ESTIMATION OF HIGH-DIMENSIONAL LOW RANK MATRICES

By Angelika Rohde and Alexandre B. Tsybakov

Universität Hamburg and CREST and Université Paris VI

Suppose that we observe entries or, more generally, linear combinations of entries of an unknown $m \times T$-matrix A corrupted by noise. We are particularly interested in the high-dimensional setting where the number mT of unknown entries can be much larger than the sample size N. Motivated by several applications, we consider estimation of matrix A under the assumption that it has small rank. This can be viewed as dimension reduction or sparsity assumption. In order to shrink towards a low-rank representation, we investigate penalized least squares estimators with a Schatten-p quasi-norm penalty term, $p \leq 1$. We study these estimators under two possible assumptions – a modified version of the restricted isometry condition and a uniform bound on the ratio "empirical norm induced by the sampling operator/Frobenius norm". The main results are stated as non-asymptotic upper bounds on the prediction risk and on the Schatten-q risk of the estimators, where $q \in [p, 2]$. The rates that we obtain for the prediction risk are of the form rm/N (for $m = T$), up to logarithmic factors, where r is the rank of A. The particular examples of multi-task learning and matrix completion are worked out in detail. The proofs are based on tools from the theory of empirical processes. As a by-product we derive bounds for the kth entropy numbers of the quasi-convex Schatten class embeddings $S_p^M \hookrightarrow S_2^M$, $p < 1$, which are of independent interest.

1. Introduction.

Consider the observations (X_i, Y_i) satisfying the model

$$Y_i = \operatorname{tr}\big(X_i' A^*\big) \; + \; \xi_i, \quad i = 1, ..., N, \tag{1}$$

where $X_i \in \mathbb{R}^{m \times T}$ are given matrices (m rows, T columns), $A^* \in \mathbb{R}^{m \times T}$ is an unknown matrix, $\operatorname{tr}(B)$ denotes the trace of square matrix B and ξ_i are i.i.d. random errors. Our aim is to estimate the matrix A^* and to predict the future Y-values based on the sample $(X_i, Y_i), i = 1, ..., N$.

We will call model (1) the *trace regression model*. Clearly, for $T = 1$ it reduces to the standard regression model. The "design" matrices X_i will be called *masks*. This name is motivated by the fact that we focus on the applications of trace regression where X_i are very sparse, i.e., contain only a small percentage of non-zero entries. Therefore, multiplication of A^* by X_i masks most of the entries of A^*. The following two examples are of particular interest.

(i) *Point masks.* For some, typically small, integer d the point masks X_i are defined

---

[*] Research partially supported by the grant ANR-06-BLAN-0194, ANR "Parcimonie" and by the PASCAL-2 Network of Excellence

AMS 2000 subject classifications: 62G05, 62F10

Keywords and phrases: High-dimensional low-rank matrices, empirical process, sparse recovery, Schatten norm, penalized least-squares estimator, quasi-convex Schatten class embeddings

as elements of the set

$$\mathcal{X}_d = \left\{ \sum_{i=1}^{d} e_{k_i}(m) e'_{l_i}(T) :\ 1 \le k_i \le m, 1 \le l_i \le T,\ \text{with } (k_i, l_i) \ne (k_{i'}, l_{i'}) \text{ for } i \ne i' \right\},$$

where $e_k(m)$ are the canonical basis vectors of $\mathbb{R}^m$. In particular, for $d = 1$ the point masks X_i are matrices that have only one non-zero entry, which equals to 1. The problem of estimation of A^* in this case becomes the problem of *matrix completion* or collaborative filtering; the observations Y_i are just some selected entries of A^* corrupted by noise. An important feature of the real-world matrix completion problems is that the number of observed entries is much smaller than the size of the matrix: $N \ll mT$, whereas mT can be very large. For example, mT is of the order of hundreds of millions for the Netflix problem.

(ii) *Column or row masks*. If X_i has only a small number d of non-zero columns or rows, it is called column or row mask, respectively. We suppose here that d is much smaller than m and T. A remarkable case $d = 1$ is covering the problem known in Statistics and Econometrics as longitudinal (or panel, or cross-section) data analysis and in Machine Learning as multi-task learning. In what follows we will designate this problem as multi-task learning, to avoid ambiguity. In the simplest version of multi-task learning, we have $N = nT$ where T is the number of tasks (for instance, in image detection each task t is associated with a particular type of visual object, e.g., face, car, chair, etc.), and n is the number of observations per task. The tasks are characterized by vectors of parameters $a_t^* \in \mathbb{R}^m$, $t = 1, \dots, T$, which constitute the columns of matrix A^*:

$$A^* = (a_1^* \cdots a_T^*).$$

The X_i are column masks, each containing only one non-zero column $\mathbf{x}^{(t,s)} \in \mathbb{R}^m$ (with the convention that $\mathbf{x}^{(t,s)}$ is the tth column):

$$X_i \in \{(0 \cdots 0 \underbrace{\mathbf{x}^{(t,s)}}_{t} 0 \cdots 0),\ t = 1, \dots, T,\ s = 1, \dots, n\}.$$

