'}$.

Now, for every i, k, α and β , since the set $S(i)$ is the subspace of $\bar{\mathcal{M}}_{\alpha}$ or $\bar{\mathcal{M}}_{\alpha}^{\perp}$ and similarly $S(k)$ is that of $\bar{\mathcal{M}}_{\beta}$ or $\bar{\mathcal{M}}_{\beta}^{\perp}$, along with the assumption in the statement, we have

$$\sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \leq \frac{\kappa_{\mathcal{L}}}{8 \binom{|I|}{2} \Lambda^2 |I|}.$$

At the same time, provided that n is greater than $128 \left(\frac{c_1 \binom{|I|}{2} \Lambda^2 |I|}{\kappa_{\mathcal{L}}} \right)^2 (t + c_3 \log p')$, we have

$$c_1 \max(\eta, \eta^2) \leq \frac{\kappa_{\mathcal{L}}}{8 \binom{|I|}{2} \Lambda^2 |I|}.$$

Combining two inequalities yields

$$\frac{1}{2} C_{\max} = \max_{i, k} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \leq \frac{\kappa_{\mathcal{L}}}{4 \binom{|I|}{2} \Lambda^2 |I|},$$

hence, from (29)

$$\frac{2}{n} \left| \sum_{\alpha < \beta} \langle X\Delta_{\alpha}, X\Delta_{\beta} \rangle \right| \leq \frac{\kappa_{\mathcal{L}}}{2} \sum_{\alpha} \|\Delta_{\alpha}\|_2^2.$$

as claimed.

C.1 Proof of Lemma 3

For any fixed i and k , in order to bound $\sigma_{\min}(X\mathcal{P}_{\bar{\mathcal{M}}_i})$ and $\sigma_{\max}(X\mathcal{P}_{\bar{\mathcal{M}}_i})$, we will appeal to the following result in non-asymptotic random matrix theory (see, for example, Theorem 5.39 in [13]):

Theorem 2. *Let M be a $n \times p$ matrix whose rows M_i are independent sub-Gaussian with its second moment $\Sigma \in \mathbb{R}^{p \times p}$. Then with probability at least $1 - 2 \exp(-c_2 \delta^2)$, we have,*

$$\begin{aligned}\sigma_{\min}(M) &\geq \sqrt{\sigma_{\min}(\Sigma)} \left(\sqrt{n} - c_1 \sqrt{p} - \delta \right), \\ \sigma_{\max}(M) &\leq \sqrt{\sigma_{\max}(\Sigma)} \left(\sqrt{n} + c_1 \sqrt{p} + \delta \right),\end{aligned}$$

where the constants c_1 and c_2 only depend on the distribution of the rows M_i .

Therefore, we have

$$\sigma_{\min}(X\mathcal{P}_{\bar{\mathcal{M}}_i}) \geq \sqrt{\sigma_{\min}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_i})} \left(\sqrt{n} - c_1 \sqrt{t} - \delta \right), \quad (31)$$

$$\sigma_{\max}(X\mathcal{P}_{\bar{\mathcal{M}}_i}) \leq \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_i})} \left(\sqrt{n} + c_1 \sqrt{t} + \delta \right), \quad (32)$$

where t is the number of non-zero columns of $X\mathcal{P}_{\bar{\mathcal{M}}_i}$, which is usually same as the number of atoms in the space $\bar{\mathcal{M}}_\alpha$. From (31) and (32), the difference between the mean and the sample mean can be derived by a straight forward way (we can extend Lemma 5.36 in [13]): For every $\delta \geq 0$, with probability at least $1 - 4 \exp(-c_2 \delta^2)$,

$$\begin{aligned}\sigma_{\max} \left[\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} - \mathcal{P}_{\bar{\mathcal{M}}_i} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_k} \right] \\ \leq \max(\eta, \eta^2) \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_i})} \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_k})}\end{aligned}$$

where $\eta = c_1 \sqrt{\frac{t}{n} + \frac{\delta}{\sqrt{n}}}$. Since $\sigma_{\max}(\cdot)$ is a norm, by the triangle inequality of a norm: $\|a\| - \|b\| \leq \|a - b\|$, we can finally have

$$\begin{aligned}\sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \left(\frac{1}{n} X^T X \right) \mathcal{P}_{\bar{\mathcal{M}}_k} \right) \\ \leq \sigma_{\max} \left(\mathcal{P}_{\bar{\mathcal{M}}_i} \Sigma \mathcal{P}_{\bar{\mathcal{M}}_k} \right) + \max(\eta, \eta^2) \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_i})} \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_k})}.\end{aligned}$$

