

On Low-rank Trace Regression under General Sampling Distribution

Nima Hamdi

*Department of Statistics
Stanford University
Stanford, CA 94305, USA*

HAMIDI@STANFORD.EDU

Mohsen Bayati

*Graduate School of Business
Stanford University
Stanford, CA 94305, USA*

BAYATI@STANFORD.EDU

Editor: Inderjit Dhillon

Abstract

In this paper, we study the trace regression when a matrix of parameters \mathbf{B}^* is estimated via the *convex relaxation of a rank-regularized regression* or via *regularized non-convex optimization*. It is known that these estimators satisfy near-optimal error bounds under assumptions on the rank, coherence, and spikiness of \mathbf{B}^* . We start by introducing a general notion of spikiness for \mathbf{B}^* that provides a generic recipe to prove the *restricted strong convexity* of the sampling operator of the trace regression and obtain near-optimal and non-asymptotic error bounds for the estimation error. Similar to the existing literature, these results require the regularization parameter to be above a certain *theory-inspired* threshold that depends on observation noise that may be unknown in practice. Next, we extend the error bounds to cases where the regularization parameter is chosen via cross-validation. This result is significant in that existing theoretical results on cross-validated estimators (Kale et al., 2011; Kumar et al., 2013; Abou-Moustafa and Szepesvari, 2017) do not apply to our setting since the estimators we study are not known to satisfy their required notion of stability. Finally, using simulations on synthetic and real data, we show that the cross-validated estimator selects a near-optimal penalty parameter and outperforms the theory-inspired approach of selecting the parameter.

Keywords: Matrix Completion, Multi-task Learning, Compressed Sensing, Low-rank Matrices, Cross-validation.

1. Introduction

We consider the problem of estimating an unknown parameter matrix $\mathbf{B}^* \in \mathbb{R}^{d_r \times d_c}$ from n noisy observations

$$Y_i = \text{Tr}(\mathbf{B}^* \mathbf{X}_i^\top) + \varepsilon_i, \quad (1.1)$$

for $i = 1, \dots, n$ where each $\varepsilon_i \in \mathbb{R}$ is a zero-mean noise and each $\mathbf{X}_i \in \mathbb{R}^{d_r \times d_c}$ is a known measurement matrix, sampled independently from a distribution Π over $\mathbb{R}^{d_r \times d_c}$. We also assume the estimation problem is high-dimensional (when $n \ll d_r \times d_c$).

Over the past decade, this problem has been studied for several families of distributions Π that span a range of applications. It is constructive to look at the following four special cases of the problem:

- **Matrix completion:** Let Π be the uniform distribution over the set of canonical basis matrices for $\mathbb{R}^{d_r \times d_c}$, i.e., over the set of all matrices that have only a single non-zero entry that is equal to 1. In this case we recover the well-known *matrix completion* problem, which estimates \mathbf{B}^* when n noisy observations of (uniformly randomly) selected entries are available (Candes and Recht, 2009; Candes and Plan, 2010; Keshavan et al., 2009, 2010). A more general version of this problem is when Π is a non-uniform probability distribution over the basis matrices (Srebro and Salakhutdinov, 2010; Negahban and Wainwright, 2011; Klopp, 2014).
- **Multi-task learning:** When the support of Π includes only matrices that have a single non-zero row, then the problem reduces to the multi-task learning problem. Specifically, one can consider a setting with n observations of multiple learning tasks, d_r different supervised learning problems, that are represented by d_r linear regression models with unknown d_c -dimensional parameters $B_1^*, \dots, B_{d_r}^*$ that form rows of \mathbf{B}^* . Equivalently, when the i_r -th row of \mathbf{X}_i is non-zero, we can assume that Y_i is a noisy observation of the i_r -th task, with a feature vector equal to the i_r -th row of \mathbf{X}_i . In multi-task learning the goal is to learn the parameters (matrix \mathbf{B}^*) by leveraging the structural properties (similarities) of the tasks (Caruana, 1997).
- **Compressed sensing via Gaussian ensembles:** If we view the matrix as a high-dimensional vector of size $d_r d_c$, then the estimation problem can be viewed as an example of the compressed sensing problem, given certain structural assumptions on \mathbf{B}^* . In this literature it is known that Gaussian ensembles, when each \mathbf{X}_i is a random matrix with entries filled with i.i.d. samples from $\mathcal{N}(0, 1)$, are a suitable family of measurement matrices (Candes and Plan, 2011).
- **Compressed sensing via factored measurements:** Consider the previous example. One drawback of the Gaussian ensembles is the need to store n large matrices which requires memory of size $O(nd_r d_c)$. Recht et al. (2010) propose factored measurements to reduce this memory requirement. They suggest using rank-1 matrices \mathbf{X}_i of the form UV^\top , where $U \in \mathbb{R}^{d_r}$ and $V \in \mathbb{R}^{d_c}$ are random vectors that reduces the memory requirement to $O(nd_r + nd_c)$.

A popular estimator, using the observations in Eq. (1.1), is given by solution to the following convex program,

$$\min_{\mathbf{B} \in \mathcal{S}} \frac{1}{n} \sum_{i=1}^n \left[Y_i - \text{Tr}(\mathbf{B}\mathbf{X}_i^\top) \right]^2 + \lambda \|\mathbf{B}\|_*, \quad (1.2)$$

where $\mathcal{S} \subseteq \mathbb{R}^{d_r \times d_c}$ is an arbitrary convex set of matrices with $\mathbf{B}^* \in \mathcal{S}$, λ is a regularization parameter, and $\|\mathbf{B}\|_*$ is the trace norm of \mathbf{B} (defined in Section 2) that favors low-rank matrices. This type of estimator was initially introduced by Candes and Recht (2009) for the noise-free version of the matrix completion problem and was later expanded to more general

cases. An admittedly incomplete list of follow-up work is (Candès and Tao, 2010; Mazumder et al., 2010; Gross, 2011; Recht, 2011; Rohde and Tsybakov, 2011; Koltchinskii et al., 2011a; Negahban and Wainwright, 2011, 2012; Klopp, 2014). Another class of estimators, studied by Srebro et al. (2005), Keshavan et al. (2009), and Keshavan et al. (2010), replaces the variable \mathbf{B} in Eq. (1.2) by \mathbf{UV}^\top , where \mathbf{U} and \mathbf{V} are explicitly low-rank matrices, and the trace-norm penalty is replaced with a ridge-type penalty term on entries of \mathbf{U} and \mathbf{V} , as shown in Eq. (3.3) of Section 3.2. These two bodies of literature provide tail bounds for the corresponding estimators, under certain assumptions on the rank, coherence, and spikiness of \mathbf{B}^* for a few classes of sampling distributions Π . We refer the reader to the detailed discussion of this literature in (Davenport and Romberg, 2016; Hastie et al., 2015), and the references therein.

Contributions. Our paper extends the above literature, and makes the following contributions:

- (i) We introduce a general notion of spikiness for \mathbf{B}^* , and construct error bounds (building on the analysis of Klopp (2014)) for the estimation error of a large family of estimators. Our main contribution is a general recipe for proving the well-known restricted strong convexity (RSC) condition, defined in Section 3.
- (ii) Next, we prove the first error bound (to the best of our knowledge) for the cross-validated version of the above family of estimators. Specifically, all bounds in for matrix estimation in the literature, as well as our bounds in (i), require the regularization parameter λ to be larger than a constant multiple of the operator norm of $\sum_{i=1}^n \varepsilon_i \mathbf{X}_i$, which is not feasible in practice due to lack of access to $\{\varepsilon_i\}_{i=1}^n$. In fact, instead of using these theory-inspired estimators, practitioners select λ via cross-validation. We prove that this cross-validated estimator satisfies similar error bounds to the ones in (i). We also show, via simulations on real and synthetic data, that the cross-validated estimator outperforms the theory-inspired estimators, and is nearly as good as the oracle estimator that chooses λ based on its knowledge of \mathbf{B}^* .

We note that existing analyses of cross-validated estimators (e.g., Kale et al., 2011; Kumar et al., 2013; Abou-Moustafa and Szepesvari, 2017) do not apply to our setting since they impose certain stability criteria on the estimation algorithm. However, establishing these criteria in our case is highly non-trivial for two reasons. First, we are studying a family of algorithms and not a specific algorithm. Second, it is not known whether stability holds even for a single low-rank matrix estimation method (including both convex and non-convex optimization).¹

- (iii) We apply our results from (i) to the four classes of problems discussed above and obtain the first such error bounds (to the best of our knowledge) for the multi-task learning problem. For matrix completion and both cases of compressed sensing (with Gaussian

1. In the convex relaxation case, it might be possible to analyze the cross-validated estimator by extending the analysis of Chetverikov et al. (2016) for the LASSO estimator. But even if such an extension is possible, it would be only for a single estimator, and not the larger family of estimators we study here. In addition, it would be a long proof (the proof for the LASSO estimator is over 30 pages), whereas our proof is only a few pages.

ensembles and with factored measurements), we obtain matching and near-optimal error bounds as the ones in the existing literature, e.g., by Negahban and Wainwright (2012), Klopp (2014), Candes and Plan (2011), and Cai and Zhang (2015).

We note that Rohde and Tsybakov (2011) and Negahban and Wainwright (2011) also consider the trace regression problem under general sampling distributions. Moreover, they only provide error bounds for the estimation error when the corresponding sampling operator satisfies the restricted isometry property (RIP) or RSC. However, they do not prove whether RIP or RSC hold for the multi-task learning problem. In fact, Rohde and Tsybakov (2011) state that their analysis cannot prove RIP for the multi-task learning problem. By contrast, we prove that RSC holds for all four classes of problems, by leveraging our unifying method of proving the RSC condition.

Koltchinskii et al. (2011a) study the trace regression problem in its general form as well. However, they study a different estimator that is not practical to implement. Specifically, Koltchinskii et al. (2011a) propose the following estimator:

$$\widehat{\mathbf{B}}_{\text{KLT}} := \arg \min_{\mathbf{B}} \left\{ \mathbb{E}[\langle \mathbf{B}, \mathbf{X} \rangle^2] - \langle \frac{2}{n} \sum_{i=1}^n Y_i \mathbf{X}_i, \mathbf{B} \rangle + \lambda \|\mathbf{B}\|_* \right\}.$$

While Koltchinskii et al. (2011a) show that $\widehat{\mathbf{B}}_{\text{KLT}}$ enjoys a general sharp oracle inequality, they need knowledge of distribution of \mathbf{X} to calculate $\mathbb{E}[\langle \mathbf{B}, \mathbf{X} \rangle^2]$. But, in our study (like majority of the literature on low-rank matrix estimation), the estimators do not need any knowledge of the distribution of \mathbf{X} and work with the data at hand, which is more practical.

Finally, for Gaussian ensembles and factored measurements, when there is no noise, our results also demonstrate that \mathbf{B}^* can be exactly recovered when the number of observations is above a certain threshold. Our recovery thresholds match the ones by Candes and Recht (2009) and Cai and Zhang (2015).

Organization of the paper. We introduce additional notation and state the precise formulation of the problem in Section 2. In Section 3 we introduce a family of estimators and prove tail bounds on their estimation error. Section 4 contains our results for the cross-validated estimator and corresponding numerical simulations. Application of our main error bounds to the aforementioned four classes of problems is given in Section 5, and exact recovery results are given in Section 6. Details of the proofs are discussed in Appendices A and B and additional simulations are provided in Appendix C.

2. Notation and Problem Formulation

Notation in bold roman capital letters (e.g., \mathbf{A}) denote matrices and non-bold italic capital letters denote vectors (e.g., V). For any positive integer m , the notation e_1, e_2, \dots, e_m refers to the standard basis for \mathbb{R}^m , and \mathbf{I}_m is the $m \times m$ identity matrix. The *trace inner product* of matrices \mathbf{A}_1 and \mathbf{A}_2 with the same dimensions is defined as

$$\langle \mathbf{A}_1, \mathbf{A}_2 \rangle := \text{Tr}(\mathbf{A}_1 \mathbf{A}_2^\top).$$

For $d_r \times d_c$ matrices $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$, let the sampling operator $\mathfrak{X} : \mathbb{R}^{d_r \times d_c} \rightarrow \mathbb{R}^n$ be given by

$$[\mathfrak{X}(\mathbf{B})]_i := \langle \mathbf{B}, \mathbf{X}_i \rangle \quad \text{for all } i \in [n],$$

where by $[k]$, we denote the set $\{1, 2, \dots, k\}$. For any two real numbers a and b , $a \vee b$ and $a \wedge b$ denote $\max(a, b)$ and $\min(a, b)$ respectively. Also, a real-valued random variable z is σ -sub-Gaussian if $\mathbb{E}[\exp(\eta z)] \leq \exp(\sigma^2 \eta^2 / 2)$ for all $\eta \in \mathbb{R}$.

For a norm² $\mathfrak{N} : \mathcal{X} \rightarrow \mathbb{R}^+ \cup \{0\}$ defined on a vector space \mathcal{X} , let $\mathfrak{N}^* : \mathcal{X} \rightarrow \mathbb{R}^+ \cup \{0, \infty\}$ be its *dual norm* defined as

$$\mathfrak{N}^*(X) = \sup_{\mathfrak{N}(Y) \leq 1} \langle X, Y \rangle \quad \text{for all } X \in \mathcal{X}.$$

In this paper, we use several different matrix norms. A brief explanation of these norms is as follows. Let \mathbf{B} be a matrix with d_r rows and d_c columns.

1. **L[∞]-norm** is defined by $\|\mathbf{B}\|_\infty := \max_{(i,j) \in [d_r] \times [d_c]} \{|\mathbf{B}_{ij}|\}$.
2. **Frobenius norm** is defined as $\|\mathbf{B}\|_F := \sqrt{\sum_{(i,j) \in [d_r] \times [d_c]} \mathbf{B}_{ij}^2}$.
3. **Operator norm** is defined as $\|\mathbf{B}\|_{\text{op}} := \sup_{\|V\|_2=1 \ \& \ V \in \mathbb{R}^{d_c}} \|\mathbf{B}V\|_2$. An alternative definition of the operator norm is given by using the *singular value decomposition* (SVD) of $\mathbf{B} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$, where \mathbf{D} is an $r \times r$ diagonal matrix and r denotes the rank of \mathbf{B} . In this case, it is well known that $\|\mathbf{B}\|_{\text{op}} = \mathbf{D}_{11}$.
4. **Trace norm** is defined as $\|\mathbf{B}\|_* := \sum_{i=1}^r \mathbf{D}_{ii}$.
5. **L_{p,q}-norm**, for $p, q \geq 1$, is defined as $\|\mathbf{B}\|_{p,q} := \left(\sum_{r=1}^{d_r} \left(\sum_{c=1}^{d_c} |\mathbf{B}_{rc}|^p \right)^{q/p} \right)^{1/q}$.
6. **L²(Π)-norm** is defined as $\|\mathbf{B}\|_{L^2(\Pi)} := \sqrt{\mathbb{E}[\langle \mathbf{B}, \mathbf{X} \rangle^2]}$, where \mathbf{X} is sampled from a probability measure Π on $\mathbb{R}^{d_r \times d_c}$.
7. **Exponential Orlicz norm** is defined for any $p \geq 1$ and probability measure Π on $\mathbb{R}^{d_r \times d_c}$ as

$$\|\mathbf{B}\|_{\psi_p(\Pi)} := \|\langle \mathbf{B}, \mathbf{X} \rangle\|_{\psi_p} = \inf \left\{ t > 0 : \mathbb{E} \left[e^{\left(\frac{|\langle \mathbf{B}, \mathbf{X} \rangle|}{t} \right)^p} - 1 \right] \leq 1 \right\},$$

where \mathbf{X} has distribution Π .

Now, we will state the main trace regression problem that is studied in this paper.

Problem 2.1 *Let \mathbf{B}^* be an unknown $d_r \times d_c$ matrix that is low-rank and has real-valued entries, specifically, $r \ll \min(d_r, d_c)$. Moreover, assume that Π is a distribution on $\mathbb{R}^{d_r \times d_c}$,*

2. \mathfrak{N} can also be a semi-norm.

and that $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ are i.i.d. samples from Π and their corresponding sampling operator is $\mathfrak{X} : \mathbb{R}^{d_r \times d_c} \rightarrow \mathbb{R}^n$. Our regression model is given by

$$Y = \mathfrak{X}(\mathbf{B}^*) + E, \tag{2.1}$$

where observation Y and noise E are both vectors in \mathbb{R}^n . Elements of E are denoted by $\varepsilon_1, \dots, \varepsilon_n$, where $\{\varepsilon_i\}_{i=1}^n$ is a sequence of independent mean-zero random variables with variance at most σ^2 . The goal is to estimate \mathbf{B}^* from the observations Y .

We also use the following two notations:

$$\Sigma := \frac{1}{n} \sum_{i=1}^n \varepsilon_i \mathbf{X}_i \quad \text{and} \quad \Sigma_R := \frac{1}{n} \sum_{i=1}^n \zeta_i \mathbf{X}_i,$$

where $\{\zeta_i\}_{i=1}^n$ is an i.i.d. sequence with a Rademacher distribution.

3. Estimation Method and Corresponding Tail Bounds

This section is dedicated to the tail bounds for the trace regression problem. The results and the proofs in this section are based on those found in Klopp (2014), with a slightly sharper analysis. For the sake of completeness, all proofs are reproduced (adapted) for our setting and are presented in Appendix B.