The column $\mathbf{x}^{(t,s)}$ is interpreted as the vector of predictor variables corresponding to sth observation for the tth task. Thus, for each $i = 1, \dots, N$ there exists a pair (t, s) with $t = 1, \dots, T,\ s = 1, \dots, n$, such that

$$\mathrm{tr}(X_i' A^*) = (a_t^*)' \mathbf{x}^{(t,s)}. \tag{2}$$

If we denote by $Y^{(t,s)}$ and $\xi^{(t,s)}$ the corresponding values Y_i and ξ_i, then the trace regression model (1) can be written as a collection of T standard regression models:

$$Y^{(t,s)} = (a_t^*)' \mathbf{x}^{(t,s)} + \xi^{(t,s)}, \quad t = 1, \dots, T,\ s = 1, \dots, n.$$

This is the usual formulation of the multi-task learning model in the literature.

For both examples given above the matrices X_i are sparse in the sense that they have only a small portion of non-zero entries. On the other hand, such a sparsity property is not necessarily granted for the target matrix A^*. Nevertheless, we can always characterize A^* by its rank $r = \text{rank}\,(A^*)$, and say that a matrix is sparse if it has small rank, cf. Recht et al.(2007). For example, the problem of estimation of a square matrix $A^* \in \mathbb{R}^{m\times m}$ is a parametric problem which is formally of dimension m^2 but it has only $(2m-r)r$ free parameters. If r is small as compared to m, then the intrinsic dimension of the problem is of the order rm. In other words, the rank sparsity assumption $r \ll m$ is a dimension reduction assumption. This assumption will be crucial for the interpretation of our results.

Estimation of high-dimensional matrices has been recently studied by several authors. To mention just a few, Meinshausen and Bühlmann (2006) investigated the Lasso in view of estimation of inverse covariance matrices and sparse high-dimensional graphical models. Bickel and Lewina (2008) considered the problem of estimating a covariance matrix of p variables from n observations by either banding or tapering the sample covariance matrix, or estimating a banded version of its inverse. Their estimates are shown to be consistent in the operator norm with the explicit rates when $(\log p)/n$ tends to zero. Ravikumar et al. (2008) studied the problem of estimating both the covariance matrix and the inverse covariance matrix by minimizing an ℓ_1-penalized log-determinant Bregman divergence. In these papers sparsity is characterized by the number of non-zero entries of a matrix. Amini and Wainwright (2009) considered principal component analysis in the high-dimensional regime "large p, small n" under the additional assumption that the maximal eigenvector is sparse.

Candès and Recht (2008), Candès and Tao (2009) considered the non-noisy setting ($\xi_i \equiv 0$) of the matrix completion problem under conditions that the singular vectors of A^* are sufficiently spread out on the unit sphere or "incoherent". They focused on exact recovery of A^*. In particular, Candès and Tao (2009) showed that under "strong incoherence condition" exact recovery is possible with high probability if $N > C\, rm \log^6 m$ with some constant $C > 0$ when we observe N entries of a square matrix $A^* \in \mathbb{R}^{m\times m}$ with locations uniformly sampled at random. Candès and Plan (2009) explored the same setting in the presence of noise, proposed estimators $\hat{A}$ of A^* and evaluated their Frobenius norm $\|\hat{A} - A^*\|_F$. They established bounds on $\|\hat{A} - A^*\|_F^2$ of the order m^3 when $A^* \in \mathbb{R}^{m\times m}$ and the noise is i.i.d. Gaussian. Moreover, they argued that, under their assumptions, even the oracle cannot achieve better rate in the Frobenius norm than rm^3/N, which is rather pessimistic. Indeed, $m^2 \gg N$ is a standard assumption in the matrix completion setting. Thus, nothing reasonable can be achieved for the Frobenius norm in matrix completion problem, and it makes sense to analyze other distance measures.