D Proof of Proposition 3

The overall proof can be easily extended from the trivial connections between $\|\cdot\|_2$ and $\|\cdot\|_F$, and between $\langle \cdot, \cdot \rangle$ and $\langle\langle \cdot, \cdot \rangle\rangle$. For the matrix parameter, however, there is a caveat in proof of Lemma 3: even after the projection, the dimension of X will not change unless the left singular vectors are the standard bases. However we can still show the inequalities (31) and (32) from the following reasoning:

$$\begin{aligned}\sigma_{\max}(X\mathcal{P}_{\bar{\mathcal{M}}_i}) &= \sigma_{\max}(XUU^\top) = \sigma_{\max}(XU) \\ &\leq \sqrt{\sigma_{\max}(\Sigma U)} \left(\sqrt{n} + c_1 \sqrt{t} + \delta \right) = \sqrt{\sigma_{\max}(\Sigma \mathcal{P}_{\bar{\mathcal{M}}_i})} \left(\sqrt{n} + c_1 \sqrt{t} + \delta \right).\end{aligned}$$

E Proof of Corollary 2

Fist we restate the condition (D3) for individual clean structures, which have been analyzed previously in [11].

$$\begin{aligned}\frac{1}{n} \|X\Delta_s\|_2^2 &\geq \kappa_1 \|\Delta_s\|_2^2 - \kappa_2 \frac{\log p}{n} \|\Delta_s\|_1^2 \quad \text{for all } \Delta_s \in \mathbb{R}^p \\ \frac{1}{n} \|X\Delta_g\|_2^2 &\geq \kappa'_1 \|\Delta_g\|_2^2 - \kappa'_2 \mathbb{E} \left[\left(\max_{t=1,2,\dots,q} \frac{\|\epsilon_{G_t}\|_{a^*}}{\sqrt{n}} \right)^2 \right] \|\Delta_g\|_{1,a}^2 \quad \text{for all } \Delta_g \in \mathbb{R}^p\end{aligned}$$

with probability at least $1 - c_1 \exp(-c_2 n)$. Note that $\|\cdot\|_{a^*}$ is dual norm of $\|\cdot\|_a$ and $\epsilon \sim N(0, I_{p \times p})$ is a standard normal vector. As noted in [11], for $a = 2$ as a special case,

$$\frac{1}{n} \|X \Delta_g\|_2^2 \geq \kappa'_1 \|\Delta_g\|_2^2 - \kappa'_2 \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log q}{n}} \right)^2 \|\Delta_g\|_{1,2}^2.$$

The next step is to choose the appropriate λ_α in order utilize Theorem 1. Since the dual norm of $\|\cdot\|_1$ is the infinity norm, we need to choose λ_s for the sparsity structure such that

$$\lambda_s \geq 2\mathcal{R}_s^*(\nabla_{\theta_s} \mathcal{L}(\theta_g^* + \theta_s^*; Z_1^n)) = 4 \left\| \frac{1}{n} X^\top w \right\|_\infty.$$

Under the conditions on the column normalization of X and sub-Gaussian of w , [11] showed that $\left\| \frac{1}{n} X^\top w \right\|_\infty \leq 2\sigma \sqrt{\frac{\log p}{n}}$ with probability at least $1 - c_1 \exp(-c_2 n \lambda_s^2)$.

Moreover, [11] proved that the choice of λ_g as in the statement satisfies:

$$\lambda_g = 8\sigma \left\{ \frac{m^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log q}{n}} \right\} \geq 4 \max_{t=1,2,\dots,q} \left\| \frac{1}{n} X_{G_t}^\top w \right\|_{a^*} = 2\mathcal{R}_g^*(\nabla_{\theta_g} \mathcal{L}(\theta_g^* + \theta_s^*; Z_1^n)).$$

with probability at least $1 - 2 \exp(-2 \log q)$.

Finally, when θ_s^* is exactly s -sparse, we use the fact

$$\Psi_s(\overline{\mathcal{M}}_s(S)) = \sup_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_1}{\|\Delta\|_2} \leq \sqrt{s} \quad (33)$$

where S is equal to the support set of θ_s^* . At the same time, for a exactly group-sparse with non-zero groups S_G with cardinality s_g , for any $\Delta \in \overline{\mathcal{M}}_g(S_G)$, we have

$$\sum_{t \in \{1,2,\dots,q\}} \|\Delta_{G_t}\|_a = \sum_{t \in S_G} \|\Delta_{G_t}\|_a \leq \sum_{t \in S_G} \|\Delta_{G_t}\|_2 \leq \sqrt{s_g} \|\Delta\|_2 \quad (34)$$

where the first inequality holds for $a \geq 2$. Hence, the claim follows from Corollary 1.