3.1 General notions of spikiness

It is a well-known fact that, in Problem 2.1, the low-rank assumption is not sufficient for estimating \mathbf{B}^* from the observations Y . For example, changing one entry of \mathbf{B}^* increases the rank of the matrix by (at most) 1, whereas it would be impossible to distinguish between these two matrices without observing the modified single entry. To remedy this difficulty, Candes and Recht (2009) and Keshavan et al. (2009) propose an incoherence assumption. If singular value decomposition (SVD) of \mathbf{B}^* is $\mathbf{U}\Sigma\mathbf{V}$, then the incoherence assumption roughly means that all rows of \mathbf{U} and \mathbf{V} have norms of the same order. Alternatively, Negahban and Wainwright (2012) study the problem under a (less restrictive) assumption that bounds the *spikiness* of the matrix \mathbf{B}^* . Here, we define a general notion of spikiness as a matrix that is sufficient for the estimation of \mathbf{B}^* under the general sampling distribution Π , and includes the spikiness studied by Negahban and Wainwright (2012) as a special case. We define the *spikiness* and *low rankness* of a matrix $\mathbf{B} \in \mathbb{R}^{d_r \times d_c}$ as

$$\text{spikiness of } \mathbf{B} := \frac{\mathfrak{N}(\mathbf{B})}{\|\mathbf{B}\|_F} \quad \text{and} \quad \text{low rankness of } \mathbf{B} := \frac{\|\mathbf{B}\|_*}{\|\mathbf{B}\|_F}.$$

The spikiness used in Negahban and Wainwright (2012) can be recovered by setting $\mathfrak{N}(\mathbf{B}) = \sqrt{d_r d_c} \|\mathbf{B}\|_\infty$. This choice of norm, however, is not suitable for many distributions of \mathbf{X} . For example, we will see in Section 5.2 in the multi-task learning setting that $\mathfrak{N}(\mathbf{B}) = 2\sqrt{d} \|\mathbf{B}\|_{2,\infty}$. In general, we will also show that in all cases, the exponential Orlicz norm can be used to guide the selection of $\mathfrak{N}(\mathbf{B})$. Specifically, $\|\cdot\|_{\psi_2(\Pi)}$ is used in Sections 5.1 to 5.3, and $\|\cdot\|_{\psi_1(\Pi)}$ is used in Section 5.4.

3.1.1 INTUITION FOR SELECTING \mathfrak{N}

Here we provide some intuition for the use of the exponential Orlicz norm to select \mathfrak{N} . Negahban and Wainwright (2011) show how an error bound can be obtained from the RSC condition (defined in Section 3.3 below) on a suitable set of matrices. The condition roughly requires that $\|\mathfrak{X}(\mathbf{B})\|_2^2/n \geq \alpha\|\mathbf{B}\|_F^2$ for a constant α . Assume that random variables $\langle \mathbf{X}_i, \mathbf{B} \rangle$ are not heavy-tailed; then $\|\mathfrak{X}(\mathbf{B})\|_2^2/n$ concentrates around its mean, $\|\mathbf{B}\|_{L^2(\Pi)}$. The Orlicz norm, which measures how heavy-tailed a distribution is, helps us construct a suitable “constraint” set of matrices where the aforementioned concentration holds simultaneously.

To be more concrete, consider the multi-tasking example (studied in Section 5.2) for the case of $d_r = d_c = d$. In particular, $\mathbf{X}_i = e_i X_i^\top$, where i is a uniformly random index in $[d]$ and X_i is a d vector with i.i.d. $\mathcal{N}(0, d)$ entries.

In this example, we have $\|\mathbf{B}\|_{L^2(\Pi)} = \|\mathbf{B}\|_F^2$ for all $\mathbf{B} \in \mathbb{R}^{d \times d}$. Now, if a fixed \mathbf{B} is such that $\langle \mathbf{X}_i, \mathbf{B} \rangle$ has a light tail, one can show that for sufficiently large n , due to concentration, $\|\mathfrak{X}(\mathbf{B})\|_2^2/n$ is at least $\|\mathbf{B}\|_F^2/2$. To see this, consider two extreme cases: let \mathbf{B}_1 be a matrix such that its first row has an ℓ^2 -norm equal to one and its remaining entries are zero, and let \mathbf{B}_2 be a matrix such that all of its rows have an ℓ_2 -norm equal to $1/\sqrt{d}$. Both of these matrices have a Frobenious norm equal to one. But, $\|\mathfrak{X}(\mathbf{B}_1)\|_2^2$ has a heavier tail than $\|\mathfrak{X}(\mathbf{B}_2)\|_2^2$, because $\langle \mathbf{X}_i, \mathbf{B}_1 \rangle$ is zero most of the time, but it is very large occasionally, whereas almost all realizations of $\langle \mathbf{X}_i, \mathbf{B}_2 \rangle$ are roughly of the same size. This intuition implies that matrices whose rows are almost the same size are more likely to satisfy RSC than the other ones. However, since X_i^\top is invariant under rotation, one can see that the only thing that matters for RSC is the *norm* of the rows. Indeed, the Orlicz norm verifies our intuition and after doing the computation (see Section 5.2 for details), one can see that $\|\mathbf{B}\|_{\psi_2(\Pi)} = O(\sqrt{d}\|\mathbf{B}\|_{2,\infty})$. We will later see in Section 5.2 that by defining $\mathfrak{N}(\mathbf{B})$ to be a constant multiple of $\sqrt{d}\|\mathbf{B}\|_{2,\infty}$, we obtain a suitable choice for $\|\mathbf{B}\|$. Note that, for the matrix completion application, one can argue similarly that, in order for a matrix to satisfy the RSC condition with high probability, it cannot have a few very large rows. However, the second part of the above argument that uses rotation invariance does not apply, and actually a similar argument to the former one implies that each row cannot also have a few very large entries. Therefore, all entries of \mathbf{B} should be roughly of the same size, which would match the spikiness notion of Negahban and Wainwright (2012).

3.2 Estimation

Before introducing the estimation approach, we state our first assumption.

Assumption 3.1 *Assume that $\mathfrak{N}(\mathbf{B}^*) \leq b^*$ for some $b^* > 0$.*

Note that in Assumption 3.1, we require only a bound on $\mathfrak{N}(\mathbf{B}^*)$ and *not* the general spikiness of \mathbf{B}^* .

Our theoretical results enjoy a certain notion of algorithmic independence. To make this point precise, we start by considering the trace-norm penalized least squares loss functions, also stated in a different format in Eq. (1.2),

$$\mathcal{L}(\mathbf{B}) := \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 + \lambda \|\mathbf{B}\|_* . \tag{3.1}$$

However, we do not necessarily need to find the global minimum of Eq. (3.1). Let $\mathcal{S} \subseteq \mathbb{R}^{d_r \times d_c}$ be an arbitrary convex set of matrices with $\mathbf{B}^* \in \mathcal{S}$. All of our bounds are stated for any $\widehat{\mathbf{B}}$ that satisfies

$$\widehat{\mathbf{B}} \in \mathcal{S} \quad \text{and} \quad \mathfrak{N}(\widehat{\mathbf{B}}) \leq b^* \quad \text{and} \quad \mathcal{L}(\widehat{\mathbf{B}}) \leq \mathcal{L}(\mathbf{B}^*). \quad (3.2)$$

While the global minimizer, $\arg \min_{\mathbf{B} \in \mathcal{S}: \mathfrak{N}(\mathbf{B}) \leq b^*} \mathcal{L}(\mathbf{B})$, satisfies Eq. (3.2), we can also achieve Eq. (3.2) by using other loss minimization problems. A notable example is the *alternating minimization* approach that aims to solve

$$(\widehat{\mathbf{U}}, \widehat{\mathbf{V}}) = \arg \min_{\mathbf{U} \in \mathbb{R}^{d_r \times r}, \mathbf{V} \in \mathbb{R}^{d_c \times r}, \mathbf{UV}^\top \in \mathcal{S}} \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{UV}^\top)\|_2^2 + \frac{\lambda}{2} (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2), \quad (3.3)$$

where r is a preselected value for the rank. If we find the minimizer of Eq. (3.3), then it is known that $\widehat{\mathbf{B}} = \widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top$ satisfies Eq. (3.2). For example, this fact is shown by Keshavan et al. (2009) and Mazumder et al. (2010).

3.3 Restricted strong convexity and the tail bounds

The upper bound for the estimation error that will be stated next relies on a *restricted strong convexity* (RSC) condition. This condition is defined below and will be shown to hold with high probability, under some assumptions on \mathbf{B}^* and Π .

Definition 3.1 (Restricted Strong Convexity Condition) *For a constraint set $\mathcal{C} \subseteq \mathbb{R}^{d_r \times d_c}$, we say that $\mathfrak{X}(\cdot)$ satisfies the RSC condition over the set \mathcal{C} if there exist constants $\alpha(\mathfrak{X}) > 0$ and $\beta(\mathfrak{X})$ such that, for all $\mathbf{B} \in \mathcal{C}$,*

$$\frac{\|\mathfrak{X}(\mathbf{B})\|_2^2}{n} \geq \alpha(\mathfrak{X}) \|\mathbf{B}\|_F^2 - \beta(\mathfrak{X}).$$

For the upper bound, we need the RSC condition to hold for a specific family of constraint sets that are parameterized by two positive parameters ν, η . Define $\mathcal{C}(\nu, \eta)$ as

$$\mathcal{C}(\nu, \eta) := \left\{ \mathbf{B} \in \mathbb{R}^{d_r \times d_c} \mid \mathfrak{N}(\mathbf{B}) = 1, \|\mathbf{B}\|_F^2 \geq \nu, \|\mathbf{B}\|_* \leq \sqrt{\eta} \|\mathbf{B}\|_F \right\}. \quad (3.4)$$

The next result provides the upper bound on the estimation error when λ is large enough and the RSC condition holds on $\mathcal{C}(\nu, \eta)$ for some constants α and β .

Theorem 3.1 *Let \mathbf{B}^* be a matrix of rank r and define $\eta := 72r$. Also assume that $\mathfrak{X}(\cdot)$ satisfies the RSC condition for $\mathcal{C}(\nu, \eta)$ as defined in Definition 3.1 with constants $\alpha = \alpha(\mathfrak{X})$ and $\beta = \beta(\mathfrak{X})$. In addition, assume that $\lambda = \lambda(n)$ is chosen such that*

$$\lambda \geq 3 \|\Sigma\|_{\text{op}}, \quad (3.5)$$

where $\Sigma = \frac{1}{n} \sum_{i=1}^n \varepsilon_i \mathbf{X}_i$. Then, for any matrix $\widehat{\mathbf{B}}$ satisfying Eq. (3.2), we have

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 \leq \left(\frac{100\lambda^2 r}{3\alpha^2} + \frac{8b^{*2}\beta}{\alpha} \right) \vee 4b^{*2}\nu. \quad (3.6)$$

To simplify the notation, we will write λ instead of $\lambda(n)$. The proof of Theorem 3.1 is given in Appendix B.1. However, here we provide a summary of the main steps to help in understanding the source of each term on the right-hand side of Eq. (3.6). First, using Eq. (3.2) and Eq. (3.5) and a series of algebraic manipulations, one can show that $\|\mathfrak{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_2^2/n \leq \lambda\sqrt{r}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F$ holds, up to constants. Combining this with the RSC condition for $(\widehat{\mathbf{B}} - \mathbf{B}^*)/\mathfrak{N}(\widehat{\mathbf{B}} - \mathbf{B}^*)$ yields the inequality $\alpha\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 - 4\beta b^{*2} \leq \lambda\sqrt{r}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F$, which can be translated, up to constants, into the term

$$\frac{\lambda^2 r}{\alpha^2} + \frac{b^{*2} \beta}{\alpha}$$

in Eq. (3.6). However, to legitimize application of the RSC condition, one needs a lower bound $\sqrt{\nu}$ on the Frobenius norm of $(\widehat{\mathbf{B}} - \mathbf{B}^*)/\mathfrak{N}(\widehat{\mathbf{B}} - \mathbf{B}^*)$ to help show that $(\widehat{\mathbf{B}} - \mathbf{B}^*)/\mathfrak{N}(\widehat{\mathbf{B}} - \mathbf{B}^*)$ is in $\mathcal{C}(\nu, \eta)$, where η is a constant multiple of r . In the absence of this lower bound, we can directly use $b^{*2}\nu$ as an upper bound for $\alpha\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2$, which explains the last term on the right-hand side of Eq. (3.6).

Note that even though the assumptions of Theorem 3.1 involve the noise and the observation matrices \mathbf{X}_i , no distributional assumption is required and the result is deterministic. However, we will employ probabilistic results later to show that the assumptions of the theorem hold. Specifically, we can show that Eq. (3.5) holds with high probability, using a version of the Bernstein tail inequality for the operator norm of matrix martingales. This is stated as Proposition B.1 in Appendix B.2 and also appears as Proposition 11 of Klopp (2014).

The other condition in Theorem 3.1, RSC for $\mathcal{C}(\nu, \eta)$, will be shown to hold with high probability by Theorem 3.2 below. Before stating this result, we need two distributional assumptions for $\mathfrak{X}(\cdot)$. Recall that the distribution (over $\mathbb{R}^{d_r \times d_c}$) from which our observation matrices $\{\mathbf{X}_i\}_{i=1}^n$ are sampled is denoted by Π .

Assumption 3.2 *For positive constants γ_{\min} and γ_{\max} , and for all $\mathbf{B} \in \mathbb{R}^{d_r \times d_c}$, the following holds:*

$$\gamma_{\min} \|\mathbf{B}\|_F^2 \leq \|\mathbf{B}\|_{L^2(\Pi)}^2 \leq \gamma_{\max} \|\mathbf{B}\|_F^2.$$

As we will see later in this section, all upper bounds include γ_{\min} in the denominator. Therefore, Assumption 3.2 provides a sufficient condition on the sampling distribution, which will make the estimation problem tractable. For example, in the special case of matrix completion, as demonstrated empirically by Srebro and Salakhutdinov (2010) and Foygel et al. (2011), when some rows or columns of \mathbf{B}^* are sampled disproportionately more frequently, the estimation becomes more difficult. The assumption also resembles the bounds on leverage scores as shown by Chen et al. (2015). The condition involving γ_{\max} will be needed in Section 4 for the analysis of the cross-validated estimator.

Assumption 3.3 *There exists $\mathfrak{c} > 0$, such that*

$$\mathbb{E}[\langle \mathbf{X}, \mathbf{B} \rangle^2 \cdot \mathbb{I}(|\langle \mathbf{X}, \mathbf{B} \rangle| \leq \mathfrak{c})] \geq \frac{1}{2} \mathbb{E}[\langle \mathbf{X}, \mathbf{B} \rangle^2],$$

for all \mathbf{B} with $\mathfrak{N}(\mathbf{B}) \leq 1$, where the expectations are with respect to Π .

Assumption 3.3 imposes a restriction on the sampling distribution such that the random variable $\langle \mathbf{X}, \mathbf{B} \rangle$ will not have a heavy tail. This allows us to show that $\langle \mathbf{X}, \mathbf{B} \rangle$ is not much smaller than its expectation, which in turn allows us to prove the RSC condition. In the special case of matrix completion, this assumption plays a similar role to the “leveraged sampling” assumption in Chen et al. (2015). The next remark shows that the variance of $\langle \mathbf{X}, \mathbf{B} \rangle$ plays an important role in proving Assumption 3.3.

Remark 3.1 *We will show later (in Corollary A.1 of Appendix A) that whenever $\text{Var}(\langle \mathbf{X}, \mathbf{B} \rangle) \approx 1$ uniformly over \mathbf{B} with $\mathfrak{N}(\mathbf{B}) = 1$, when $\mathfrak{N}(\cdot) = \|\cdot\|_{\psi_p(\Pi)}$ and $p \geq 1$, then \mathfrak{c} is a small constant that does not depend on the dimensions.*

The next result shows that a slightly more general form of the RSC condition holds with high probability.

Theorem 3.2 (Restricted Strong Convexity) *Define*

$$\mathcal{C}'(\theta, \eta) := \left\{ \mathbf{A} \in \mathbb{R}^{d_r \times d_c} \mid \mathfrak{N}(\mathbf{A}) = 1, \|\mathbf{A}\|_{L^2(\Pi)}^2 \geq \theta, \|\mathbf{A}\|_* \leq \sqrt{\eta} \|\mathbf{A}\|_F \right\}.$$

If Assumptions 3.2 and 3.3 hold, then for all $\mathbf{A} \in \mathcal{C}'(\theta, \eta)$, the inequality

$$\frac{1}{n} \|\mathfrak{X}(\mathbf{A})\|_2^2 \geq \frac{1}{4} \|\mathbf{A}\|_{L^2(\Pi)}^2 - \frac{93\eta\mathfrak{c}^2}{\gamma_{\min}} \mathbb{E}[\|\Sigma_R\|_{\text{op}}]^2 \quad (3.7)$$

holds with probability greater than $1 - 2 \exp\left(-\frac{Cn\theta}{\mathfrak{c}^2}\right)$ where $C > 0$ is an absolute constant, provided that $Cn\theta > \mathfrak{c}^2$, and the equality $\Sigma_R := \frac{1}{n} \sum_{i=1}^n \zeta_i \mathbf{X}_i$ holds where $\{\zeta_i\}_{i=1}^n$ is an i.i.d. sequence with a Rademacher distribution.

Remark 3.2 *Our proof of Theorem 3.2, stated in Appendix B.3, follows a similar strategy as in (Klopp, 2014). However, in our general setting the random variable $\langle \mathbf{X}, \mathbf{A} \rangle$, where $\mathbf{A} \in \mathcal{C}(\nu, \eta)$, can be large by a factor of \sqrt{d} , e.g., in the multi-task learning case. In this situation, a direct application of the approach in (Klopp, 2014) leads to an extra factor \sqrt{d} in $\beta(\mathfrak{X})$ that also appears in the final upper bound of Theorem 3.1. In order to avoid this extra factor, we sharpen the application of the Massart concentration inequality, as shown in the proof of Lemma B.2.*

Note that Theorem 3.2 states RSC holds for $\mathcal{C}'(\theta, \eta)$ which is slightly different than the set $\mathcal{C}(\nu, \eta)$ defined in Eq. (3.4). But, using Assumption 3.2, we can see that

$$\mathcal{C}(\nu, \eta) \subseteq \mathcal{C}'(\gamma_{\min} \nu, \eta).$$

Therefore, the following variant of the RSC condition holds.

Corollary 3.1 *If Assumptions 3.2 and 3.3 hold, with probability greater than $1 - 2 \exp\left(-Cn\gamma_{\min} \nu / \mathfrak{c}^2\right)$, the inequality*

$$\frac{1}{n} \|\mathfrak{X}(\mathbf{A})\|_2^2 \geq \frac{\gamma_{\min}}{4} \|\mathbf{A}\|_F^2 - \frac{93\eta\mathfrak{c}^2}{\gamma_{\min}} \mathbb{E}[\|\Sigma_R\|_{\text{op}}]^2 \quad (3.8)$$

holds for all $\mathbf{A} \in \mathcal{C}(\nu, \eta)$, where $C > 0$ is an absolute constant satisfying $Cn\gamma_{\min} \nu > \mathfrak{c}^2$, and $\Sigma_R := \frac{1}{n} \sum_{i=1}^n \zeta_i \mathbf{X}_i$ with $\{\zeta_i\}_{i=1}^n$ is an i.i.d. sequence with Rademacher distribution.

We conclude this section by stating the following corollary. This corollary combines the RSC condition (the version in Corollary 3.1) and the general deterministic error bound (Theorem 3.1) to obtain the following probabilistic error bound.