In this paper we consider the class of Schatten-p estimators $\hat{A}$, i.e., the penalized least squares estimators with a penalty proportional to Schatten-p norm, cf. (4). We study their convergence properties with respect to the prediction risk

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 = N^{-1} \sum_{i=1}^{N} \text{tr}^2(X_i'(\hat{A} - A^*))$$

and to the Schatten-q risk. For the prediction risk, we prove three types of results. First, under mild assumptions on the masks X_i we obtain bounds involving the Schatten-p norm of A^*. Second, essentially under no assumption on X_i, we show that Schatten-p estimators with p sufficiently close 0 achieve faster rates of convergence. This result is proved for square matrices $A^* \in \mathbb{R}^{m\times m}$ whose eigenvalues are not exponentially large. Third, we obtain bounds under the Restricted Isometry (RI) condition on the masks X_i, which is a rather strong condition, and with no assumption on A^*. The bounds for the Schatten-q risk of $\hat{A}$ are derived only in this third case, because in the first and second case the assumptions on X_i are so mild that identifiability is not guaranteed.

The main goal of this paper is to establish achievable rates of convergence of the estimators in the trace regression problem. For the reason of space, we do not provide lower bounds showing that the obtained rates cannot be improved. We rather give here a heuristic motivation. Assume for simplicity that A^* is a square $m \times m$ matrix with rank $(A^*) = r$. As mentioned above, the *intrinsic dimension* (the number of parameters to be estimated to recover A^*) is then $(2m - r)r$, which is of the order $\sim rm$ if $r \ll m$. It can be shown that under suitable assumptions including, for example, the normality of ξ_i, (1) is a regular model with Fisher information matrix $I(\theta)$ where θ is the vector of $\sim rm$ parameters to be estimated. For the lower bounds, we leave only the parameters, which are entries of the matrices with orthonormal rows and columns in the singular value decomposition of A^*. Due to the linearity of the model, the prediction risk $\hat{d}_{2,N}\big(\tilde{A}, A^*\big)^2$ of an estimator $\tilde{A}$ can be expressed, up to a constant factor, in the form $(\hat{\theta} - \theta)' I(\theta)(\hat{\theta} - \theta)$ where $\hat{\theta} = \hat{\theta}(\tilde{A})$ is some statistic. The standard minimax lower bound for the latter quantity in a regular model is $\frac{\text{dimension}}{\text{sample size}} \sim rm/N$. Thus, the rate that we would like to obtain for the prediction risk is of the order rm/N. The dimension reduction principle is expressed here by the fact that, if the rank r is small, we have $rm/N \ll m^2/N$, where m^2/N is the rate of estimation for a matrix of full rank. As discussed above, the case $m^2 \gg N$ is often of interest, so that the rate m^2/N makes no sense, whereas rm/N is meaningful. The main message of this paper is to show that the suitably tuned Schatten estimators attain the rate rm/N up to logarithmic factors.

Finally, it is useful to compare the results for matrix estimation when the sparsity is expressed by the rank with those for the high-dimensional vector estimation when the sparsity is expressed by the number of non-zero components of the vector. For the vector estimation we have the linear model

$$Y_i \;=\; X_i'\beta \;+ \xi_i, \quad i = 1, ..., N,$$

where $X_i \in \mathbb{R}^p$, $\beta \in \mathbb{R}^p$ and, for example, ξ_i are i.i.d. $\mathcal{N}(0,1)$ random variables. Consider the high-dimensional case, $p \gg N$. (This is analogous to the assumption $m^2 \gg N$ in the matrix problem and means that the nominal dimension is much larger than the sample size.) The sparsity assumption for the vector case has the form $s \ll N$, where s is the number of non-zero components, or the *intrinsic dimension* of β. Let $\hat{\beta}$ be an estimator of β. Then the optimal rate of convergence of the prediction risk $N^{-1}\sum_{i=1}^N (X_i'(\hat{\beta} - \beta))^2$ on the class of vectors β with given s is of the order s/N, up to logarithmic factors. This rate is shown to be attained, up to logarithmic factors, for many estimators, such as the BIC, the Lasso, the Dantzig selector, Sparse Exponential Weighting etc., cf., e.g., Bunea et al. (2007), Koltchinskii (2008), Bickel et al. (2009), Dalalyan and Tsybakov (2008). Note that this rate is of the form $\frac{\text{intrinsic dimension}}{\text{sample size}} = \frac{s}{N}$, up to a logarithmic factor. The general interpretation is therefore completely analogous to that of the matrix case. An interesting difference is that the logarithmic risk inflation factor is inevitable in the vector case (cf. Foster and George, 1994), but not in the matrix problem, as our results reveal.

The paper is organized as follows. In Section 2 we introduce the notation, some definitions and basic facts about the Schatten quasi-norms. Section 3 contains the definition of Schatten-p estimator and the main results about its rates of convergence expressed via bounds on the stochastic term of the risk. Section 4 concretizes different bounds on the stochastic term. In Section 5 they are combined with the theorems from Section 3 to provide final bounds for two examples: (i) multi-task learning and (ii) matrix completion with i.i.d. entries sampled uniformly at random. Sections 6 and 7 are devoted to the proofs.