F Proof of Corollary 3

For $n \geq \max(p, m)$, the RSC conditions for each structure are satisfied with $g_\alpha = 0$:

$$\frac{1}{n} \|X \Delta_\alpha\|_F = \frac{1}{n} \sum_{j=1}^p \|(X \Delta_\alpha)_j\|_2^2 \geq \frac{\sigma_{\min}(X^\top X)}{n} \|\Delta_\alpha\|_F^2.$$

Then, by Theorem 5.39 in [13], we can easily verify that

$$\frac{1}{\sqrt{n}} \sigma_{\min}(X) \geq \sqrt{\sigma_{\min}(\Sigma)} \left(1 - c \sqrt{\frac{p}{n}} - c' \sqrt{\frac{1}{n} \log \frac{1}{\delta}} \right) \quad (35)$$

with probability at least $1 - 2\delta$.

At the same time, the subspace compatibility constants for each structure can be easily extended from (33) and (34):

$$\begin{aligned} \Psi_s(\overline{\mathcal{M}}_s) &= \sup_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_1}{\|\Delta\|_F} \leq \sqrt{s}, \\ \Psi_r(\overline{\mathcal{M}}_r) &= \sum_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_{r,a}}{\|\Delta\|_F} \leq \sqrt{s_r}, \\ \Psi_c(\overline{\mathcal{M}}_c) &= \sum_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_{c,a}}{\|\Delta\|_F} \leq \sqrt{s_c}. \end{aligned}$$

Hence, in order to leverage Corollary 1, it remains to show $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha}\mathcal{L}(\theta^*; Z_1^n))$ for each structure. Under the conditions of $\frac{\|X_j\|_2}{\sqrt{n}} \leq 1$, $\sigma_{\max}(X) \leq \sqrt{n}$, and the sub-Gaussian observation noise, we can simply extend the results in [11]:

$$\lambda_s = 8\sigma\sqrt{\frac{\log p + \log m}{n}} \geq 4\left\|\frac{1}{n}X^\top W\right\|_\infty = 2\mathcal{R}_s^*(\nabla_{\Theta_s}\mathcal{L}(\Theta_r^* + \Theta_c^* + \Theta_s^*; Z_1^n))$$

with probability at least $1 - c_1 \exp(-c_2(\log p + \log m))$. Moreover,

$$\lambda_r = 8\sigma\left\{\frac{m^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log p}{n}}\right\} \geq 4\max_{t=1,2,\dots,p}\left\|\frac{1}{n}X^\top W_{t,*}\right\|_{a^*} = 2\mathcal{R}_r^*(\nabla_{\Theta_r}\mathcal{L}(\Theta_r^* + \Theta_c^* + \Theta_s^*; Z_1^n)).$$

with probability at least $1 - 2\exp(-2\log p)$. Similarly,

$$\lambda_c = 8\sigma\left\{\frac{p^{1-1/a}}{\sqrt{n}} + \sqrt{\frac{\log m}{n}}\right\} \geq 4\max_{t=1,2,\dots,m}\left\|\frac{1}{n}X^\top W_{*,t}\right\|_{a^*} = 2\mathcal{R}_c^*(\nabla_{\Theta_c}\mathcal{L}(\Theta_r^* + \Theta_c^* + \Theta_s^*; Z_1^n)).$$

with probability at least $1 - 2\exp(-2\log m)$.

G Proof of Corollary 4

Since the PCA model can be understood as the special case of (7) with $X = I_{p \times p}$, the restricted strong convexities for both structures are trivially satisfied with $\kappa_{\mathcal{L}} = 1$ and $g_\alpha = 0$.

As in the previous case, the subspace compatibility can be easily derived as:

$$\begin{aligned}\Psi_\Theta(\overline{\mathcal{M}}_\Theta) &= \sup_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_1}{\|\Delta\|_F} \leq \sqrt{r}, \\ \Psi_\Gamma(\overline{\mathcal{M}}_\Gamma) &= \sup_{\Delta \in \mathcal{M} \setminus \{0\}} \frac{\|\Delta\|_1}{\|\Delta\|_F} \leq \sqrt{s}.\end{aligned}$$

Hence, it again remains to show $\lambda_\alpha \geq 2\mathcal{R}_\alpha^*(\nabla_{\theta_\alpha}\mathcal{L}(\theta^*; Z_1^n))$ for each structure; we can directly appeal to the results for clean models [1]:

$$\lambda_\Theta = 16\sqrt{\|\Sigma\|_2}\sqrt{\frac{p}{n}} \geq 4\|W\|_2 = 2\mathcal{R}_\Theta^*(\nabla_\Theta\mathcal{L}(\Theta + \Gamma; Z_1^n))$$

with probability at least $1 - 2\exp(-c_1 p)$. Also,

$$\lambda_\Gamma = 32\rho(\Sigma)\sqrt{\frac{\log p}{n}} \geq 4\|W\|_\infty = 2\mathcal{R}_\Gamma^*(\nabla_\Gamma\mathcal{L}(\Theta + \Gamma; Z_1^n))$$

with probability at least $1 - 2\exp(-c_2 \log p)$, which completes the proof.