Corollary 3.2 *If Assumptions 3.1 to 3.3 hold, and let $\lambda = \lambda(n)$ be larger than $C \mathbb{E}[\|\Sigma_R\|_{\text{op}}] b^* \mathbf{c}$ where C is an arbitrary (and positive) constant, then we have that*

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 \leq \frac{C' \lambda^2 r}{\gamma_{\min}^2}$$

holds with probability at least

$$1 - \mathbb{P}(\lambda < 3\|\Sigma\|_{\text{op}}) - 2 \exp\left(-\frac{C'' n \lambda^2 r}{\mathbf{c}^2 b^{*2} \gamma_{\min}}\right) \quad (3.9)$$

for numerical constants $C', C'' > 0$.

Proof. First, we denote the threshold for λ by $\lambda_1 = \lambda_1(n) = C \mathbb{E}[\|\Sigma_R\|_{\text{op}}] b^* \mathbf{c}$. Now, by defining

$$\alpha := \frac{\gamma_{\min}}{4}, \quad \beta := \frac{6696 \cdot r \mathbf{c}^2}{\gamma_{\min}} \mathbb{E}[\|\Sigma_R\|_{\text{op}}]^2, \quad \text{and} \quad \nu := \frac{\lambda_1^2 r}{\gamma_{\min}^2 b^{*2}},$$

we observe that

$$\begin{aligned} \left(\frac{100\lambda^2 r}{3\alpha^2} + \frac{8b^{*2}\beta}{\alpha}\right) \vee 4b^{*2}\nu &= \left(\frac{1600\lambda^2 r}{3\gamma_{\min}^2} + \frac{32 \times 6696\lambda_1^2 r}{C^2\gamma_{\min}^2}\right) \vee \frac{4\lambda_1^2 r}{\gamma_{\min}^2} \\ &\leq \frac{C' \lambda^2 r}{\gamma_{\min}^2}, \end{aligned}$$

for a sufficiently large constant $C' > 0$ (in fact, we would need $C' \geq 534 + 2.15 \times 10^5 \times C^{-2}$). The rest follows immediately from Theorem 3.1 and Corollary 3.1. We note that the condition of Corollary 3.1 can be shown to hold by taking C'' such that $C'' n \lambda^2 r > \mathbf{c}^2 b^{*2} \gamma_{\min}$. \square

Remark 3.3 *While Corollary 3.2 relies on two conditions for λ , namely, that $\lambda \geq \lambda_1$ and that $\mathbb{P}(\lambda < 3\|\Sigma\|_{\text{op}})$ is small, only the latter condition is needed to obtain a tail bound like Eq. (3.6) in Theorem 3.1. The additional condition $\lambda \geq \lambda_1$ only helps to make the upper bound simpler, i.e., by allowing us to obtain $C' \lambda^2 r / \gamma_{\min}^2$ instead of the right-hand side of Eq. (3.6).*

Remark 3.4 (Sampling complexity) *The final upper bound obtained by Corollary 3.2 depends on λ , which is required to be larger than $C \mathbb{E}[\|\Sigma_R\|_{\text{op}}] b^* \mathbf{c}$. We will show in Section 5 that $\mathbb{E}[\|\Sigma_R\|_{\text{op}}]$ is generally of order $\sqrt{d \log d/n}$, which yields an upper bound of order $rd \log/n$. We will also show that the probability of the upper bound, i.e., Eq. (3.9), is close to 1 when $n \geq d \log^3 d$ (Sections 5.1 to 5.3) or $n \geq d \log^4 d$ (Section 5.4), up to constants. Similarly, in the case of exact recovery, in Section 6, we will show that the sampling complexity is determined by a certain inequality that depends on $\mathbb{E}[\|\Sigma_R\|_{\text{op}}]$, Eq. (6.3), and is satisfied when (up to constants) $n \geq rd$ (for compressed sensing with Gaussian ensembles) and when $n \geq rd \log(d)$ (for compressed sensing with factored measurements).*

Remark 3.5 (Lower bound and optimality) *While Corollary 3.2 provides a general upper bound that applies to several families of problems, as we will showcase in Section 5, one may wonder whether the provided upper bound is tight. We will show in Section 5 that the upper bound indeed matches provable lower bounds (known in the literature). Specifically, we will show that the upper bound of Corollary 3.2 is the same as that of Corollary 1 of Negahban and Wainwright (2012), for the matrix completion problem, and the same as the upper bound of Theorem 2.4 of Candes and Plan (2011), for the compressed sensing problem. In both of these papers, the bounds are shown to be optimal, which means that the upper bound of Corollary 3.2 is also optimal for both classes of problems. But, it is an interesting open question whether one can prove a general lower bound, in the same spirit as the general upper bound of Corollary 3.2.*

4. Tail Bound for the Cross-validated Estimator

One of the assumptions required for the tail bounds in Section 3 for $\widehat{\mathbf{B}} = \widehat{\mathbf{B}}(\lambda)$ is that λ should be larger than $3\|\boldsymbol{\Sigma}\|_{\text{op}}$. However, in practice we do not have access to the latter parameter, which relies on knowledge of the noise values $\{\varepsilon_i\}_{i=1}^n$. Therefore, practitioners often use cross-validation to tune parameter λ . In this section, we prove that if λ is selected via cross-validation, $\widehat{\mathbf{B}}(\lambda)$ has similar tail bounds to those in Section 3. This would provide theoretical support for the selection of λ via cross-validation for our family of estimators.

Let $\{(\mathbf{X}_i, y_i)\}_{i=1}^n$ be a set of observations and denote by K the number of cross-validation folds. Let $\{\mathcal{I}_k\}_{k \in [K]}$ be a set of disjoint subsets of $[n]$ where $\cup_{k \in [K]} \mathcal{I}_k = [n]$. Also, we define $\mathcal{I}_{-k} := [n] \setminus \mathcal{I}_k$. Letting $n_k := |\mathcal{I}_k|$, we have

$$n = n_1 + \dots + n_K.$$

Let $\mathfrak{X}_k(\cdot)$ and $\mathfrak{X}_{-k}(\cdot)$ be sampling operators for $\{\mathbf{X}_i\}_{i \in \mathcal{I}_k}$ and $\{\mathbf{X}_i\}_{i \in \mathcal{I}_{-k}}$, respectively. Similarly, E_k and E_{-k} are noise vectors on the corresponding sets of indices. Finally, Y_k and Y_{-k} denote the response vectors corresponding to $\mathfrak{X}_k(\mathbf{B}^*) + E_k$ and $\mathfrak{X}_{-k}(\mathbf{B}^*) + E_{-k}$, respectively. In our analysis, we assume that each partition contains a large fraction of all samples; namely, we assume that $n_k \geq N/(2K)$ for all $k \in [K]$.

Also, throughout this section, we assume that for any $\lambda > 0$, the estimators $\widehat{\mathbf{B}}_{-k}(\lambda)$ satisfy Eq. (3.2) for observations $(\mathfrak{X}_{-k}(\mathbf{B}^*), Y_{-k})$ for each $k \in [K]$. Note that this means that in $\mathcal{L}(\widehat{\mathbf{B}}_{-k}(\lambda))$, by Eq. (3.1), the proper scaling $1/(n - n_k)$ is used instead of $1/n$.

Define

$$\widehat{E}(\lambda) := \sum_{k=1}^K \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda))\|_2^2.$$

For any fixed λ , it can be observed that $\|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda))\|_2^2$ is an unbiased estimate of the prediction error for $\widehat{\mathbf{B}}_{-k}(\lambda)$. For every λ we also define the estimator $\widehat{\mathbf{B}}_{\text{cv}}(\lambda)$ as follows:

$$\widehat{\mathbf{B}}_{\text{cv}}(\lambda) := \sum_{k=1}^K \frac{n_k}{n} \cdot \widehat{\mathbf{B}}_{-k}(\lambda). \quad (4.1)$$

Cross-validation works by starting with a set $\Lambda = \{\lambda_1, \lambda_2, \dots, \lambda_L\}$ of potential (positive) regularization parameters and then choosing

$$\widehat{\lambda}_{\text{cv}} \in \arg \min_{\lambda \in \Lambda} \widehat{E}(\lambda).$$

Then, the K -fold cross-validated estimator with respect to Λ is $\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}})$.

Next, we state two main results. First, in Theorem 4.1, we show a bound for $\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda})$, where $\widehat{\lambda}$ can be any value in Λ . Then, in Theorem 4.2, we combine Theorem 4.1 with Corollary 3.2 to obtain the main result of this section, which is an explicit tail bound for $\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}})$.

Theorem 4.1 *Let $\Lambda = \{\lambda_1, \lambda_2, \dots, \lambda_L\}$ be a set of positive regularization parameters, $\widehat{\mathbf{B}}_{\text{cv}}$ be defined as in Eq. (4.1), and $\widehat{\lambda}$ be a random variable such that $\widehat{\lambda} \in \Lambda$ almost surely. Moreover, define*

$$\bar{\sigma}^2 = \frac{1}{n} \sum_{i \in [n]} \text{Var}(\varepsilon_i),$$

and assume that $(\varepsilon_i)_{i=1}^n$ are independent zero-mean σ^2 -sub-Gaussian random variables. Then, for any $t > 0$, we have

$$\|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 \leq \widehat{E}(\widehat{\lambda}) - \bar{\sigma}^2 + t, \quad (4.2)$$

with probability at least

$$1 - 6KL \exp \left[-C \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) \cdot \frac{n}{K} \right], \quad (4.3)$$

where $C > 0$ is a numerical constant. More generally, for all $\ell \in [L]$ and $k \in [K]$, in addition to Eq. (4.2), the following inequality is satisfied with the same probability as in Eq. (4.3),

$$\left| \frac{1}{n_k} \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda_\ell))\|_2^2 - \left[\|\widehat{\mathbf{B}}_{-k}(\lambda_\ell) - \mathbf{B}^*\|_{L^2(\Pi)}^2 + \bar{\sigma}_k^2 \right] \right| \leq t,$$

where $\bar{\sigma}_k^2 = (1/n_k) \sum_{i \in \mathcal{I}_k} \text{Var}(\varepsilon_i)$.

In order to prove Theorem 4.1, we need to state and prove the following lemma.

Lemma 4.1 *Let \mathbf{B} and \mathbf{B}^* be two $d_r \times d_c$ matrices with $\|\mathbf{B} - \mathbf{B}^*\|_{\psi_2(\Pi)} \leq 2b^*$, and let $(\mathbf{X}_i)_{i=1}^n$ be a sequence of i.i.d. samples drawn from Π . Let $(\varepsilon_i)_{i=1}^n$ be a sequence of independent zero-mean σ^2 -sub-Gaussian random variables and let $\bar{\sigma}^2$ be defined as follows*

$$\bar{\sigma}^2 := \frac{1}{n} \sum_{i=1}^n \text{Var}(\varepsilon_i).$$

Then, for any $t > 0$, the inequality

$$\left| \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 - \left(\|\mathbf{B} - \mathbf{B}^*\|_{L^2(\Pi)}^2 + \bar{\sigma}^2 \right) \right| \leq t,$$

holds with probability at least

$$1 - 6 \exp \left[-C \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) n \right],$$

where $C > 0$ is a numerical constant.

Proof. Recall from Section 2 that the vector of all noise values $\{\varepsilon_i\}_{i=1}^n$ is denoted by E . Note that,

$$\begin{aligned} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 &= \|E - \mathfrak{X}(\mathbf{B} - \mathbf{B}^*)\|_2^2 \\ &= \|E\|_2^2 + \|\mathfrak{X}(\mathbf{B} - \mathbf{B}^*)\|_2^2 - 2\langle E, \mathfrak{X}(\mathbf{B} - \mathbf{B}^*) \rangle. \end{aligned}$$

Next, using our Lemma A.4 as well as Lemma 5.14 of Vershynin (2010), for each $i \in [n]$, we have

$$\left\| \varepsilon_i^2 - \mathbb{E}[\varepsilon_i^2] \right\|_{\psi_1} \leq 2\|\varepsilon_i^2\|_{\psi_1} \leq 4\|\varepsilon_i\|_{\psi_2}^2 \leq 2\sigma^2.$$

Also, $\|\mathbf{B} - \mathbf{B}^*\|_{\psi_2(\Pi)} \leq 2b^*$ means that $\|\langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle\|_{\psi_2} \leq 2b^*$, which in turn yields that $\langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle^2$ is subexponential. We can now follow a similar logic as the above and obtain

$$\left\| \langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle^2 - \mathbb{E}[\langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle^2] \right\|_{\psi_1} \leq 16b^{*2}.$$

In the same way, $\varepsilon_i \langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle$, which is a product of two sub-Gaussian random variables, becomes subexponential with a zero mean:

$$\|\varepsilon_i \langle \mathbf{X}_i, \mathbf{B} - \mathbf{B}^* \rangle\|_{\psi_1} \leq 8\sigma b^*.$$

It follows from Corollary 5.17 of Vershynin (2010) that, by defining

$$\begin{aligned} \mathcal{E}_1(t) &:= \left\{ \left| \|E\|_2^2 - \mathbb{E}[\|E\|_2^2] \right| \geq nt \right\}, \\ \mathcal{E}_2(t) &:= \left\{ \left| \|\mathfrak{X}(\mathbf{B} - \mathbf{B}^*)\|_2^2 - \mathbb{E}[\|\mathfrak{X}(\mathbf{B} - \mathbf{B}^*)\|_2^2] \right| \geq nt \right\}, \\ \mathcal{E}_3(t) &:= \left\{ \left| 2\langle E, \mathfrak{X}(\mathbf{B} - \mathbf{B}^*) \rangle - \mathbb{E}[2\langle E, \mathfrak{X}(\mathbf{B} - \mathbf{B}^*) \rangle] \right| \geq nt \right\}, \end{aligned}$$

we have

$$\mathbb{P}(\mathcal{E}_j) \leq 2 \exp \left[-C \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) n \right],$$

for $j \in \{1, 2, 3\}$ where $C > 0$ is a numerical constant. Applying the union bound, we get

$$\mathbb{P} \left(\left| \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 - \left[\|\mathbf{B} - \mathbf{B}^*\|_{L^2(\Pi)}^2 + \bar{\sigma}^2 \right] \right| \geq t \right) \leq 6 \exp \left[C \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) n \right],$$

for some (different) numerical constant $C > 0$. \square

Now we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. For all $\ell \in [L]$ and $k \in [K]$, define the bad event $\mathcal{E}_{\ell,k}$ to be

$$\mathcal{E}_{\ell,k} := \left\{ \left| \frac{1}{n_k} \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda_\ell))\|_2^2 - \left[\|\widehat{\mathbf{B}}_{-k}(\lambda_\ell) - \mathbf{B}^*\|_{L^2(\Pi)}^2 + \bar{\sigma}_k^2 \right] \right| > t \right\}.$$

It follows from Lemma 4.1, the union bound, and the assumption $n_k \geq n/(2K)$, for all $k \in [K]$, that the bad event satisfies

$$\mathbb{P} \left(\bigcup_{k \in [K]} \bigcup_{\ell \in [L]} \mathcal{E}_{\ell,k} \right) \leq 6KL \exp \left(-c \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) \cdot \frac{n}{K} \right),$$

for some constant $c \geq 0$.

Now, in the complement of the bad event, it follows from the convexity of $\|\cdot\|_{L^2(\Pi)}^2$ that

$$\begin{aligned} \|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 &\leq \sum_{k=1}^K \frac{n_k}{n} \cdot \|\widehat{\mathbf{B}}_{-k}(\widehat{\lambda}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 \\ &\leq \sum_{k=1}^K \frac{n_k}{n} \cdot \left[\frac{1}{n_k} \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\widehat{\lambda}))\|_2^2 - \bar{\sigma}_k^2 + t \right] \\ &= \widehat{E}(\widehat{\lambda}) - \bar{\sigma}^2 + t, \end{aligned}$$

which is the desired result. \square

Before stating the next result, we also define notations $\Sigma_{R,-k}$ and Σ_{-k} as follows:

$$\Sigma_{-k} = \sum_{i \in \mathcal{I}_{-k}} \varepsilon_i \mathbf{X}_i \quad \text{and} \quad \Sigma_{R,-k} = \sum_{i \in \mathcal{I}_{-k}} \zeta_i \mathbf{X}_i,$$

where, as in Section 3, $\{\zeta_i\}_{i \in [n]}$ are i.i.d. Rademacher random variables.

Now, we are ready to state the main result of this section, which is obtained by combining Theorem 4.1 with Corollary 3.2.

Theorem 4.2 *Let Assumptions 3.1 to 3.3 hold, and let $\ell_0 \in [L]$ be such that λ_{ℓ_0} (in Λ) is larger than $Cb^*\mathfrak{c} \max_{k \in [K]} \mathbb{E}[\|\Sigma_{R,-k}\|_{\text{op}}]$, where C is an arbitrary (and positive) constant, and \mathfrak{c} is defined as in Section 3. Also assume that Λ , $\widehat{\lambda}_{\text{cv}}$, and $\widehat{\mathbf{B}}_{\text{cv}}$ are defined as above. In addition, assume that $\{\varepsilon_i\}_{i=1}^n$ are independent zero-mean σ^2 -sub-Gaussian random variables. Then, for all $t > 0$, we have*

$$\|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 \leq \frac{C_1 \gamma_{\max} \lambda_{\ell_0}^2 r}{\gamma_{\min}^2} + 2t$$

with probability at least

$$\begin{aligned} &1 - 6KL \exp \left[-C_2 \min \left(\frac{t^2}{\sigma^4 \vee b^{*4}}, \frac{t}{\sigma^2 \vee b^{*2}} \right) \cdot \frac{n}{K} \right] \\ &- \sum_{k \in [K]} \mathbb{P}(\lambda_{\ell_0} \geq 3\|\Sigma_{-k}\|_{\text{op}}) \\ &- K \exp \left(-\frac{C_3 n \lambda_{\ell_0}^2 r}{\mathfrak{c}^2 b^{*2} \gamma_{\min}} \right), \end{aligned} \tag{4.4}$$

where C_1 , C_2 , and C_3 are positive constants.