2. Notation and preliminaries.

We will write $|\cdot|_2$ for the Euclidean norm in $\mathbb{R}^d$ for any integer d. For any matrix $A \in \mathbb{R}^{m\times T}$, we denote by $A_{(j,\cdot)}$ for $1 \le j \le m$ its jth row and write $A_{(\cdot,k)}$ for its kth column, $1 \le k \le T$. We denote by $\sigma_1(A) \ge \sigma_2(A) \ge \cdots \ge 0$ the singular values of A. The (quasi-) norm of some (quasi-) Banach space $\mathcal{B}$ is canonically denoted by $\|\cdot\|_{\mathcal{B}}$. In particular, for any matrix $A \in \mathbb{R}^{m\times T}$ and $0 < p < \infty$ we consider the Schatten (quasi-)norms

$$\|A\|_{S_p} = \Bigg(\sum_{j=1}^{\min(m,T)} \sigma_j(A)^p \Bigg)^{1/p} \quad \text{and} \quad \|A\|_{S_\infty} = \sigma_1(A).$$

The Schatten spaces S_p are defined as spaces of all matrices $A \in \mathbb{R}^{m\times T}$ equipped with quasi-norm $\|A\|_{S_p}$. The Schatten-2 norm coincides with the Frobenius norm:

$$\|A\|_{S_2} = \sqrt{\operatorname{tr}(A'A)} = \Big(\sum_{i,j} a_{ij}^2 \Big)^{1/2}$$

where a_{ij} denote the elements of matrix $A \in \mathbb{R}^{m\times T}$. Recall that for $0 < p < 1$ the Schatten spaces S_p are not normed but only quasi-normed, and $\|\cdot\|_{S_p}^p$ satisfies the inequality

$$\|A+B\|_{S_p}^p \le \|A\|_{S_p}^p + \|B\|_{S_p}^p \tag{3}$$

for any $0 < p \le 1$ and any two matrices $A, B \in \mathbb{R}^{m\times T}$, cf., e.g., Gohberg and Krein (1969).

Let $\mathcal{L} : \mathbb{R}^{m\times T} \to \mathbb{R}^N$ be the linear mapping defined by

$$A \mapsto \big(\operatorname{tr}(X_1'A), ..., \operatorname{tr}(X_N'A)\big)/\sqrt{N}.$$

Then

$$|\mathcal{L}(A)|_2^2 = N^{-1} \sum_{i=1}^N \operatorname{tr}^2(X_i'A).$$

We say that $\mathcal{L}$ satisfies the *Restricted Isometry condition* RI (r,ν) for some integer $1 \le r \le \min(m,T)$ and some $0 < \nu < \infty$ if there exists a constant $\delta_r \in (0,1)$ such that

$$(1-\delta_r)\|A\|_{S_2} \;\le\; \nu|\mathcal{L}(A)|_2 \;\le\; (1+\delta_r)\|A\|_{S_2}$$

for all matrices $A \in \mathbb{R}^{m\times T}$ of rank at most r.

A difference of this condition from the Restricted Isometry condition introduced by Candes and Tao (2005) in the vector case or from its analog for the matrix case suggested by Recht et al. (2007), is in the presence of factor ν. This factor is introduced to account for the fact that the masks X_i are typically very sparse, so that they do not induce isometries with coefficient close to one. Indeed, ν will be large in the examples that we consider below.

3. The estimator and its rates of convergence. In this paper we study the estimator $\hat{A}$ defined as a solution of the minimization problem

$$\min_{A\in\mathbb{R}^{m\times T}}\left(\frac{1}{N}\sum_{i=1}^{N}\big(Y_i-\mathrm{tr}(X_i'A)\big)^2\,+\,\lambda\|A\|_{S_p}^p\right) \tag{4}$$

with some fixed $0<p\le 1$ and $\lambda>0$. In other words, we consider a penalized least squares estimator with a penalty proportional to Schatten-p norm. The case $p=1$ is of outstanding interest since the minimization problem is then convex and thus can be efficiently solved in polynomial time. We call $\hat{A}$ the *Schatten-p estimator*. Such estimators have been recently considered by many authors motivated by applications to multi-task learning and collaborative filtering (Argyriou et al. (2007, 2008, 2009), Bach (2008), Abernethy et al.(2009)). These papers discussed connections of (4) to other related minimization problems, along with characterizations of the solutions and computational issues, mainly focusing on the convex case $p=1$. Also for the non-convex case ($0<p<1$), Argyriou et al. (2007, 2008) suggested an algorithm of approximate computation of Schatten-p estimator. However, for $0<p<1$ the methods can find only a local minimum in (4), so that Schatten estimators with such p remain for the moment mainly of theoretical value. In particular, analyzing these estimators reveals, which rates of convergence can, in principle, be attained.