Proof. The definition of $\widehat{\lambda}_{\text{cv}}$, together with Theorem 4.1, Corollary 3.2, Assumption 3.2, and the union bound, yields

$$\begin{aligned} \|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 &\leq \widehat{E}(\widehat{\lambda}_{\text{cv}}) - \bar{\sigma}^2 + t \\ &\leq \widehat{E}(\lambda_{\ell_0}) - \bar{\sigma}^2 + t \end{aligned}$$

$$\begin{aligned}
 &= \sum_{k=1}^K \frac{n_k}{n} \cdot \left[\frac{1}{n_k} \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda_{\ell_0}))\|_2^2 - \bar{\sigma}_k^2 + t \right] \\
 &\leq \sum_{k=1}^K \frac{n_k}{n} \cdot \left[\|\widehat{\mathbf{B}}_{-k}(\lambda_{\ell_0}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 + 2t \right] \\
 &\leq \sum_{k=1}^K \frac{n_k}{n} \cdot \left[\frac{C_1 \gamma_{\max} \lambda_{\ell_0}^2 r}{\gamma_{\min}^2} + 2t \right] \\
 &= \frac{C_1 \gamma_{\max} \lambda_{\ell_0}^2 r}{\gamma_{\min}^2} + 2t
 \end{aligned}$$

with the probability stated in Eq. (4.4). Note that we also used the fact that $|\mathcal{I}_{-k}| \geq n[1 - 1/(2K)] \geq n/2$ in the last term of Eq. (4.4). \square

Remark 4.1 While the tail bound of Theorem 4.2 is stated for $\|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}}) - \mathbf{B}^*\|_{L^2(\Pi)}^2$, it is straightforward to use Assumption 3.2 to obtain the following bound on $\|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}}) - \mathbf{B}^*\|_{\text{F}}^2$ with the same probability

$$\|\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda}_{\text{cv}}) - \mathbf{B}^*\|_{\text{F}}^2 \leq \frac{C_1 \gamma_{\max} \lambda_{\ell_0}^2 r}{\gamma_{\min}^3} + \frac{2t}{\gamma_{\min}}.$$

Remark 4.2 In light of Remark 3.5, and the fact that the above upper bound is the same as that of Corollary 3.2 up to constants, our bounds for the cross-validated estimator are also tight for the matrix completion and the compressed sensing settings.

Remark 4.3 (Variants of cross-validation) Recall that Theorem 4.1 shows that the test error $(1/n_k) \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda_{\ell}))\|_2^2$ concentrates around its expectation $\|\widehat{\mathbf{B}}_{-k}(\lambda_{\ell}) - \mathbf{B}^*\|_{L^2(\Pi)}^2 + \bar{\sigma}_k^2$, for every fold k and every regularization parameter λ_{ℓ} . Therefore, one can obtain a similar bound as in Theorem 4.2 when the cross-validation approach is modified. One alternative is to choose the optimum penalty using a single training and test split. Another option is to create K random splits of the data into training and test sets and, similar to Eq. (4.1), define $\widehat{\mathbf{B}}_{\text{cv}}(\widehat{\lambda})$ as a weighted average of the estimator for each fold. The proof of Theorem 4.1 stays valid since it requires only the convexity of $\|\cdot\|_{L^2(\Pi)}^2$ and the concentration of the test error for every split. However, the approach that uses a single split is expected to have a higher variance. Finally, often in practice after $\widehat{\lambda}_{\text{cv}}$ is selected via some form of cross-validation, the estimator is refitted on the entire training set with the selected penalty parameter. Our current theoretical analysis does not provide a guarantee for this version of cross-validation since the independence of ε_i and λ in our analysis would break down. However, we do compare our approach with this refitting version in the simulations of Appendix C.

Remark 4.4 (Selecting the λ -sequence) It is important to note that the upper bound of Theorem 4.2 depends on λ_{ℓ_0} . The bound is practically relevant if λ_{ℓ_0} is close to the theoretically optimal one. However, one can ensure that, up to constants, such an λ_{ℓ_0} exists in the sequence Λ by selecting Λ as follows: let λ_{\max} be the minimum real number for which the only minimizer of the convex program is zero. It can be easily shown that $\lambda_{\max} = \|\sum_{i=1}^n Y_i \mathbf{X}_i\|_{\text{op}}$. Then, we

select the sequence of values of $(\lambda)_{i=1}^L$ as follows: $\lambda_1 := \lambda_{\max}$, $\lambda_{\ell+1} = \lambda_\ell/2$, and L is taken large enough such that $\lambda_L = 2^{-L+1}\lambda_{\max}$ is close to 0, thereby guaranteeing that the optimal penalty, up to a factor of 2, is included in the sequence. Note that such a large L impacts the upper bound only double logarithmically. Specifically, a larger L requires selecting a larger t to make the tail probability $6KL \exp[-C_2 \min(t^2/(\sigma^4 \vee b^{*4}), t/(\sigma^2 \vee b^{*2})) \cdot (n/K)]$ small, which means, up to constants, that the term $2t$ in the upper bound of Theorem 4.2 is

$$\sigma^2 \vee b^{*2} \max \left(\frac{\log[K \log \frac{\lambda_{\max}}{\lambda_L}]}{n/K}, \sqrt{\frac{\log[K \log \frac{\lambda_{\max}}{\lambda_L}]}{n/K}} \right).$$

This is the reason, in practice, why even $L = 100$ leads to a good cross-validated estimator.

In the next section, we numerically evaluate the effectiveness of this cross-validated estimator.

4.1 Simulations

In this section, we test the empirical performance of the cross-validated estimator, using synthetic and real data.

For the synthetic data setting, we generate a $d \times d$ matrix \mathbf{B}^* of rank r . Following a similar approach as in (Keshavan et al., 2010), we first generate $d \times r$ matrices \mathbf{B}_L^* and \mathbf{B}_R^* with independent standard normal entries and then set $\mathbf{B}^* := \mathbf{B}_L^* \cdot \mathbf{B}_R^{*\top}$. For the real data setting, we use movie ratings from the MovieLens³ data set, from which we take the d most frequently watched movies and the d users that rated the most movies. For d equal to 50 or 100, this leads to a matrix with over 93% observed ratings and we impute the remaining entries with the average of the observed ratings. We also center the matrix and scale it so that the overall mean and variance of the entries match the synthetic data, and refer to the resulting rating as matrix \mathbf{B}^* . For the distribution of observations, Π , we consider the matrix completion case. Specifically, for each $i \in [n]$, r_i and c_i are integers in $[d]$, selected independently and uniformly at random. Then, $\mathbf{X}_i = e_{r_i} e_{c_i}^\top$. This leads to n observations $Y_i = \mathbf{B}^*_{r_i c_i} + \varepsilon_i$, where ε_i are taken to be i.i.d. standard normal random variables.

Given these observations, we compare the estimation error of the following five different estimators:

1. **theory-1**, **theory-2**, and **theory-3** estimators solve the convex program Eq. (1.2) for a given value of $\lambda = \lambda_0$ that is motivated by the theoretical results. Specifically, by Remark 3.3, we need $\lambda_0 \geq 3\|\boldsymbol{\Sigma}\|_{\text{op}}$ to hold with high probability, which means that we select λ_0 so that $\lambda_0 \geq 3\|\boldsymbol{\Sigma}\|_{\text{op}}$ holds with probability 0.9. For each sample of size n , we find λ_0 by generating 1000 independent data sets of the same size and then, for the **theory-3** estimator, we choose the 100th biggest value of $3\|\boldsymbol{\Sigma}\|_{\text{op}}$. While the main theory-inspired estimator is **theory-3**, as we will see below, it performs very poorly since the constant 3 as multiplier of $\|\boldsymbol{\Sigma}\|_{\text{op}}$ may be too conservative. Therefore, we also consider two other variants of this estimator where constant 3 is replaced with constants 1 and 2 and denote these estimators by **theory-1** and **theory-2**, respectively. Overall, we highlight that these three estimators cannot be used in practice since they need access to $\|\boldsymbol{\Sigma}\|_{\text{op}}$.

3. The file `ratings.csv` from <https://www.kaggle.com/datasets/grouplens/movielens-20m-dataset>.

2. The `oracle` estimator solves the convex program Eq. (1.2) over a set of regularization parameters Λ . Then, the estimate is obtained by picking the matrix $\widehat{\mathbf{B}}$ that has the minimum distance to the ground truth matrix \mathbf{B}^* under the Frobenius norm. The set Λ that is used in this estimator is obtained as follows: by Remark 4.4, λ_{\max} is the minimum real number for which the only minimizer of the convex program is zero, i.e., $\lambda_{\max} = \|\sum_{i=1}^n Y_i \mathbf{X}_i\|_{\text{op}}$. Then, we set $\lambda_{\min} := \min(\lambda_0/10, 10^{-5})$, and then the sequence of values of $(\lambda)_{i=1}^L$ is generated as follows: $\lambda_1 := \lambda_{\max}$ and $\lambda_{\ell+1} = \lambda_{\ell}/2$ such that L is the smallest integer with $\lambda_L \leq \lambda_{\min}$.
3. The `cv` estimator is introduced at the beginning of this section. We use 10 folds (i.e., $K = 10$) and use a set of regularization parameters $\Lambda' = \{\lambda\}_{i=1 \in [L']}$ constructed exactly like the ones in the `oracle` estimator; however, since `cv` does not have access to λ_0 , L' is the smallest integer with $\lambda_{L'} \leq 0.001\lambda_{\max}$.

Finally, for each of these estimators, we compute the *relative error* of the estimate $\widehat{\mathbf{B}}$ from the ground truth \mathbf{B}^* defined as $\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 / \|\mathbf{B}^*\|_{\text{F}}^2$ for a range of n . The results, averaged over 100 runs with 2SE error bars, are shown in Figure 1, for two instances $(d, r) = (50, 2)$ and $(d, r) = (100, 3)$. We can see that `cv` performs close to the `oracle` and outperforms the theoretical estimators. We also present the distribution of the selected λ by each method, over all 100 runs, in Figure 2. There is no variation for the selected λ in the theory-inspired estimators as we selected their λ_0 at the beginning of the experiment, using 1,000 data sets. It is also noticeable that the penalty selected by `cv` is close to the one selected by `oracle`, while the ones suggested by the theory are larger.

Sensitivity to a different noise level or cross-validation approach. Recall that in the above simulations the noise standard deviation is 1. We repeat the above simulations for two more values of the noise standard deviation, 0.5 and 2. In addition, and motivated by Remark 4.3, we also test three more variations of cross-validation. The results, presented in Appendix C, validate the above observations.⁴

5. Applications to Noisy Recovery

In this section, we will show the benefit of proving Corollary 3.2, which provides an upper bound with a more general norm $\mathfrak{N}(\cdot)$. We will look at four different special cases for the distribution Π and in two cases (matrix completion and compressed sensing with Gaussian ensembles) we will recover existing results by Negahban and Wainwright (2012), Klopp (2014), and Candes and Recht (2009). For the other two (multi-task learning and compressed sensing with factored measurements), stated as open problems in (Rohde and Tsybakov, 2011) and (Recht et al., 2010), we will obtain the first (to the best of our knowledge) such results. Overall, in order to apply Corollary 3.2 in each case, we need to perform the following:

1. Choose a norm $\mathfrak{N}(\cdot)$. In the examples below, we will be using $\|\mathbf{B}\|_{\psi_p(\Pi)}$ for parameters p that are determined in each case.

4. All codes are available in this repository <https://github.com/mohsenbayati/cv-impute>.




Figure 1: Comparison of the relative error (i.e., $\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\mathbb{F}}^2 / \|\mathbf{B}^*\|_{\mathbb{F}}^2$) of the proposed estimators, for the synthetic and MovieLens data.

2. Compute $\|\mathbf{B}\|_{L^2(\Pi)}$ to find an appropriate constant γ_{\min} for which Assumption 3.2 holds.
3. Compute $\mathfrak{N}(\mathbf{B}^*)$ to obtain the constant b^* .
4. Choose an appropriate constant \mathfrak{c} such that Assumption 3.3 holds.
5. Apply Proposition B.1 of Appendix B.2 to obtain a bound for $\mathbb{P}(\lambda < 3\|\boldsymbol{\Sigma}\|_{\text{op}})$ as well as calculate $\mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}}]$.

To simplify the notation, we assume that $d_r = d_c = d$ throughout this section; however, it is easy to see that the arguments hold for $d_r \neq d_c$ as well. We also assume, for simplicity, that $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ for all $i \in [n]$.




Figure 2: Comparison of the selected penalty parameter λ of the proposed estimators, for the synthetic and MovieLens data.

5.1 Matrix completion

Let \mathbf{B}^* be a $d \times d$ matrix and recall that e_1, e_2, \dots, e_d denotes the standard basis for \mathbb{R}^d . Let, also, for each $i \in [n]$, r_i and c_i be integers in $[d]$, selected independently and uniformly at random. Then, let $\mathbf{X}_i = \xi_i \cdot e_{r_i} e_{c_i}^\top$ where, for each i , ξ_i is an independent $4d^2$ -sub-Gaussian random variable that is also independent of r_j and c_j , $j \in [n]$. If we set $\xi_i := d$ almost surely, then $\|\xi_i\|_{\psi_2} = d/\sqrt{\log 2} \leq 2d$. This corresponds to the problem studied in (Negahban and Wainwright, 2012). Here we show the bounds for the slightly more general case of $\xi_i \sim \mathcal{N}(0, d^2)$.

First, note that

$$\|\mathbf{B}\|_{L^2(\Pi)} = \|\mathbf{B}\|_F,$$

which means that γ_{\min} is equal to 1. In order to find a suitable norm $\mathfrak{N}(\cdot)$, we next study $\|\mathbf{B}\|_{\psi_2(\Pi)}$ to see how heavy-tailed $\langle \mathbf{B}, \mathbf{X}_i \rangle$ is. We have

$$\begin{aligned} \mathbb{E} \left[\exp \left(\frac{|\langle \mathbf{B}, \mathbf{X}_i \rangle|^2}{4d^2 \|\mathbf{B}\|_\infty^2} \right) \right] &= \frac{1}{d^2} \sum_{j,k=1}^d \mathbb{E} \left[\exp \left(\frac{\xi_i^2 \mathbf{B}_{jk}^2}{4d^2 \|\mathbf{B}\|_\infty^2} \right) \right] \\ &= \frac{1}{d^2} \sum_{j,k=1}^d \left[\frac{1}{\sqrt{\left(1 - \frac{\mathbf{B}_{jk}^2}{2\|\mathbf{B}\|_\infty^2}\right)_+}} \right] \\ &\leq 2, \end{aligned}$$

where the second equality follows from Lemma A.1 of Appendix A. Therefore,

$$\|\mathbf{B}\|_{\psi_2(\Pi)} \leq 2d\|\mathbf{B}\|_\infty,$$

which guides our selection of $\mathfrak{N}(\cdot) = d\|\cdot\|_\infty$ and $b^* = 2d\|\mathbf{B}^*\|_\infty$. We can now see that $\mathbf{c} = 9$ fulfils Assumption 3.3. The reason for this is that, given $\mathfrak{N}(\mathbf{B}) = d\|\mathbf{B}\|_\infty = 1$, we can condition on r_i and c_i and use Corollaries A.1 and A.2 of Appendix A to obtain

$$\begin{aligned} \mathbb{E}[\xi_i^2 \mathbf{B}_{r_i c_i}^2 \cdot \mathbb{I}(|\xi_i \mathbf{B}_{r_i c_i}| \leq 9) \mid r_i, c_i] &\geq \mathbb{E} \left[\xi_i^2 \mathbf{B}_{r_i c_i}^2 \cdot \mathbb{I} \left(|\xi_i \mathbf{B}_{r_i c_i}| \leq 5\sqrt{\frac{8}{3}} d \|\mathbf{B}_{r_i c_i}\| \right) \mid r_i, c_i \right] \\ &\geq \frac{\mathbb{E}[\xi_i^2 \mathbf{B}_{r_i c_i}^2 \mid r_i, c_i]}{2}. \end{aligned}$$

Therefore, we can take the expectation with respect to r_i and c_i and use the tower property to show that Assumption 3.3 holds for $\mathbf{c} = 9$.

The next step is to use Proposition B.1 of Appendix B.2, with specification $\mathbf{Z}_i := \varepsilon_i \mathbf{X}_i$, to find a tail bound inequality for $\mathbb{P}(\lambda < 3\|\boldsymbol{\Sigma}\|_{\text{op}})$. Define $\delta := d\sigma e/(e-1)$ and let G_1 and G_2 be two independent standard normal random variables. It follows that

$$\begin{aligned} \mathbb{E} \left[\exp \left(\frac{\|\mathbf{Z}_i\|_{\text{op}}}{\delta} \right) \right] &= \mathbb{E} \left[\exp \left(\frac{d\sigma |G_1 G_2|}{\delta} \right) \right] \\ &\leq \mathbb{E} \left[\exp \left(\frac{d\sigma (G_1^2 + G_2^2)}{2\delta} \right) \right] \\ &= \mathbb{E} \left[\exp \left(\frac{d\sigma G_1^2}{2\delta} \right) \right]^2 \\ &= \mathbb{E} \left[\exp \left(\frac{(e-1)G_1^2}{2e} \right) \right]^2 \\ &= \left[\frac{1}{\sqrt{1 - \frac{e-1}{e}}} \right]^2 = e. \end{aligned}$$

Next, notice that

$$\mathbb{E}[\mathbf{Z}_i \mathbf{Z}_i^\top] = \mathbb{E}[\mathbf{Z}_i^\top \mathbf{Z}_i] = d\sigma^2 \mathbf{I}_d.$$

Therefore, applying Proposition B.1 of Appendix B.2 gives

$$\mathbb{P}(\lambda < 3\|\boldsymbol{\Sigma}\|_{\text{op}}) \leq 2d \exp \left[-\frac{Cn\lambda}{d\sigma} \left(\frac{\lambda}{\sigma} \wedge \frac{1}{\log d} \right) \right], \quad (5.1)$$

for some constant C .

We can follow the same argument for $\zeta_i \mathbf{X}_i$, and use Corollary B.1 of Appendix B.2 to obtain

$$\mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}}] \leq C_1 \sqrt{\frac{d \log d}{n}}, \quad (5.2)$$

provided that $n \geq C_2 d \log^3 d$ for constants C_1 and C_2 . We can now combine Eq. (5.1), Eq. (5.2), and Corollary 3.2 to obtain the following result: for any $\lambda \geq C_3 b^* \sqrt{d \log d/n}$ and $n \geq C_2 d \log^3 d$, the inequality

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 \leq C_4 \lambda^2 r$$

holds with probability at least

$$1 - 2d \exp \left[-\frac{Cn\lambda}{\sigma d} \left(\frac{\lambda}{\sigma} \wedge \frac{1}{\log d} \right) \right] - \exp \left(-\frac{C_5 n \lambda^2 r}{b^{*2}} \right).$$

In particular, setting

$$\lambda = C_6 (\sigma \vee b^*) \sqrt{\frac{\rho d}{n}}, \quad (5.3)$$

for some $\rho \geq \log d$, we have that

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 \leq C_7 (\sigma^2 \vee b^{*2}) \frac{\rho dr}{n} \quad (5.4)$$

with probability at least $1 - \exp(-C_8 \rho)$ whenever $n \geq C_2 d \log^3 d$. This result closely recovers Corollary 1 in Negahban and Wainwright (2012) whenever $\rho = \log d$.

5.2 Multi-task learning

As in Section 5.1, let \mathbf{B}^* be a $d \times d$ matrix and, for each $i \in [n]$, let r_i be an integer in $[n]$, selected independently and uniformly at random. Then let $\mathbf{X}_i = e_{r_i} \cdot X_i^\top$ where, for each i , X_i is an independent $\mathcal{N}(0, d \cdot \mathbf{I}_d)$ random vector that is also independent of $\{r_j\}_{j=1}^n$. It follows that

$$\|\mathbf{B}\|_{L^2(\Pi)} = \|\mathbf{B}\|_{\text{F}},$$

which means that $\gamma_{\min} = 1$. Also,

$$\|\mathbf{B}\|_{\psi_2(\Pi)} \leq 2\sqrt{d} \|\mathbf{B}\|_{2, \infty}.$$

To see the latter, for any $\mathbf{X} \sim \Pi$ we follow the same steps as in the previous section and obtain,

$$\begin{aligned} \mathbb{E}\left[\exp\left(\frac{|\langle \mathbf{B}, \mathbf{X} \rangle|^2}{4d\|\mathbf{B}\|_{2,\infty}^2}\right)\right] &= \frac{1}{d} \sum_{j=1}^d \mathbb{E}[X_j] \exp\left(\frac{|\langle \mathbf{B}, e_j \cdot X_i^\top \rangle|^2}{4d\|\mathbf{B}\|_{2,\infty}^2}\right) \\ &= \frac{1}{d} \sum_{j=1}^d \left[\frac{1}{\sqrt{\left(1 - \frac{\|B_j\|_2^2}{2\|\mathbf{B}\|_{2,\infty}^2}\right)_+}} \right] \\ &\leq 2, \end{aligned} \tag{5.5}$$

where B_j denotes the j^{th} row of \mathbf{B} and Eq. (5.5) uses Lemma A.1 of Appendix A, since $\langle B_j, X_i \rangle \sim \mathcal{N}(0, d\|B_j\|_2^2)$. The final step uses $\|B_j\|_2 \leq \|\mathbf{B}\|_{2,\infty}$, which follows from the definition of $\|\mathbf{B}\|_{2,\infty}$. As in Section 5.1, we use this to set $\mathfrak{N}(\cdot) = 2\sqrt{d}\|\mathbf{B}^*\|_{2,\infty}$, which means that we can select b^* to be equal to $2\sqrt{d}\|\mathbf{B}^*\|_{2,\infty}$. Also, as in Section 5.1, we can condition on the random variable r_i to show that

$$\mathbb{E}[(B_{r_i} X_i)^2 \mathbb{I}(|B_{r_i} X_i| \leq 9) \mid r_i] \geq \frac{1}{2} \mathbb{E}[(B_{r_i} X_i)^2 \mid r_i],$$

which means that $\mathfrak{c} = 9$ satisfies the requirement of Assumption 3.3.

Let \mathbf{Z}_i be as in Section 5.1, $\delta = d\sigma e/(e-1)$, and $(G_i)_{i=0}^d$ be a sequence of $d+1$ independent standard normal random variables. We see that

$$\begin{aligned} \mathbb{E}\left[\exp\left(\frac{\|\mathbf{Z}_i\|_{\text{op}}}{\delta}\right)\right] &= \mathbb{E}\left[\exp\left(\frac{\|\mathbf{Z}_i\|_{\text{F}}}{\delta}\right)\right] \\ &= \mathbb{E}\left[\exp\left(\frac{\sigma d |G_0| \sqrt{\frac{G_1^2 + \dots + G_d^2}{d}}}{\delta}\right)\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{\sigma d \left(G_0^2 + \frac{G_1^2 + \dots + G_d^2}{d}\right)}{2\delta}\right)\right] \\ &= \left(1 - \frac{\sigma d}{\delta}\right)^{-\frac{1}{2}} \cdot \left(1 - \frac{\sigma}{\delta}\right)^{-\frac{d}{2}} \\ &= \left(1 - \frac{e-1}{e}\right)^{-\frac{1}{2}} \cdot \left(1 - \frac{e-1}{ed}\right)^{-\frac{d}{2}} \\ &\leq \sqrt{e} \cdot e^{\frac{e-1}{2e}} \tag{using } 1+x \leq e^x \\ &\leq e. \end{aligned}$$

Furthermore, we have

$$\mathbb{E}[\mathbf{Z}_i \mathbf{Z}_i^\top] = d\sigma^2 \mathbf{I}_d \quad \text{and} \quad \mathbb{E}[\mathbf{Z}_i^\top \mathbf{Z}_i] = \sigma^2 \mathbf{I}_d.$$

This implies that Eq. (5.1) and Eq. (5.2) hold in this case as well. Since γ_{\min} and \mathfrak{c} are the same as Section 5.1, we conclude that Eq. (5.4) holds, with the same probability as in Section 5.1, when $n \geq C_2 d \log^3 d$.

5.3 Compressed sensing via Gaussian ensembles

Let \mathbf{B}^* be a $d \times d$ matrix. Let each \mathbf{X}_i be a random matrix with entries filled with i.i.d. samples drawn from $\mathcal{N}(0, 1)$. It follows that

$$\|\mathbf{B}\|_{L^2(\Pi)} = \|\mathbf{B}\|_F,$$

which means that $\gamma_{\min} = 1$, and

$$\|\mathbf{B}\|_{\psi_2(\Pi)} \leq 2\|\mathbf{B}\|_F.$$

To see the latter, since $\langle \mathbf{B}, \mathbf{X}_i \rangle \sim \mathcal{N}(0, \|\mathbf{B}\|_F^2)$ by Lemma A.1 of Appendix A, it follows that

$$\mathbb{E}\left[\exp\left(\frac{|\langle \mathbf{B}, \mathbf{X}_i \rangle|^2}{4\|\mathbf{B}\|_F^2}\right)\right] = \left[\frac{1}{\sqrt{(1 - \frac{1}{2})_+}}\right] \leq 2.$$

Therefore, as before, we can use $\mathfrak{N}(\cdot) = \|\cdot\|_{\psi_2(\Pi)}$ and $\mathfrak{c} = 9$. Therefore, a similar argument to those of Sections 5.1 and 5.2 shows that Eq. (5.4) holds in this setting as well. This bound closely recovers the bound of Candes and Plan (2011).

5.4 Compressed sensing via factored measurements

Recht et al. (2010) propose factored measurements to alleviate the need for a memory of size nd^2 to carry computations in the compressed sensing applications with large dimensions. The idea is to use measurement matrices of the form $\mathbf{X}_i = UV^\top$, where U and V are random vectors of length d . Even though UV^\top is a $d \times d$ matrix, we only need a memory of size $\mathcal{O}(nd)$ to store all the input, which is a significant improvement over the Gaussian ensembles of Section 5.3. We will now study this problem when U and V are both $\mathcal{N}(0, \mathbf{I}_d)$ vectors that are independent of each other. In this case we have

$$\begin{aligned} \|\mathbf{B}\|_{L^2(\Pi)}^2 &= \mathbb{E}\left[\langle \mathbf{B}, UV^\top \rangle^2\right] \\ &= \mathbb{E}\left[(V^\top \mathbf{B}^\top U)(U^\top \mathbf{B} V)\right] \\ &= \mathbb{E}\left[V^\top \mathbf{B}^\top \cdot \mathbb{E}[UU^\top \mid V] \cdot \mathbf{B} V\right] \\ &= \mathbb{E}\left[V^\top \mathbf{B}^\top \mathbf{B} V\right] \\ &= \mathbb{E}\left[\text{Tr}(V^\top \mathbf{B}^\top \mathbf{B} V)\right] \\ &= \mathbb{E}\left[\text{Tr}(\mathbf{B} V V^\top \mathbf{B}^\top)\right] \\ &= \text{Tr}(\mathbf{B} \mathbb{E}[V V^\top] \mathbf{B}^\top) \\ &= \text{Tr}(\mathbf{B} \mathbf{B}^\top) \\ &= \|\mathbf{B}\|_F^2, \end{aligned}$$

which means that the parameter γ_{\min} can be selected to be 1 again. Next, let $\mathbf{B} = \mathbf{O}_1 \mathbf{D} \mathbf{O}_2^\top$ be the singular value decomposition of \mathbf{B} . Then, we get

$$\langle \mathbf{B}, UV^\top \rangle = U^\top \mathbf{B} V$$

$$\begin{aligned}
 &= U^\top \mathbf{O}_1 \mathbf{D} \mathbf{O}_2^\top V \\
 &= (\mathbf{O}_1^\top U)^\top \mathbf{D} (\mathbf{O}_2^\top V).
 \end{aligned}$$

As the distribution of U and V is invariant by the multiplication of unitary matrices, for any $t > 0$, we have

$$\begin{aligned}
 \mathbb{E} \left[\exp \left(\frac{\langle \mathbf{B}, UV^\top \rangle}{t} \right) \right] &= \mathbb{E} \left[\exp \left(\frac{\langle \mathbf{D}, UV^\top \rangle}{t} \right) \right] \\
 &= \mathbb{E} \left[\exp \left(\frac{U^\top \mathbf{D} V}{t} \right) \right] \\
 &= \mathbb{E}[U] \mathbb{E}[V] \exp \left(\frac{U^\top \mathbf{D} V}{t} \right) U \\
 &= \mathbb{E} \left[\exp \left(\frac{\|U^\top \mathbf{D}\|^2}{2t^2} \right) \right] \\
 &= \mathbb{E} \left[\exp \left(\frac{\sum_{i=1}^d U_i^2 \mathbf{D}_{ii}^2}{2t^2} \right) \right] \\
 &= \prod_{i=1}^d \mathbb{E} \left[\exp \left(\frac{U_i^2 \mathbf{D}_{ii}^2}{2t^2} \right) \right] \\
 &= \prod_{i=1}^d \frac{1}{\sqrt{\left(1 - \frac{\mathbf{D}_{ii}^2}{t^2}\right)_+}}.
 \end{aligned}$$

By Lemma A.3 of Appendix A, the necessary condition for $\|\langle \mathbf{B}, UV^\top \rangle\|_{\psi_1} \leq t$ to hold is

$$\mathbb{E} \left[\exp \left(\frac{\langle \mathbf{B}, UV^\top \rangle}{t} \right) \right] \leq 2. \tag{5.6}$$

This, in particular, implies that

$$\frac{1}{\sqrt{\left(1 - \frac{\mathbf{D}_{ii}^2}{t^2}\right)_+}} \leq 2$$

or equivalently

$$\frac{\mathbf{D}_{ii}^2}{t^2} \leq \frac{3}{4},$$

for all $i \in [d]$. By taking the derivatives and by applying the concavity of logarithm, we can observe that $-2x \leq \log(1-x) \leq -x$ for all $x \in [0, \frac{3}{4}]$. This implies that, whenever Eq. (5.6) holds, we have

$$\frac{\mathbf{D}_{ii}^2}{2t^2} \leq \log \left(1 - \frac{\mathbf{D}_{ii}^2}{t^2} \right)_+^{-\frac{1}{2}} \leq \frac{\mathbf{D}_{ii}^2}{t^2},$$

and thus

$$\exp\left(\frac{\sum_{i=1}^d \mathbf{D}_{ii}^2}{2t^2}\right) \leq \mathbb{E}[\exp\left(\frac{\langle \mathbf{B}, UV^\top \rangle}{t}\right)] \leq \exp\left(\frac{\sum_{i=1}^d \mathbf{D}_{ii}^2}{t^2}\right).$$

Using $\|\mathbf{B}\|_{\mathbb{F}}^2 = \sum_{i=1}^d \mathbf{D}_{ii}^2$, we can simplify the above to

$$\exp\left(\frac{\|\mathbf{B}\|_{\mathbb{F}}^2}{2t^2}\right) \leq \mathbb{E}[\exp\left(\frac{\langle \mathbf{B}, UV^\top \rangle}{t}\right)] \leq \exp\left(\frac{\|\mathbf{B}\|_{\mathbb{F}}^2}{t^2}\right). \quad (5.7)$$

Combining the above results, we conclude that $\|\langle \mathbf{B}, \mathbf{X} \rangle\|_{\psi_1} \leq t$ implies

$$\frac{2\mathbf{D}_{11}}{\sqrt{3}} \leq t \quad \text{and} \quad \frac{\|\mathbf{B}\|_{\mathbb{F}}}{\sqrt{2\log 2}} \leq t.$$

Therefore, we have

$$\frac{1}{\sqrt{2\log 2}} \|\mathbf{B}\|_{\mathbb{F}} \leq \|\mathbf{B}\|_{\psi_1(\Pi)}. \quad (5.8)$$

Next, define

$$t := \max\left\{\frac{2\mathbf{D}_{11}}{\sqrt{3}}, \frac{\|\mathbf{B}\|_{\mathbb{F}}}{\sqrt{\log 2}}\right\} = \frac{\|\mathbf{B}\|_{\mathbb{F}}}{\sqrt{\log 2}}.$$

Since $D_{11}^2/t^2 \leq 3/4$, we can use Eq. (5.7) to obtain

$$\mathbb{E}\left[\exp\left(\frac{\langle \mathbf{B}, UV^\top \rangle}{t}\right)\right] \leq 2.$$

Using Lemma A.3 of Appendix A, we conclude that

$$\|\mathbf{B}\|_{\psi_1(\Pi)} \leq \frac{8}{\sqrt{\log 2}} \|\mathbf{B}\|_{\mathbb{F}}.$$

Now, by setting $\mathfrak{N}(\cdot) = \|\cdot\|_{\psi_1(\Pi)}$, given that the ratio $\|\mathbf{B}\|_{\psi_1(\Pi)}/\|\mathbf{B}\|_{\mathbb{F}}$ is at most $8/\sqrt{\log(2)}$, we apply Corollary A.1 of Appendix A to $Z = \langle \mathbf{B}, \mathbf{X} \rangle$ to see that $\mathfrak{c} = 53$ satisfies Assumption 3.3.

Now, in order to bound $\mathbb{P}(\lambda < 3\|\Sigma\|_{\text{op}})$ we need to use a truncation argument for the noise. Specifically, let

$$\bar{\varepsilon}_i := \varepsilon_i \mathbb{I}\left[|\varepsilon_i| \leq C_\varepsilon \sigma \sqrt{\log d}\right],$$

for a large enough constant C_ε . Next, using a union bound and defining $\bar{\Sigma} := \sum_{i=1}^n \bar{\varepsilon}_i \mathbf{X}_i$, we have

$$\begin{aligned} \mathbb{P}(\lambda < 3\|\Sigma\|_{\text{op}}) &\leq \mathbb{P}(\lambda < 3\|\bar{\Sigma}\|_{\text{op}}) + \sum_{i=1}^n \mathbb{P}\left(|\varepsilon_i| > C_\varepsilon \sigma \sqrt{\log d}\right) \\ &\leq \mathbb{P}(\lambda < 3\|\bar{\Sigma}\|_{\text{op}}) + 2ne^{-\frac{C_\varepsilon^2}{2} \log d}. \end{aligned}$$

Now, we define $\delta := 2C_\varepsilon\sigma d\sqrt{\log d}$ and $\mathbf{Z}_i = \bar{\varepsilon}_i\mathbf{X}_i$, as in Section 5.1, and use Proposition B.1. Let $(G_i)_{i=1}^{2d}$ be a sequence of $2d + 1$ independent standard normal random variables. We can use similar steps as in Sections 5.1 and 5.2 to obtain

$$\begin{aligned} \mathbb{E}\left[\exp\left(\frac{\|\mathbf{Z}_i\|_{\text{op}}}{\delta}\right)\right] &= \mathbb{E}\left[\exp\left(\frac{|\bar{\varepsilon}_i|\sqrt{\sum_{j=1}^d G_j^2}\sqrt{\sum_{j=d+1}^{2d} G_j^2}}{\delta}\right)\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{|\bar{\varepsilon}_i|\sum_{j=1}^{2d} G_j^2}{2\delta}\right)\right] \\ &\leq \mathbb{E}\left[\mathbb{E}\left[\exp\left(\frac{|\bar{\varepsilon}_i|G_1^2}{2\delta}\right)\right]^{2d}\middle|\bar{\varepsilon}_i\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{G_1^2}{4d}\right)\right]^{2d} \\ &\leq \left(\frac{1}{1-\frac{1}{2d}}\right)^{2d} \\ &\leq e. \end{aligned}$$

Furthermore, we have

$$\mathbb{E}[\mathbf{Z}_i\mathbf{Z}_i^\top] \preceq d\sigma^2\mathbf{I}_d \quad \text{and} \quad \mathbb{E}[\mathbf{Z}_i^\top\mathbf{Z}_i] \preceq \sigma^2\mathbf{I}_d.$$

Therefore, the following slight variation of Eq. (5.1) holds:

$$\mathbb{P}(\lambda < 3\|\boldsymbol{\Sigma}\|_{\text{op}}) \leq 2d\exp\left[-\frac{Cn\lambda}{d\sigma}\left(\frac{\lambda}{\sigma} \wedge \frac{1}{\log^{3/2}d}\right)\right] + 2ne^{-\frac{C_\varepsilon^2}{2}\log d}. \quad (5.9)$$

However, Eq. (5.2) stays unchanged since for ζ_i , unlike for ε_i , we do not need to use any truncation. This means that we can define λ as in Eq. (5.3) and obtain a bound as in Eq. (5.4) with probability at least $1 - \exp(-C_8\rho)$ whenever $\rho \geq \log d$ and $n \geq C_2d\log^4 d$. This bound recovers Theorem 2.3 of Cai and Zhang (2015); however, their bound works for $n = O(rd)$ which is smaller than our bound when $r < \log^4 d$.