The statistical properties of Schatten estimators are not yet well understood. To our knowledge, the only previous study is that of Bach (2008) showing that for $p=1$, under some condition on $X_i's$ (analogous to strong irrepresentability condition in the vector case, cf. Meinshausen and Bühlmann (2006), Zhao and Yu (2006)), $\mathrm{rank}(A^*)$ is consistently recovered by $\mathrm{rank}(\hat{A})$ when m,T are fixed and $N\to\infty$. Our results are of a different kind. They are non-asymptotic and meaningful in the case $mT\gg N>\max(m,T)$. Furthermore, we do not consider the recovery of the rank, but rather the estimation and prediction properties of Schatten-p estimators.

In what follows we will use the notation

$$\hat{d}_{2,N}(A,A')=|\mathcal{L}(A-A')|_2$$

where A and A' are any matrices in $\mathbb{R}^{m\times T}$. Unless the reverse is explicitly stated, we will tacitly assume that the X_i are non-random matrices.

We start with a simple observation. By definition of $\hat{A}$,

$$\frac{1}{N}\sum_{i=1}^{N}\big(Y_i-\mathrm{tr}(X_i'\hat{A})\big)^2\,+\,\lambda\|\hat{A}\|_{S_p}^p\le\frac{1}{N}\sum_{i=1}^{N}\big(Y_i-\mathrm{tr}(X_i'A^*)\big)^2\,+\,\lambda\|A^*\|_{S_p}^p.$$

Using that $Y_i=\mathrm{tr}(X_i'A^*)+\xi_i$, this can be transformed by a simple algebra to:

$$\hat{d}_{2,N}(\hat{A},A^*)^2\;\le\;\frac{2}{N}\sum_{i=1}^{N}\xi_i\mathrm{tr}\Big((\hat{A}-A^*)'X_i\Big)\;+\;\lambda\Big(\|\,A^*\,\|_{S_p}^{\,p}\;-\|\,\hat{A}\,\|_{S_p}^{\,p}\Big). \tag{5}$$

In the sequel, inequality (5) will be referred to as basic inequality and the random variable

$$\frac{1}{N}\sum_{i=1}^{N}\xi_i\mathrm{tr}\Big((\hat{A}-A^*)'X_i\Big)$$

will be called the stochastic term. In Section 4 we will show that under suitable conditions for any $0 < p \leq 1$ the stochastic term can be bounded with probability close to 1 as follows:

(6)

$$\Big| \frac{1}{N} \sum_{i=1}^{N} \xi_i \mathrm{tr}\big(X_i'(\hat{A} - A^*)\big) \Big| \;\leq\; \frac{\delta}{2} I_{\{0<p<1\}} \hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 + \tau\delta^{p-1} \|\hat{A} - A^*\|_{S_p}^p, \;\forall \delta > 0,$$

where $I_{\{\cdot\}}$ denotes the indicator function and $0 < \tau < \infty$ is a parameter depending on m, T and N. We will derive explicitly the values of τ for different examples.

From (5) and (6) with $\delta = 1/2$ and $\lambda = 4\tau$ we get

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \;\leq\; 8\tau\Big(\|\hat{A} - A^*\|_{S_p}^p + \|A^*\|_{S_p}^p \; - \|\hat{A}\|_{S_p}^p\Big).$$

This and (3) yield

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \;\leq\; 16\tau\|A^*\|_{S_p}^p. \tag{7}$$

Thus, we have the following result.

THEOREM 1. *Let $A^* \in \mathbb{R}^{m\times T}$, and let $0 < p \leq 1$. Assume that (6) holds with probability at least $1-\varepsilon$ for some $\varepsilon > 0$ and $0 < \tau < \infty$. Let $\hat{A}$ be the Schatten-p estimator defined as a minimizer of (4) with $\lambda = 4\tau$. Then (7) holds with probability at least $1 - \varepsilon$.*

The bound (7) depends on the magnitude of the elements of A^* via $\|A^*\|_{S_p}$. The next theorem shows that under the RI condition this dependence can be avoided, and only the rank of A^* affects the rate of convergence. Moreover, for small τ the rate becomes faster.