6. Applications to Exact Recovery

In this section we study the trace regression problem when there is no noise. It is known that, under certain assumptions, it is possible to recover the true matrix \mathbf{B}^* exactly, with high probability (Candes and Recht, 2009; Keshavan et al., 2009). The discussion of Section 3.4 in (Negahban and Wainwright, 2012), in the context of the matrix-completion problem, we can deduce that that upper bounds that rely on the spikiness of \mathbf{B}^* are not strong enough to guarantee exact recovery of \mathbf{B}^* , even in the noiseless setting. However, we will show in this section that the methodology from Section 3 can be used to prove exact recovery for the two cases of compressed sensing studied in previous sections (Sections 5.3 and 5.4). After that, we will conclude this section with a brief discussion on exact recovery for the multi-task learning case (Section 5.2).

For any arbitrary sampling operator $\mathfrak{X}(\cdot)$, let \mathcal{S} be defined as follows:

$$\mathcal{S} := \left\{ \mathbf{B} \in \mathbb{R}^{d_r \times d_c} : \mathfrak{X}(\mathbf{B}) = Y \right\}.$$

Using $\sigma_\varepsilon = 0$ and the linear model Eq. (2.1), one can verify that $\mathbf{B}^* \in \mathcal{S}$ and thus that \mathcal{S} is not empty. The definition of \mathcal{S} implies that \mathcal{S} is an affine space and is thus convex. Next, note that, for any $\mathbf{B} \in \mathcal{S}$, the following identity holds:

$$\mathcal{L}(\mathbf{B}) = \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 + \lambda \|\mathbf{B}\|_* = \lambda \|\mathbf{B}\|_*.$$

Therefore, the minimizers of the optimization problem

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B})\|_2^2 + \lambda \|\mathbf{B}\|_* \\ & \text{subject to} && \mathfrak{X}(\mathbf{B}) = Y, \end{aligned}$$

are also the minimizers of

$$\begin{aligned} & \text{minimize} && \|\mathbf{B}\|_* \\ & \text{subject to} && \mathfrak{X}(\mathbf{B}) = Y. \end{aligned}$$

Note that, in the above formulation, the convex problem does not depend on λ anymore, and so λ can be chosen arbitrarily. In the noiseless setting, $\|\Sigma\|_{\text{op}} = 0$, and so any $\lambda > 0$ satisfies Eq. (3.5). Therefore, if the RSC condition holds for $\mathfrak{X}(\cdot)$ with parameters α and β on the set $\mathcal{C}(0, \eta)$, Theorem 3.1 leads to

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 \leq \frac{8b^{*2}\beta}{\alpha}. \quad (6.1)$$

Now, defining ν_0 as

$$\nu_0 := \inf_{\mathbf{B} \neq 0} \frac{\|\mathbf{B}\|_{\text{F}}^2}{\mathfrak{N}(\mathbf{B})^2}, \quad (6.2)$$

one can easily observe that $\mathcal{C}(\nu_0, \eta) = \mathcal{C}(0, \eta)$. Moreover, assume that n is large enough so that

$$\mathbb{E} \left[\|\Sigma_R\|_{\text{op}} \right]^2 \leq \frac{\gamma_{\min}^2 \nu_0}{800\epsilon^2 \eta}, \quad (6.3)$$

where Σ_R is defined in Section 3. Combining this with Corollary 3.1, we obtain that

$$\begin{aligned} \frac{\|\mathfrak{X}(\mathbf{A})\|_2^2}{n} & \geq \frac{\gamma_{\min}}{4} \|\mathbf{A}\|_{\text{F}}^2 - \frac{93\eta\epsilon^2}{\gamma_{\min}} \mathbb{E} \left[\|\Sigma_R\|_{\text{op}} \right]^2 \\ & \geq \frac{\gamma_{\min}}{4} \|\mathbf{A}\|_{\text{F}}^2 - \frac{\gamma_{\min} \nu_0}{8} \\ & \geq \frac{\gamma_{\min}}{4} \|\mathbf{A}\|_{\text{F}}^2 - \frac{\gamma_{\min}}{8} \|\mathbf{A}\|_{\text{F}}^2 \\ & = \frac{\gamma_{\min}}{8} \|\mathbf{A}\|_{\text{F}}^2, \end{aligned}$$

with probability at least

$$1 - 2 \exp\left(-\frac{Cn\gamma_{\min}\nu_0}{\mathfrak{c}^2}\right),$$

for all $\mathbf{A} \in \mathcal{C}(\nu_0, \eta)$. This shows that $\mathfrak{X}(\cdot)$ satisfies the RSC condition with $\alpha = \gamma_{\min}/8$ and $\beta = 0$. As a result, from Eq. (6.1), we can deduce the following proposition.

Proposition 6.1 *Let ν_0 and n be as in Eq. (6.2) and Eq. (6.3). Then, the unique minimizer of the constraint optimization problem*

$$\begin{aligned} & \text{minimize} && \|\mathbf{B}\|_* \\ & \text{subject to} && \mathfrak{X}(\mathbf{B}) = Y \end{aligned}$$

is $\mathbf{B} = \mathbf{B}^*$, and hence exact recovery of \mathbf{B}^* is possible with probability at least $1 - 2 \exp\left(-\frac{Cn\gamma_{\min}\nu_0}{\mathfrak{c}^2}\right)$.

Next, we can use the above proposition to prove that exact recovery is possible for the two problems of compressed sensing with Gaussian ensembles (Section 5.3) and compressed sensing with factored measurements (Section 5.4). Note that in both cases, we have $\gamma_{\min} = 1$, $\nu_0 \geq 0.1$, and $\mathfrak{c} \leq 60$.

Therefore, in order to use Proposition 6.1, all we need to do is to find a lower bound for n such that Eq. (6.3) holds. We will study each case in turn.

1. **Compressed sensing with Gaussian ensembles.** Here, since entries of \mathbf{X}_i 's are i.i.d. $\mathcal{N}(0, 1)$ random variables, the entries of Σ_R are i.i.d. $\mathcal{N}(0, 1/n)$ random variables. We can thus use Theorem 5.32 of Vershynin (2010) to get

$$\mathbb{E}[\|\Sigma_R\|_{\text{op}}] \leq \frac{2\sqrt{d}}{\sqrt{n}}.$$

Therefore, Eq. (6.3) is satisfied if $n \geq Crd$, where $C > 0$ is a large enough constant.

2. **Compressed sensing with factored measurements.** Here, the observation matrices are of the form $U_i V_i^\top$, where U_i and V_i are independent vectors distributed according to $\mathcal{N}(0, \mathbf{I}_d)$. Note that we have $\|\mathbf{X}_i\|_{\text{op}} = \|U_i\|_2 \|V_i\|_2 \leq \|U_i\|_2^2 + \|V_i\|_2^2$. Then,

$$\begin{aligned} \|\|\mathbf{X}_i\|_{\text{op}}\|_{\psi_1} &\leq \|\| \|U_i\|_2^2 + \|V_i\|_2^2 \|_{\psi_1} \\ &\leq 2\|\| \|U_i\|_2^2 \|_{\psi_1} \\ &= O(d). \end{aligned}$$

By Eq. (3.9) of Koltchinskii et al. (2011b) we have

$$\mathbb{E}[\|\Sigma_R\|_{\text{op}}] = O\left(\sqrt{\frac{d \log(2d)}{n}} \vee \frac{d \log(2d)^2}{n}\right).$$

We can thus infer that Eq. (6.3) holds for all $n \geq Crd \log(d)$ where C is a large enough constant.

Therefore, Proposition 6.1 guarantees that, for n satisfying the conditions stated above, exact recovery is possible in each of the two aforementioned settings, with probability at least $1 - \exp(-C'n)$, where $C' > 0$ is a numerical constant.

Implications for multi-task learning. We can apply Proposition 6.1 to the multi-task learning case (Section 5.2) as well. We have $\gamma_{\min} = 1$ and $\mathfrak{c} = 9$, but, for ν_0 , we have

$$\begin{aligned} \nu_0 &= \inf_{\mathbf{B} \neq 0} \frac{\|\mathbf{B}\|_{\text{F}}^2}{\mathfrak{N}(\mathbf{B})^2} \\ &= \inf_{\mathbf{B} \neq 0} \frac{\|\mathbf{B}\|_{\text{F}}^2}{4d\|\mathbf{B}\|_{2,\infty}^2} \\ &= \frac{1}{4d}, \end{aligned}$$

where the infimum is achieved if and only if \mathbf{B} has exactly one non-zero row. Notice that, in contrast to the previous examples where we had $\nu_0 \geq 0.1$, ν_0 depends on the dimensions of the matrix.

It is straightforward to verify that

$$\begin{aligned} \|\|\mathbf{X}_i\|_{\text{op}}\|_{\psi_2} &= \left\| \|X_i\|_2 \right\|_{\psi_2} \\ &= O(\sqrt{d}). \end{aligned}$$

Therefore, a similar argument to the above shows that

$$\mathbb{E}[\|\Sigma_R\|_{\text{op}}] = O\left(\sqrt{\frac{d \log(2d)}{n}} \vee \frac{\sqrt{d} \log(2d)^{\frac{3}{2}}}{n}\right).$$

This, in turn, implies that, in order for Eq. (6.3) to hold, it suffices to have $n \geq Crd^2\sqrt{d}\log(2d)$, for a large enough constant C . In this case, Proposition 6.1 shows that exact recovery is possible with probability at least $1 - 2\exp(-\frac{C'n}{d})$. However, this result is trivial, since $n \geq Crd^2\sqrt{d}\log(2d)$ means that with high probability, each row is observed at least d times, and so each row can be reconstructed separately (without using the low-rank assumption). This result can not be improved without further assumptions, as it is possible in a rank-2 matrix that all rows are equal to each other except for one row, and that row can be reconstructed only if at least d observations are made for that row. Since this must hold for all rows, at least d^2 observations would be needed. Nonetheless, one can expect that with stronger assumptions than generalized spikiness, such as incoherence, the number of required observations can be reduced to $rd\log(d)$.

Acknowledgments

The authors gratefully acknowledge support of the National Science Foundation (CAREER award CMMI: 1554140) and Stanford Data Science Initiative.

Appendix A. Auxiliary proofs

Lemma A.1 *Let Z be an $\mathcal{N}(0, \sigma^2)$ random variable. Then, for all $\eta > 0$,*

$$\mathbb{E}\left[e^{\eta Z^2}\right] = \frac{1}{\sqrt{(1 - 2\sigma^2\eta)_+}}.$$

Proof. Easily follows by using the formula $\int_{-\infty}^{\infty} \exp(-\frac{t^2}{2a^2}) dt = \sqrt{2\pi a^2}$. \square

Lemma A.2 *Let Z be a nonnegative random variable such that $\|Z\|_{\psi_p} = \nu$ holds for some $p \geq 1$, and assume that $c > 0$ is given. Then, we have*

$$\mathbb{E}[Z^2 \cdot \mathbb{I}(Z \geq c)] \leq (2c^2 + 4c\nu + 4\nu^2) \cdot \exp\left(-\frac{c^p}{\nu^p}\right).$$

Proof. Without loss of generality, we can assume that Z has a density function $f(z)$ and $\nu = 1$. Moreover, let $F(z) := \mathbb{P}(Z \leq z)$ be the cumulative distribution function of Z . The assumption that $\|Z\|_{\psi_p} = 1$ together with the Markov inequality yields

$$F(z) \geq 1 - 2 \exp(-z^p).$$

Therefore,

$$\begin{aligned} \mathbb{E}[Z^2 \cdot \mathbb{I}(Z \geq c)] &= \int_c^{\infty} z^2 f(z) dz \\ &= \int_c^{\infty} (-z^2)[-f(z)] dz \\ &= c^2 [1 - F(c)] + \int_c^{\infty} 2z[1 - F(z)] dz \\ &\leq 2c^2 \cdot \exp(-c^p) + \int_c^{\infty} 4z \exp(-z^p) dz. \end{aligned}$$

Next, note that the function $f_c(p)$ defined as

$$f_c(p) := \frac{\int_c^{\infty} z \exp(-z^p) dz}{(c+1) \exp(-c^p)} = \frac{\int_c^{\infty} z \exp(c^p - z^p) dz}{c+1},$$

is decreasing in p . Therefore, we have

$$\begin{aligned} \int_c^{\infty} z \exp(-z^p) dz &\leq (c+1) \exp(-c^p) f_c(1) \\ &= (c+1) \exp(-c^p), \end{aligned}$$

using $f_c(1) = 1$. Therefore, we have

$$\mathbb{E}[Z^2 \cdot \mathbb{I}(Z \geq c)] \leq (2c^2 + 4c + 4) \cdot \exp(-c^p),$$

which is the desired result. \square

Corollary A.1 *Let Z be a random variable such that $\|Z\|_{\psi_p} = \nu$ holds for some $p \geq 1$ and $\mathbb{E}[Z^2] = \sigma^2$. Then, for*

$$c_{\sigma,p} := \nu \cdot \max \left\{ 5, \left[10 \log \left(\frac{2\nu^2}{\sigma^2} \right) \right]^{\frac{1}{p}} \right\}, \quad (\text{A.1})$$

we have

$$\mathbb{E}[Z^2 \cdot \mathbb{I}(|Z| \leq c_{\sigma,p})] \geq \frac{\mathbb{E}[Z^2]}{2}.$$

Proof. Without loss of generality, we can assume that $\nu = 1$. Using Lemma A.2 for $|Z|$ and any $c \geq 5$, we have

$$\begin{aligned} \mathbb{E}[Z^2 \cdot \mathbb{I}(|Z| \leq c)] &\geq \mathbb{E}[Z^2] - \mathbb{E}[Z^2 \cdot \mathbb{I}(|Z| \geq c)] \\ &\geq \mathbb{E}[Z^2] - 3c^2 \cdot \exp(-c^p). \end{aligned}$$

Next, it is easy to show that, for any $c \geq 5$,

$$3c^2 \cdot \exp\left(-\frac{9c^p}{10}\right) \leq 3c^2 \cdot \exp\left(-\frac{9c}{10}\right) \leq 1.$$

Therefore, letting c_σ be defined as in Eq. (A.1), we get

$$\begin{aligned} 3c_{\sigma,p}^2 \cdot \exp(-c_{\sigma,p}^p) &= 3c_{\sigma,p}^2 \cdot \exp\left(-\frac{9c_{\sigma,p}^p}{10}\right) \exp\left(-\frac{c_{\sigma,p}^p}{10}\right) \\ &\leq \exp\left(-\frac{c_{\sigma,p}^p}{10}\right) \\ &\leq \frac{\sigma^2}{2}, \end{aligned}$$

which completes the proof of this corollary. \square

Corollary A.2 *Let Z be an $\mathcal{N}(0, \sigma^2)$ random variable. Then, the constant $c_{\sigma,2}$ defined in Eq. (A.1) satisfies $c_{\sigma,2} \leq 5\|Z\|_{\psi_2}$.*

Proof. By Lemma A.1, we obtain $\nu = \|Z\|_{\psi_2} = \sqrt{8\sigma^2/3}$, which means that $\nu^2/\sigma^2 = 8/3$. The rest follows from Corollary A.1. \square

The Orlicz norm of a random variable is defined in terms of the absolute value of that random variable, and it is usually easier to work with the random variable than with its absolute value. The next lemma relates the Orlicz norm to the moment-generating function of a random variable.

Lemma A.3 *Let X be a zero-mean random variable and*

$$\alpha := \inf \left\{ t > 0 : \max \left\{ \mathbb{E} \left[\exp \left(\frac{X}{t} \right) \right], \mathbb{E} \left[\exp \left(-\frac{X}{t} \right) \right] \right\} \leq 2 \right\}.$$

Then, we have

$$\alpha \leq \|X\|_{\psi_1} \leq 8\alpha.$$

Proof. The first inequality, $\alpha \leq \|X\|_{\psi_1}$, follows from the monotonicity of the exponential function. For the second inequality, note that for any $t > 0$,

$$\mathbb{E} \left[\exp \left(\frac{|X|}{t} \right) \right] = \mathbb{E} \left[\exp \left(\frac{|X| - \mathbb{E}[|X|]}{t} \right) \right] \cdot \exp \left(\frac{\mathbb{E}[|X|]}{t} \right). \quad (\text{A.2})$$

Now, the union bound and Markov inequality lead to the following tail bound for $|X|$:

$$\mathbb{P}(|X| \geq x) \leq \mathbb{P}(X \geq x) + \mathbb{P}(-X \geq x)$$

$$\leq 4 \exp\left(-\frac{x}{\alpha}\right).$$

Hence, we have

$$\begin{aligned} \mathbb{E}[|X|] &= \int_0^\infty \mathbb{P}(|X| \geq x) dx \\ &\leq 4\alpha. \end{aligned}$$

Next, assuming that X' is an independent copy of X and ζ is a Rademacher random variable independent of X and X' , we have

$$\begin{aligned} \mathbb{E}\left[\exp\left(\frac{|X| - \mathbb{E}[|X|]}{t}\right)\right] &= \mathbb{E}\left[\exp\left(\frac{|X| - \mathbb{E}[|X'|]}{t}\right)\right] \\ &\leq \mathbb{E}\left[\exp\left(\frac{|X| - |X'|}{t}\right)\right] && \text{(by Jensen's inequality)} \\ &= \mathbb{E}\left[\exp\left(\frac{\zeta(|X| - |X'|)}{t}\right)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\exp\left(\frac{\zeta(|X| - |X'|)}{t}\right) \middle| X, X'\right]\right] \\ &\stackrel{(*)}{\leq} \mathbb{E}\left[\mathbb{E}\left[\exp\left(\frac{\zeta|X - X'|}{t}\right) \middle| X, X'\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\exp\left(\frac{\zeta(X - X')}{t}\right) \middle| X, X'\right]\right] \\ &= \mathbb{E}\left[\exp\left(\frac{\zeta(X - X')}{t}\right)\right] \\ &= \frac{1}{2}\mathbb{E}\left[\exp\left(\frac{X - X'}{t}\right)\right] + \frac{1}{2}\mathbb{E}\left[\exp\left(\frac{X' - X}{t}\right)\right] \\ &= \mathbb{E}\left[\exp\left(\frac{X - X'}{t}\right)\right] \\ &= \mathbb{E}\left[\exp\left(\frac{X}{t}\right)\right] \cdot \mathbb{E}\left[\exp\left(-\frac{X}{t}\right)\right], \end{aligned}$$

where (*) follows from $||a| - |b|| \leq |a - b|$ and the fact that the function $z \mapsto \frac{1}{2}(\exp(-z) + \exp(z))$ is increasing in z when $z > 0$.