THEOREM 2. *Let $A^* \in \mathbb{R}^{m\times T}$ with $\mathrm{rank}(A^*) \leq r$, and let $0 < p \leq 1$. Assume that (6) holds with probability at least $1-\varepsilon$ for some $\varepsilon > 0$ and $0 < \tau < \infty$. Assume also that the Restricted Isometry condition RI $((2+a)r),\nu)$ holds with some $0 < \nu < \infty$, with a sufficiently large $a = a(p)$ depending only on p and with $0 < \delta_{(2+a)r} \leq \delta_0$ for a sufficiently small $\delta_0 = \delta_0(p)$ depending only on p.*

Let $\hat{A}$ be the Schatten-p estimator defined as a minimizer of (4) with $\lambda = 4\tau$. Then with probability at least $1-\varepsilon$ we have

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \;\leq\; C_1 r \tau^{\frac{2}{2-p}} \nu^{\frac{2p}{2-p}}, \tag{8}$$

$$\|\hat{A} - A^*\|_{S_q}^q \;\leq\; C_2 r \tau^{\frac{q}{2-p}} \nu^{\frac{2q}{2-p}}, \quad \forall\, q \in [p, 2], \tag{9}$$

where C_1 and C_2 are constants, C_1 depends only on p and C_2 depends on p and q.

Proof of Theorem 2 is given in Section 6. The values $a = a(p)$ and $\delta_0(p)$ can be deduced from the proof. In particular, for $p = 1$ it is sufficient to take $a = 19$.

We can compare Theorem 1 with Theorem 2 assuming that $p = 1$ and all the singular values of A^* are uniformly bounded by a constant. Then the bound (7) is of the order $r\tau$, whereas (8) is of the order $r\tau^2\nu$. Clearly, (8) decreases faster as $\tau \to 0$ but there is an additional scaling factor ν.

In general, the bounds of Theorem 2 depend on the positive powers of factor ν which can be large. For example, we will see below that in the multi-task learning setup ν is of the order $\sqrt{T}$. The resulting bounds turn out to be rather rough. The reason is that the masks X_i are sparse. The sparser are X_i, the larger is ν. The extreme situation corresponds to matrix completion problems. Indeed, let the X_i be point masks with $d = 1$ which are i.i.d. and uniformly distributed on

$$\Big\{ e_k(m)e_l'(T) : 1 \le k \le m,\ 1 \le l \le T \Big\}$$

(this was the setting considered in Candès and Recht (2008), Candès and Tao (2009), Candès and Plan (2009)). Then, introducing the notation $\delta_{kl}^{(i)} = I_{\{X_i = e_k(m)e_l'(T)\}}$, for any matrix $A \in \mathbb{R}^{m\times T}$ we have

$$\begin{aligned} mT|\mathcal{L}(A)|_2^2 = \frac{mT}{N}\sum_{i=1}^N \mathrm{tr}^2\big(X_i'A\big) \ &= \ \frac{mT}{N}\sum_{i=1}^N \sum_{k,l} a_{kl}^2\,\delta_{kl}^{(i)} \\ &= \ \sum_{k,l} a_{kl}^2 \Big(\frac{mT}{N}\sum_{i=1}^N \delta_{kl}^{(i)}\Big) \end{aligned} \tag{10}$$

with probability 1. Note that $\mathbb{E}\Big(\frac{mT}{N}\sum_{i=1}^N \delta_{kl}^{(i)}\Big) = 1$ for all k, l, and $\sum_{k,l} a_{kl}^2 = \|A\|_{S_2}^2$. So, the RI condition, if it holds, is naturally scaled by $\nu = \sqrt{mT}$, which is a very large value. What is more, to make this condition hold with probability close to 1, we need at least that the random variables $\frac{mT}{N}\sum_{i=1}^N \delta_{kl}^{(i)}$ were close to their expectations. But this is only possible if N is larger than mT, since $\delta_{kl}^{(i)}$ are i.i.d. Bernoulli$(1/(mT))$ variables. So, nothing can be done under the requirement $N \ll mT$ which is intrinsic for matrix completion problems. We see that Theorem 2 and, in general, the argument based on the restricted isometry or related conditions is not well adapted for such settings.

Motivated by this, we suggest another approach described in the next theorem. For simplicity we focus on the symmetric case $m = T$ and Gaussian errors ξ_i. We will also need the following condition.

Assumption 1. *There exists a constant* $c_0 < \infty$ *such that*

$$|\mathcal{L}(A)|_2^2 \ \le \ c_0\|A\|_{S_2}^2$$

for all matrices $A \in \mathbb{R}^{m\times T}$.

Note that Assumption 1 is not restrictive when the masks X_i are sparse. For example, (10) immediately implies that for the above special case of matrix completion problem Assumption 1 holds almost surely with $c_0 = 1$.

Theorem 3. *Let* $\xi_1, \dots, \xi_N$ *be i.i.d.* $\mathcal{N}(0, \sigma^2)$ *random variables, and assume that* $m = T > 1$, $N > \mathrm{e}\, m$, *and Assumption 1 hold. Let* $A^* \in \mathbb{R}^{m\times m}$ *with* $\mathrm{rank}(A^*) \le r$ *and the maximal singular value* $\sigma_1(A^*) \le (N/m)^{C^*}$ *for some* $0 < C^* < \infty$. *Set* $p = (\log(N/m))^{-1}$, $c_\kappa = (2\kappa - 1)(2\kappa)\kappa^{-1/(2\kappa-1)}$ *where* $\kappa = (2-p)/(2-2p)$ *and*

$$\tau = c_\kappa(\vartheta/p)^{1-p/2}\Big(\frac{m}{N}\Big)^{1-p/2} \tag{11}$$

for some $\vartheta \geq C^2$ *and* C *a universal positive constant independent of* r*,* m *and* N*. Then the Schatten-p estimator* $\hat{A}$ *defined as a minimizer of (4) with* $\lambda = 4\tau$ *satisfies:*

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \leq C_3 \vartheta \, \frac{rm}{N} \log\left(\frac{N}{m}\right) \tag{12}$$

with probability at least $1 - C\, \exp(-\vartheta m/C^2)$ *where the positive constant* C_3 *is independent of* r*,* m *and* N*.*

Proof. By Lemma 2, inequality (6) holds with probability at least $1 - C\, \exp(-\vartheta m/C^2)$. We then use (7) and note that, under our choice of p,

$$\tau \leq c\vartheta \left(\frac{m}{Np}\right)$$

for some constant $c < \infty$, which does not depend on m and N, and

$$\|A^*\|_{S_p}^p \leq r[\sigma_1(A^*)]^p \leq r\left(\frac{N}{m}\right)^{C^* p} = \exp(C^*)\, r.$$

□

Finally, we give the following theorem quantifying the rates of convergence of the prediction risk it terms of the Schatten norms of A^*.

Theorem 4. *Let* $\xi_1, \dots, \xi_N$ *be i.i.d.* $\mathcal{N}(0, \sigma^2)$ *random variables and let Assumption 1 hold. Then the Schatten-p estimator* $\hat{A}$ *has the following properties.*

(i) Let $A^* \in \mathbb{R}^{m \times T}$*,* $p = 1$*, and* $\lambda = 4\tau$ *where*

$$\tau = 8\sigma \sqrt{\frac{c_0(m+T)}{N}}.$$

Then

$$\hat{d}_{2,N}\left(\hat{A}, A^*\right)^2 \ \leq \ C\sigma c_0^{1/2} \|A^*\|_{S_1} \sqrt{\frac{m+T}{N}} \tag{13}$$

with probability at least $1 - 2\exp\{-(2 - \log 5)(m+T)\}$ *where* $C > 0$ *is an absolute constant.*

(ii) Let $A^* \in \mathbb{R}^{m \times m}$*,* $0 < p < 1$*, and* $\lambda = 4\tau$ *where* τ *is given in (11). Then*

$$\hat{d}_{2,N}\left(\hat{A}, A^*\right)^2 \ \leq \ C\|A^*\|_{S_p}^p \left(\frac{m}{N}\right)^{1-p/2} \tag{14}$$

with probability at least $1 - C\, \exp(-\vartheta m/C^2)$ *where the constant* $C > 0$ *is independent of* r*,* m *and* N*.*

Proof is straightforward in view of Theorem 1 and Lemmas 4, 2.

4. Bounds on the stochastic term. In this section we present bounds on the stochastic term leading to (6). The proofs are given in Section 7. The first bound is obtained under the Bernstein condition and it applies for the case $p = 1$. We say that the random variables ξ_i, $i = 1, \ldots, N$, satisfy the Bernstein condition if

$$\max_{1 \le i \le N} \mathbb{E}|\xi_i|^l \;\le\; \frac{1}{2} l!\, \sigma^2 H^{l-2}, \quad l = 2, 3 \ldots, \tag{15}$$

with some finite constants σ and H.

Lemma 1. *Let the i.i.d. zero-mean random variables ξ_i satisfy the Bernstein condition (15). Let also either*

$$\max_{1 \le j \le m} \frac{1}{N} \sum_{i=1}^{N} |X_{i(j,\cdot)}|_2^2 \;\le\; S_{row}^2 \quad \textit{and} \tag{16}$$

$$\max_{1 \le j \le m, 1 \le i \le N} |X_{i(j,\cdot)}|_2 \;\le\; H_{row} \tag{17}$$

or the conditions

$$\max_{1 \le k \le T} \frac{1}{N} \sum_{i=1}^{N} |X_{i(\cdot,k)}|_2^2 \;\le\; S_{col}^2 \quad \textit{and} \tag{18}$$