Therefore, from the above inequalities, we can deduce that

$$\mathbb{E}\left[\exp\left(\frac{|X|}{t}\right)\right] \leq \mathbb{E}\left[\exp\left(\frac{X}{t}\right)\right] \cdot \mathbb{E}\left[\exp\left(-\frac{X}{t}\right)\right] \cdot \exp\left(\frac{4\alpha}{t}\right).$$

Now, by setting $t := 8\alpha$ and using Jensen's inequality, we get

$$\begin{aligned} \mathbb{E}\left[\exp\left(\frac{|X|}{t}\right)\right] &\leq \mathbb{E}\left[\exp\left(\frac{X}{\alpha}\right)\right]^{\frac{1}{8}} \cdot \mathbb{E}\left[\exp\left(-\frac{X}{\alpha}\right)\right]^{\frac{1}{8}} \cdot \exp\left(\frac{1}{2}\right) \\ &\leq \exp\left(\frac{\log 2 + 2}{4}\right) \end{aligned}$$

$$\leq 2.$$

This implies that $\|X\|_{\psi_1} \leq t = 8\alpha$. □

Lemma A.4 *For any subexponential random variable X , we have*

$$\|\mathbb{E}[X]\|_{\psi_1} \leq \|X\|_{\psi_1}.$$

Proof.

$$\begin{aligned} \|\mathbb{E}[X]\|_{\psi_1} &= \inf \left\{ t > 0 : \exp \left(\frac{|\mathbb{E}[X]|}{t} \right) \leq 2 \right\} \\ &= \frac{|\mathbb{E}[X]|}{\log 2} \\ &\leq \frac{\mathbb{E}[|X|]}{\log 2} \\ &= \frac{\|X\|_{\psi_1} \cdot \log \exp \mathbb{E} \left[\frac{|X|}{\|X\|_{\psi_1}} \right]}{\log 2} \\ &\leq \frac{\|X\|_{\psi_1} \cdot \log \mathbb{E} \left[\exp \frac{|X|}{\|X\|_{\psi_1}} \right]}{\log 2} \\ &\leq \frac{\|X\|_{\psi_1} \cdot \log 2}{\log 2} \\ &= \|X\|_{\psi_1}. \end{aligned}$$

□

Appendix B. Trace regression proofs, adapted from (Klopp, 2014)

B.1 Proof of Theorem 3.1

Proof. First, it follows from Eq. (3.2) that

$$\frac{1}{n} \|Y - \mathfrak{X}(\widehat{\mathbf{B}})\|_2^2 + \lambda \|\widehat{\mathbf{B}}\|_* \leq \frac{1}{n} \|Y - \mathfrak{X}(\mathbf{B}^*)\|_2^2 + \lambda \|\mathbf{B}^*\|_*.$$

By substituting Y for $\mathfrak{X}(\mathbf{B}^*) + E$ and doing some algebra, we have

$$\frac{1}{n} \|\mathfrak{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_2^2 + 2\langle \Sigma, \mathbf{B}^* - \widehat{\mathbf{B}} \rangle + \lambda \|\widehat{\mathbf{B}}\|_* \leq \lambda \|\mathbf{B}^*\|_*.$$

Then, using the duality between the operator norm and the trace norm, we get

$$\frac{1}{n} \|\mathfrak{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_2^2 + \lambda \|\widehat{\mathbf{B}}\|_* \leq 2\|\Sigma\|_{\text{op}} \cdot \|\mathbf{B}^* - \widehat{\mathbf{B}}\|_* + \lambda \|\mathbf{B}^*\|_*. \quad (\text{B.1})$$

For a given set of vectors S , we denote by \mathbf{P}_S the orthogonal projection on the linear subspace spanned by elements of S (i.e., $\mathbf{P}_S = \sum_{i=1}^k u_i u_i^\top$ if $\{u_1, \dots, u_k\}$ is an orthonormal basis for S). For matrix $\mathbf{B} \in \mathbb{R}^{d_r \times d_c}$, let $S_r(\mathbf{B})$ and $S_c(\mathbf{B})$ be the linear subspace spanned

by the left and right orthonormal singular vectors of \mathbf{B} , respectively. Then, for $\mathbf{A} \in \mathbb{R}^{d_r \times d_c}$ define

$$\mathbf{P}_{\mathbf{B}}(\mathbf{A}) := \mathbf{A} - \mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A}) \quad \text{and} \quad \mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A}) := \mathbf{P}_{S_r^{\perp}(\mathbf{B})} \mathbf{A} \mathbf{P}_{S_c^{\perp}(\mathbf{B})}.$$

We can alternatively express $\mathbf{P}_{\mathbf{B}}(\mathbf{A})$ as

$$\begin{aligned} \mathbf{P}_{\mathbf{B}}(\mathbf{A}) &= \mathbf{A} - \mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A}) \\ &= \mathbf{P}_{S_r(\mathbf{B})} \mathbf{A} + \mathbf{P}_{S_r^{\perp}(\mathbf{B})} \mathbf{A} - \mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A}) \\ &= \mathbf{P}_{S_r(\mathbf{B})} \mathbf{A} + \mathbf{P}_{S_r^{\perp}(\mathbf{B})} [\mathbf{A} - \mathbf{A} \mathbf{P}_{S_c^{\perp}(\mathbf{B})}] \\ &= \mathbf{P}_{S_r(\mathbf{B})} \mathbf{A} + \mathbf{P}_{S_r^{\perp}(\mathbf{B})} \mathbf{A} \mathbf{P}_{S_c(\mathbf{B})}. \end{aligned} \tag{B.2}$$

In particular, since $S_r(\mathbf{B})$ and $S_c(\mathbf{B})$ both have dimension $\text{rank}(\mathbf{B})$, it follows from Eq. (B.2) that

$$\text{rank}(\mathbf{P}_{\mathbf{B}}(\mathbf{A})) \leq 2 \text{rank}(\mathbf{B}). \tag{B.3}$$

Moreover, the definition of $\mathbf{P}_{\mathbf{B}}^{\perp}$ implies that the left and right singular vectors of $\mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A})$ are orthogonal to those of \mathbf{B} . We thus have

$$\|\mathbf{B} + \mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A})\|_* = \|\mathbf{B}\|_* + \|\mathbf{P}_{\mathbf{B}}^{\perp}(\mathbf{A})\|_*.$$

Set $\mathbf{B} := \mathbf{B}^*$ and $\mathbf{A} := \widehat{\mathbf{B}} - \mathbf{B}^*$; then the above equality entails

$$\|\mathbf{B}^* + \mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* = \|\mathbf{B}^*\|_* + \|\mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*. \tag{B.4}$$

We can then use the latter to get the following inequality:

$$\begin{aligned} \|\widehat{\mathbf{B}}\|_* &= \|\mathbf{B}^* + \widehat{\mathbf{B}} - \mathbf{B}^*\|_* \\ &= \|\mathbf{B}^* + \mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*) + \mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \\ &\geq \|\mathbf{B}^* + \mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* - \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \\ &= \|\mathbf{B}^*\|_* + \|\mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* - \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*. \end{aligned} \tag{B.5}$$

Combining Eq. (B.1) with Eq. (B.5), we get

$$\begin{aligned} \frac{1}{n} \|\mathfrak{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_2^2 &\leq 2 \|\Sigma\|_{\text{op}} \cdot \|\mathbf{B}^* - \widehat{\mathbf{B}}\|_* + \lambda \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* - \lambda \|\mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \\ &\leq (2 \|\Sigma\|_{\text{op}} + \lambda) \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* + (2 \|\Sigma\|_{\text{op}} - \lambda) \|\mathbf{P}_{\mathbf{B}^*}^{\perp}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \\ &\leq \frac{5}{3} \lambda \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*, \end{aligned}$$

where, in the last inequality, we have used Eq. (3.5). Now, using this and the fact that $\text{rank}(\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)) \leq 2 \text{rank}(\mathbf{B}^*)$ (by Eq. (B.3)), we can apply the Cauchy–Schwartz inequality to singular values of $\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)$ to obtain

$$\begin{aligned} \frac{1}{n} \|\mathfrak{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_2^2 &\leq \frac{5}{3} \lambda \sqrt{2 \text{rank}(\mathbf{B}^*)} \|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_{\text{F}} \\ &\leq \frac{5}{3} \lambda \sqrt{2 \text{rank}(\mathbf{B}^*)} \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}. \end{aligned} \tag{B.6}$$

The next lemma makes a connection between $\widehat{\mathbf{B}}$ and the constraint set $\mathcal{C}(\nu, \eta)$.

Lemma B.1 *If $\lambda \geq 3\|\Sigma\|_{\text{op}}$, then*

$$\|\mathbf{P}_{\mathbf{B}^*}^\perp(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \leq 5\|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*.$$

Proof. Note that $\|\mathfrak{X}(\cdot)\|_2^2$ is a convex function. We can then use the convexity at \mathbf{B}^* to get

$$\begin{aligned} \frac{1}{n}\|\mathfrak{X}(\widehat{\mathbf{B}})\|_2^2 - \frac{1}{n}\|\mathfrak{X}(\mathbf{B}^*)\|_2^2 &\geq -\frac{2}{n}\sum_{i=1}^n (y_i - \langle \mathbf{X}_i, \mathbf{B}^* \rangle) \langle \mathbf{X}_i, \widehat{\mathbf{B}} - \mathbf{B}^* \rangle \\ &= -2\langle \Sigma, \widehat{\mathbf{B}} - \mathbf{B}^* \rangle \\ &\geq -2\|\Sigma\|_{\text{op}}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_* \\ &\geq -\frac{2\lambda}{3}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_*. \end{aligned}$$

Combining this with Eq. (3.2) and Eq. (B.5), we have

$$\begin{aligned} \frac{2\lambda}{3}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_* &\geq \frac{1}{n}\|\mathfrak{X}(\mathbf{B}^*)\|_2^2 - \frac{1}{n}\|\mathfrak{X}(\widehat{\mathbf{B}})\|_2^2 \\ &\geq \lambda\|\widehat{\mathbf{B}}\|_* - \lambda\|\mathbf{B}^*\|_* \\ &\geq \lambda\|\mathbf{P}_{\mathbf{B}^*}^\perp(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* - \lambda\|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*. \end{aligned}$$

By the triangle inequality, we have

$$\|\mathbf{P}_{\mathbf{B}^*}^\perp(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \leq 5\|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_*.$$

This finishes proof of Lemma B.1. □

Lemma B.1, the triangle inequality, and Eq. (B.3) imply that

$$\begin{aligned} \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_* &\leq 6\|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_* \\ &\leq \sqrt{72\text{rank}(\mathbf{B}^*)}\|\mathbf{P}_{\mathbf{B}^*}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_{\text{F}} \\ &\leq \sqrt{72\text{rank}(\mathbf{B}^*)}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}. \end{aligned}$$

Next, define $b := \mathfrak{N}(\widehat{\mathbf{B}} - \mathbf{B}^*)$ and $\mathbf{A} := \frac{1}{b}(\widehat{\mathbf{B}} - \mathbf{B}^*)$. We then have that

$$\mathfrak{N}(\mathbf{A}) = 1 \quad \text{and} \quad \|\mathbf{A}\|_* \leq \sqrt{72\text{rank}(\mathbf{B}^*)}\|\mathbf{A}\|_{\text{F}}.$$

In order to finish the proof, we consider the following two cases:

Case 1: If $\|\mathbf{A}\|_{\text{F}}^2 < \nu$, then

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 < 4b^2\nu.$$

Case 2: Otherwise, $\mathbf{A} \in \mathcal{C}(\nu, \eta)$.

We can now use the RSC condition and Eq. (B.6) to get

$$\alpha \frac{\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2}{b^2} - \beta \leq \frac{\|\mathfrak{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_2^2}{nb^2},$$

which leads to

$$\begin{aligned} \alpha \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 - 4b^{*2}\beta &\leq \frac{\|\mathfrak{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_2^2}{n} \\ &\leq \frac{5\lambda\sqrt{2\text{rank}(\mathbf{B}^*)}}{3} \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}} \\ &\leq \frac{50\lambda^2\text{rank}(\mathbf{B}^*)}{3\alpha} + \frac{\alpha}{2} \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2. \end{aligned}$$

Therefore, we have

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}^2 \leq \frac{100\lambda^2\text{rank}(\mathbf{B}^*)}{3\alpha^2} + \frac{8b^{*2}\beta}{\alpha},$$

which completes the proof of Theorem 3.1. \square

B.2 Matrix Bernstein inequality

The next proposition is a variant of the Bernstein inequality (Proposition 11 of Klopp (2014)).

Proposition B.1 *Let $(\mathbf{Z}_i)_{i=1}^n$ be a sequence of $d_r \times d_c$ independent random matrices with zero mean, such that*

$$\mathbb{E}[\exp\left(\frac{\|\mathbf{Z}_i\|_{\text{op}}}{\delta}\right)] \leq e \quad \forall i \in [n],$$

and

$$\sigma_{\mathbf{Z}} = \max \left\{ \left\| \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbf{Z}_i \mathbf{Z}_i^{\top}] \right\|_{\text{op}}, \left\| \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbf{Z}_i^{\top} \mathbf{Z}_i] \right\|_{\text{op}} \right\}^{\frac{1}{2}},$$

for some positive values δ and $\sigma_{\mathbf{Z}}$. Then, there exists a numerical constant $C > 0$ such that, for all $t > 0$,

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \right\|_{\text{op}} \leq C \max \left\{ \sigma_{\mathbf{Z}} \sqrt{\frac{t + \log(d)}{n}}, \delta \left(\log \frac{\delta}{\sigma_{\mathbf{Z}}} \right) \frac{t + \log(d)}{n} \right\}, \quad (\text{B.7})$$

with probability at least $1 - \exp(-t)$, where $d = d_r + d_c$.

We also state the following corollary of the matrix Bernstein inequality.

Corollary B.1 *If Eq. (B.7) holds and $n \geq \frac{\delta^2}{C\sigma_{\mathbf{Z}}^2} \log d \left(\log \frac{\delta}{\sigma_{\mathbf{Z}}} \right)^2$, then*

$$\mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \right\|_{\text{op}} \right] \leq C' \sigma_{\mathbf{Z}} \sqrt{\frac{2e \log d}{n}},$$

where $C' > 0$ is a numerical constant.

This corollary has been proved for the case of $\mathbf{Z} = \zeta_i \mathbf{X}_i$ in Klopp (2014). The proof can be adapted to the general case as well.

B.3 Proof of Theorem 3.2

Proof. First, we reproduce a slightly modified version of proof Lemma 12 in (Klopp, 2014), adapted to our setting. Set

$$\beta := \frac{93\eta\mathfrak{c}^2}{\gamma_{\min}} \mathbb{E}[\|\Sigma_R\|_{\text{op}}]^2.$$

By \mathcal{B} , we denote a bad event defined as

$$\mathcal{B} := \left\{ \exists \mathbf{A} \in \mathcal{C}'(\theta, \eta) \text{ such that } \frac{1}{2}\|\mathbf{A}\|_{L^2(\Pi)}^2 - \frac{1}{n}\|\mathfrak{X}(\mathbf{A})\|_2^2 > \frac{1}{4}\|\mathbf{A}\|_{L^2(\Pi)}^2 + \beta \right\}.$$

We thus need to bound the probability of this event. Set $\xi = 6/5$. Then, for $T > 0$, we define

$$\mathcal{C}'(\theta, \eta, T) := \{\mathbf{A} \in \mathcal{C}'(\theta, \eta) T \leq \|\mathbf{A}\|_{L^2(\Pi)}^2 < \xi T\}.$$

Clearly, we have

$$\mathcal{C}'(\theta, \theta) = \bigcup_{l=1}^{\infty} \mathcal{C}'(\theta, \eta, \xi^{l-1}\theta).$$

Now, if the event \mathcal{B} holds for some $\mathbf{A} \in \mathcal{C}'(\theta, \eta)$, then $\mathbf{A} \in \mathcal{C}'(\theta, \eta, \xi^{l-1}\theta)$ for some $l \in \mathbb{N}$. In this case, we have

$$\begin{aligned} \frac{1}{2}\|\mathbf{A}\|_{L^2(\Pi)}^2 - \frac{1}{n}\|\mathfrak{X}(\mathbf{A})\|_2^2 &> \frac{1}{4}\|\mathbf{A}\|_{L^2(\Pi)}^2 + \beta \\ &\geq \frac{1}{4}\xi^{l-1}\theta + \beta \\ &= \frac{5}{24}\xi^l\theta + \beta. \end{aligned}$$

Next, we define the event \mathcal{B}_l as

$$\mathcal{B}_l := \left\{ \exists \mathbf{A} \in \mathcal{C}'(\theta, \eta, \xi^{l-1}\theta) \text{ such that } \frac{1}{2}\|\mathbf{A}\|_{L^2(\Pi)}^2 - \frac{1}{n}\|\mathfrak{X}(\mathbf{A})\|_2^2 > \frac{5}{24}\xi^l\theta + \beta \right\}.$$

It follows that

$$\mathcal{B} \subseteq \bigcup_{l=1}^{\infty} \mathcal{B}_l.$$

The following lemma helps us obtain an upper bound for the probability of each of these events \mathcal{B}_l .

Lemma B.2 *Define*

$$Z_T := \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left\{ \frac{1}{2}\|\mathbf{A}\|_{L^2(\Pi)}^2 - \frac{1}{n}\|\mathfrak{X}(\mathbf{A})\|_2^2 \right\}.$$

Then, assuming that $(\mathbf{X}_i)_{i=1}^n$ are i.i.d. samples drawn from Π , we get

$$\mathbb{P}(Z_T \geq \frac{5\xi T}{24} + \beta) \leq \exp\left(-\frac{Cn\xi T}{\mathfrak{c}^2}\right), \quad (\text{B.8})$$

for some numerical constant $C > 0$.

Proof. The proof is similar to the proof of Lemma 14 in (Klopp, 2014). For a $d_r \times d_c$ matrix \mathbf{A} , define

$$f(\mathbf{X}; \mathbf{A}) := \langle \mathbf{X}, \mathbf{A} \rangle^2 \cdot \mathbb{I}(|\langle \mathbf{X}, \mathbf{A} \rangle| \leq \mathfrak{c}).$$

Next, denote

$$W_T := \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \frac{1}{n} \sum_{i=1}^n \{\mathbb{E}[f(\mathbf{X}_i; \mathbf{A})] - f(\mathbf{X}_i; \mathbf{A})\},$$

$$\widetilde{W}_T := \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left| \frac{1}{n} \sum_{i=1}^n \{\mathbb{E}[f(\mathbf{X}_i; \mathbf{A})] - f(\mathbf{X}_i; \mathbf{A})\} \right|.$$

It follows from Assumption 3.3 (where \mathfrak{c} is defined) that $Z_T \leq W_T$, and clearly $W_T \leq \widetilde{W}_T$. Therefore, we have

$$\mathbb{P}(Z_T \geq t) \leq \mathbb{P}(\widetilde{W}_T \geq t)$$

for all t . Therefore, if we prove that Eq. (B.8) holds when Z_T is replaced with \widetilde{W}_T , we will be done. In the sequel, we will aim to prove this using Massart's inequality (e.g., Theorem 3 of Massart (2000)). In order to invoke Massart's inequality, we need bounds for $\mathbb{E}[\widetilde{W}_T]$ and $\text{Var}(\widetilde{W}_T)$.

First, we find an upper bound for $\mathbb{E}[\widetilde{W}_T]$. It follows from the symmetrization argument (e.g., Lemma 6.3 of Ledoux and Talagrand (1991)) that

$$\mathbb{E}[\widetilde{W}_T] \leq 2 \mathbb{E} \left[\sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left| \frac{1}{n} \sum_{i=1}^n \zeta_i f(\mathbf{X}_i; \mathbf{A}) \right| \right], \quad (\text{B.9})$$

where $(\zeta_i)_{i=1}^n$ is a sequence of i.i.d. Radamacher random variables. Note that Lemma 6.3 of Ledoux and Talagrand (1991) requires the use of a convex function and a norm. Here, the convex function is the identity function and the norm is the infinity norm applied to an infinite-dimensional vector (indexed by $\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)$).

Next, we will use the contraction inequality (e.g., Theorem 4.4 of Ledoux and Talagrand (1991)). First, we write $f(\mathbf{X}_i; \mathbf{A}) = \alpha_i \langle \mathbf{X}_i, \mathbf{A} \rangle$, where $\alpha_i = \langle \mathbf{X}_i, \mathbf{A} \rangle \cdot \mathbb{I}(|\langle \mathbf{X}_i, \mathbf{A} \rangle| \leq \mathfrak{c})$. By definition, $|\alpha_i| \leq \mathfrak{c}$. Now, for every realization of the random variables $\mathbf{X}_1, \dots, \mathbf{X}_n$ we can apply Theorem 4.4 in Ledoux and Talagrand (1991) to obtain

$$\mathbb{E}[\zeta] \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left| \frac{1}{n} \sum_{i=1}^n \zeta_i f(\mathbf{X}_i; \mathbf{A}) \right| \leq \mathfrak{c} \mathbb{E}[\zeta] \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left| \frac{1}{n} \sum_{i=1}^n \zeta_i \langle \mathbf{X}_i, \mathbf{A} \rangle \right|.$$

Now, taking expectation of both sides with respect to \mathbf{X}_i 's, using the tower property, and combining with Eq. (B.9), we obtain

$$\begin{aligned} \mathbb{E}[\widetilde{W}_T] &\leq 8\mathfrak{c} \mathbb{E} \left[\sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \left| \frac{1}{n} \sum_{i=1}^n \zeta_i \langle \mathbf{X}_i, \mathbf{A} \rangle \right| \right] \\ &\leq 8\mathfrak{c} \mathbb{E} \left[\|\Sigma_R\|_{\text{op}} \sup_{\mathbf{A} \in \mathcal{C}'(\theta, \eta, T)} \|\mathbf{A}\|_* \right] \end{aligned}$$

$$\begin{aligned}
 &\leq 8\mathbf{c}\sqrt{\eta}\mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}} \sup_{\mathbf{A}\in\mathcal{C}'(\theta,\eta,T)} \|\mathbf{A}\|_{\text{F}}] \\
 &\leq 8\mathbf{c}\sqrt{\frac{\eta}{\gamma_{\min}}}\mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}} \sup_{\mathbf{A}\in\mathcal{C}'(\theta,\eta,T)} \|\mathbf{A}\|_{L^2(\Pi)}] \\
 &\leq 8\mathbf{c}\sqrt{\frac{\eta\xi T}{\gamma_{\min}}}\mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}}].
 \end{aligned}$$

In the above, we also used the definition of $\mathcal{C}'(\theta, \eta, T)$ as well as Assumption 3.2. We can now use $2ab \leq a^2 + b^2$ to get

$$\mathbb{E}[\widetilde{W}_T] \leq \frac{8}{9} \left(\frac{5\xi T}{24} \right) + \frac{87\eta\mathbf{c}^2}{\gamma_{\min}} \mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}}]^2.$$

Next, we turn to find an upper bound for a certain variance term, specifically, the term $\sum_{i=1}^n \text{Var}(f(\mathbf{X}_i; \mathbf{A})/n)$. Since

$$\begin{aligned}
 \text{Var}(f(\mathbf{X}_i; \mathbf{A})) &\leq \mathbb{E}[f(\mathbf{X}_i; \mathbf{A})^2] \\
 &\leq \mathbf{c}^2 \cdot \mathbb{E}[\langle \mathbf{X}_i, \mathbf{A} \rangle^2] \\
 &= \mathbf{c}^2 \cdot \|\mathbf{A}\|_{L^2(\Pi)}^2,
 \end{aligned}$$

it follows that

$$\begin{aligned}
 \sup_{\mathbf{A}\in\mathcal{C}'(\theta,\eta,T)} \sum_{i\in[n]} \frac{1}{n^2} \text{Var}(f(\mathbf{X}_i; \mathbf{A})) &\leq \frac{\mathbf{c}^2}{n} \cdot \sup_{\mathbf{A}\in\mathcal{C}'(\theta,\eta,T)} \|\mathbf{A}\|_{L^2(\Pi)}^2 \\
 &\leq \frac{\xi T \mathbf{c}^2}{n}.
 \end{aligned}$$

Finally, noting that $\frac{1}{n}f(\mathbf{X}_i; \mathbf{A}) \leq \frac{1}{n}\mathbf{c}^2$ almost surely, we can use Massart's inequality (e.g., Theorem 3 of Massart (2000)) to conclude that

$$\begin{aligned}
 \mathbb{P}(\widetilde{W}_T \geq \frac{5\xi T}{24} + \beta) &= \mathbb{P}(\widetilde{W}_T \geq \frac{5\xi T}{24} + \frac{93\eta\mathbf{c}^2}{\gamma_{\min}} \mathbb{E}[\|\boldsymbol{\Sigma}_R\|_{\text{op}}]^2) \\
 &\leq \mathbb{P}(\widetilde{W}_T \geq \frac{18}{17} \mathbb{E}[\widetilde{W}_T] + \frac{1}{17} \left(\frac{5\xi T}{24} \right)) \\
 &\leq \exp\left(-\frac{Cn\xi T}{\mathbf{c}^2}\right)
 \end{aligned}$$

for some numerical constant $C_0 > 0$. □

Lemma B.2 entails that

$$\begin{aligned}
 \mathbb{P}(\mathcal{B}_l) &\leq \exp\left(-\frac{C_0 n \xi^l \theta}{\mathbf{c}^2}\right) \\
 &\leq \exp\left(-\frac{C_0 n l \log(\xi) \theta}{\mathbf{c}^2}\right).
 \end{aligned}$$

Therefore, by setting the numerical constant $C > 0$ appropriately, the union bound implies that

$$\begin{aligned} \mathbb{P}(\mathcal{B}) &\leq \sum_{l=1}^{\infty} \mathbb{P}(\mathcal{B}_l) \\ &\leq \sum_{l=1}^{\infty} \exp\left(-\frac{Cnl\theta}{\mathfrak{c}^2}\right) \\ &= \frac{\exp\left(-\frac{Cn\theta}{\mathfrak{c}^2}\right)}{1 - \exp\left(-\frac{Cn\theta}{\mathfrak{c}^2}\right)}. \end{aligned}$$

Finally, assuming that $Cn\theta > \mathfrak{c}^2$, we get that

$$\mathbb{P}(\mathcal{B}) \leq 2 \exp\left(-\frac{Cn\theta}{\mathfrak{c}^2}\right),$$

which completes the proof of Theorem 3.2. \square

Appendix C. Additional simulations

In addition to the four estimators described in Section 4.1, in this appendix we consider the following variants of the cross-validated estimator.

1. **cv 1 fold** follows the same procedure as the **cv** estimator, except that it uses only a single training and validation split. Specifically, this estimator is equal to $\widehat{\mathbf{B}}_{-k}(\widehat{\lambda}'_{\text{cv}})$ for a random k , $k \in [K]$, where

$$\widehat{\lambda}'_{\text{cv}} \in \arg \min_{\lambda \in \Lambda} \|Y_k - \mathfrak{X}_k(\widehat{\mathbf{B}}_{-k}(\lambda))\|_2^2.$$

The motivation to consider this estimator goes back to Remark 4.3 where we discussed that that **cv 1 fold** has the same theoretical properties as **cv**, but one expects the latter to benefit from averaging and to perform better empirically.


2. **cv refit** follows the same procedure as the **cv** estimator to choose the optimal penalty $\widehat{\lambda}_{\text{cv}}$. However, it uses all of the data and refits the estimator. Specifically, **cv refit** is a solution to Eq. (3.2) when all of the observations $\{(\mathbf{X}_i, y_i)\}_{i=1}^n$ and $\widehat{\lambda}_{\text{cv}}$ are used. While **cv refit** is not covered by our theory (by Remark 4.3), it is a common way to perform cross-validation in practice.
3. **cv overfit** is constructed as follows. For a sequence of Λ , as in **cv** the estimator $\widehat{\mathbf{B}}(\lambda)$ is calculated using the entire observations $\{(\mathbf{X}_i, y_i)\}_{i=1}^n$. Next, the penalty $\widehat{\lambda}''_{\text{cv}}$ is selected such that the *in-sample* error on the same set of observations is minimized. Thus,

$$\widehat{\lambda}''_{\text{cv}} \in \arg \min_{\lambda \in \Lambda} \|Y - \mathfrak{X}(\widehat{\mathbf{B}}(\lambda))\|_2^2.$$


The **cv overfit** estimator $\widehat{\mathbf{B}}(\widehat{\lambda}''_{\text{cv}})$ uses the same observations to fit $\widehat{\mathbf{B}}$ and calculate the error, which is technically overfitted to the observations. The reason to include this estimator is to have a baseline for the effect of overfitting.

We then repeat the simulations of Section 4.1 for all (seven) estimators, for three different noise levels, when the standard deviation of the noise belongs to $\{0.5, 1, 2\}$, and the results are presented in Figures 3 to 6. The results are generally aligned with those in Figure 1 and Figure 2. However, a few interesting observations about various cross-validated estimators are worth highlighting:


- `cv refit` performs very well in general and nearly ties with `oracle`. But for small values of n , `cv` slightly outperforms `cv refit`, and this effect is more visible in the high-noise regime.
- `cv 1 fold` performs well compared to the theory inspired estimators, but generally loses to `cv`, which underscores the benefit of averaging. It is clear from Figures 5 to 6 that `cv 1 fold` and `cv` both select a penalty close to the one `oracle` selects, but the selection has higher variance for `cv 1 fold`, which could explain its empirical inferiority to `cv`.
- As expected, the penalty parameter selected by `cv overfit` is close to 0 due to overfitting. But the relative error of `cv overfit` demonstrates a more nuanced behavior. The estimator performs very poorly in the high-noise regime (its relative error is above 1 and the curve is not visible in the plots), but in the low-to-medium noise regime, it performs well when n is not too large. There are two explanations for this behavior. First, when n is large, the impact of overfitting is increased as the estimator is perfectly fitting a larger number of noisy observations. Second, in the low-to-medium noise regime, the counterintuitively good performance of `cv overfit` can be explained by the exact recovery results that are theoretically correct when the noise is zero. In these settings, the optimum penalty is indeed 0; hence one can expect that, at λ equal to 0, the relative error only gradually deteriorates when the noise is increased.




(a) MovieLens, $d = 50$, $\sigma = 0.5$




(b) MovieLens, $d = 100$, $\sigma = 0.5$




(c) MovieLens, $d = 50$, $\sigma = 1$



(d) MovieLens, $d = 100$, $\sigma = 1$




(e) MovieLens, $d = 50$, $\sigma = 2$




(f) Synthetic, $d = 100$, $\sigma = 2$


Figure 6: Comparison of the selected penalty parameter λ for the proposed estimators on movielens data, for varying noise level and cross-validation methodology.




(a) Synthetic, $(d, r) = (50, 2)$, $\sigma = 0.5$



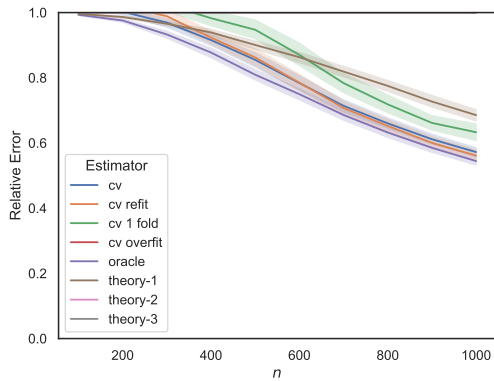
(b) Synthetic, $(d, r) = (100, 3)$, $\sigma = 0.5$




(c) Synthetic, $(d, r) = (50, 2)$, $\sigma = 1$



(d) Synthetic, $(d, r) = (100, 3)$, $\sigma = 1$

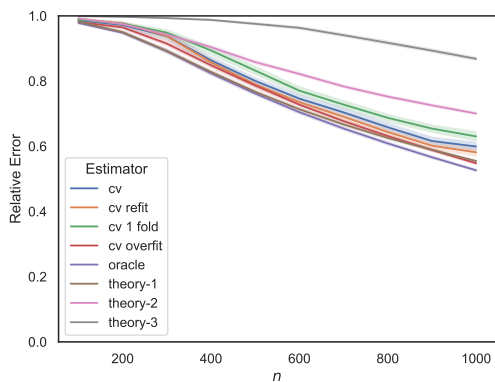


(e) Synthetic, $(d, r) = (50, 2)$, $\sigma = 2$




(f) Synthetic, $(d, r) = (100, 3)$, $\sigma = 2$


Figure 3: Comparison of the relative error (i.e., $\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 / \|\mathbf{B}^*\|_F^2$) for the proposed estimators on synthetic data, for varying noise level and cross-validation methodology.




(a) MovieLens, $d = 50$, $\sigma = 0.5$




(b) MovieLens, $d = 100$, $\sigma = 0.5$




(c) MovieLens, $d = 50$, $\sigma = 1$



(d) MovieLens, $d = 100$, $\sigma = 1$




(e) MovieLens, $d = 50$, $\sigma = 2$




(f) MovieLens, $d = 100$, $\sigma = 2$


Figure 4: Comparison of the relative error (i.e., $\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 / \|\mathbf{B}^*\|_F^2$) for the proposed estimators on movielens data, for varying noise level and cross-validation methodology.




(a) Synthetic, $(d, r) = (50, 2)$, $\sigma = 0.5$




(b) Synthetic, $(d, r) = (100, 3)$, $\sigma = 0.5$




(c) Synthetic, $(d, r) = (50, 2)$, $\sigma = 1$



(d) Synthetic, $(d, r) = (100, 3)$, $\sigma = 1$



(e) Synthetic, $(d, r) = (50, 2)$, $\sigma = 2$



(f) Synthetic, $(d, r) = (100, 3)$, $\sigma = 2$

Figure 5: Comparison of the selected penalty parameter λ for the proposed estimators on synthetic data, for varying noise level and cross-validation methodology

References

- Karim Abou-Moustafa and Csaba Szepesvari. An a Priori Exponential Tail Bound for k-Folds Cross-Validation. *arXiv e-prints*, art. arXiv:1706.05801, Jun 2017.
- T. Tony Cai and Anru Zhang. Rop: Matrix recovery via rank-one projections. *Ann. Statist.*, 43(1):102–138, 02 2015.
- E. J. Candes and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inf. Theor.*, 57(4):2342–2359, April 2011. ISSN 0018-9448.
- Emmanuel Candes and Benjamin Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2009.
- Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- Rich Caruana. Multitask learning. *Machine Learning*, 28(1):41–75, 1997.
- Yudong Chen, Srinadh Bhojanapalli, Sujay Sanghavi, and Rachel Ward. Completing any low-rank matrix, provably. *Journal of Machine Learning Research*, 16(94):2999–3034, 2015.
- Denis Chetverikov, Zhipeng Liao, and Victor Chernozhukov. On cross-validated Lasso. *arXiv e-prints*, art. arXiv:1605.02214, May 2016.
- M. A. Davenport and J. Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- Rina Foygel, Ruslan Salakhutdinov, Ohad Shamir, and Nathan Srebro. Learning with the weighted trace-norm under arbitrary sampling distributions. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, page 2133–2141, Red Hook, NY, USA, 2011. Curran Associates Inc.
- David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC, 2015. ISBN 1498712169.
- Satyen Kale, Ravi Kumar, and Sergei Vassilvitskii. Cross-validation and mean-square stability. In *Proceedings of 2nd Symposium on Innovations in Computer Science (ICS)*, 2011.
- Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, 2009.

- Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11(69):2057–2078, 2010.
- Olga Klopp. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.
- Vladimir Koltchinskii, Karim Lounici, and Alexandre B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302 – 2329, 2011a.
- Vladimir Koltchinskii et al. Von neumann entropy penalization and low-rank matrix estimation. *The Annals of Statistics*, 39(6):2936–2973, 2011b.
- Ravi Kumar, Daniel Lokshtanov, Sergei Vassilvitskii, and Andrea Vattani. Near-optimal bounds for cross-validation via loss stability. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 27–35. PMLR, 2013.
- M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. A Series of Modern Surveys in Mathematics Series. Springer, 1991. ISBN 9783540520139.
- Pascal Massart. About the constants in Talagrand’s concentration inequalities for empirical processes. *The Annals of Probability*, 28(2):863 – 884, 2000.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11(80):2287–2322, 2010.
- Sahand Negahban and Martin J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, 39(2):1069 – 1097, 2011. doi: 10.1214/10-AOS850.
- Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13:1665–1697, 2012.
- Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.
- Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- Angelika Rohde and Alexandre B. Tsybakov. Estimation of high-dimensional low-rank matrices. *The Annals of Statistics*, 39(2):887 – 930, 2011.
- Nathan Srebro and Ruslan R Salakhutdinov. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2056–2064. Curran Associates, Inc., 2010.

Nathan Srebro, Noga Alon, and Tommi S. Jaakkola. Generalization error bounds for collaborative prediction with low-rank matrices. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1321–1328. 2005.

Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.