$$\max_{1 \le k \le T, 1 \le i \le N} |X_{i(\cdot,k)}|_2 \;\le\; H_{col} \tag{19}$$

hold true with some constants $S_{row}, H_{row}, S_{col}, H_{col}$. Let $D > 1$. Then, respectively, with probability at least $1 - 2/m^{D-1}$ or at least $1 - 2/T^{D-1}$ we have

$$\Big| \frac{1}{N} \sum_{i=1}^{N} \xi_i \mathrm{tr}\big(X_i'(\hat{A} - A^*)\big) \Big| \;\le\; \tau \|\hat{A} - A^*\|_{S_1} \tag{20}$$

where $\tau = C_{row}\sqrt{m(\log m)/N}$ if (16) and (17) hold or $\tau = C_{col}\sqrt{T(\log T)/N}\}$ if (18) and (19) hold. Here

$$C_{row} = \left(\sqrt{2D\sigma^2 S_{row}^2} + 2DH_{row}H\sqrt{\frac{\log m}{N}} \right), \;\; C_{col} = \left(\sqrt{2D\sigma^2 S_{col}^2} + 2DH_{col}H\sqrt{\frac{\log T}{N}} \right).$$

The second bound applies in the case $0 < p < 1$. It is given in the following lemma.

Lemma 2. *Let $\xi_1, \ldots, \xi_N$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables, $0 < p < 1$ and $m = T (\equiv M)$. Let Assumption 1 hold with some constant $c_0 < \infty$. Set $c_\kappa = (2\kappa - 1)(2\kappa)\kappa^{-1/(2\kappa-1)}$ where $\kappa = (2-p)/(2-2p)$. Then for any fixed $\delta > 0$, $\vartheta \ge C^2$ and*

$$\tau = c_\kappa (\vartheta/p)^{1-p/2} \Big(\frac{M}{N}\Big)^{1-p/2}$$

we have

$$\Big| \frac{1}{N} \sum_{i=1}^{N} \xi_i \mathrm{tr}\big(X_i'(\hat{A} - A^*)\big) \Big| \;\le\; \frac{\delta}{2} \hat{d}_{2,N}(\hat{A}, A^*)^2 + \tau\, \delta^{p-1} \big\| \hat{A} - A^* \big\|_{S_p}^p \tag{21}$$

with probability at least $1 - C\, \exp(-\vartheta M/C^2)$ for some constant $C = C(p, c_0, \sigma^2) > 0$ which is independent of M and N and satisfies $\sup_{0 < p \le q} C(p, c_0, \sigma) < \infty$ for all $q < 1$.

We finally give a bound on the stochastic term when $p = 1$ for a special case of matrix completion problem discussed in the previous section and considered by Candès and Recht (2008), Candès and Tao (2009) in the noiseless case and by Candès and Plan (2009) in the presence of noise. The following lemma allows one to treat the random noise under different conditions than in Candès and Plan (2009), and shows that there are some unusual effects.

LEMMA 3. *Let the i.i.d. zero-mean random variables ξ_i satisfy the Bernstein condition (15). Assume that $mT(m+T) > N$ and that X_i are point masks, which are i.i.d. uniformly distributed on*

$$\Big\{ e_k(m)e_l'(T) : 1 \le k \le m,\, 1 \le l \le T \Big\}$$

and independent from $\xi_1, \dots, \xi_N$. Then, for any $D \ge 2$ and

$$\tau = \big(4\sigma\sqrt{10D} + 8HD\big)\frac{m+T}{N} \tag{22}$$

we have (20) with probability at least $1 - 4\exp\{-(2-\log 5)(m+T)\}$.

Note that the concentration rate in Lemma 3 is faster than in the first lemma of this section. Indeed, it is exponential rather than polynomial in the dimension. The same effect is observed in the next lemma which provides a value of τ under Assumption 1.

LEMMA 4. *Let $\xi_1, ..., \xi_N$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables, and let Assumption 1 hold. Then, for any*

$$\tau \ge 8\sigma\sqrt{\frac{c_0(m+T)}{N}} \tag{23}$$

we have (20) with probability at least $1 - 2\exp\{-(2-\log 5)(m+T)\}$.

Lemma 4 can be compared with Lemma 1. First note that the assumptions of Lemma 1 are not directly comparable with Assumption 1, provided that we consider the latter with c_0 independent of m and T. If m and T have the same order of magnitude, the bound of Lemma 4 is better. On the other hand, if m and T differ dramatically, for example, $m \gg T$, then Lemma 1 can provide a significant improvement. Indeed, the "column" version of Lemma 1 guarantees the rate $\tau \sim \sqrt{\frac{T\log T}{N}}$ which in this case is much smaller than $\sqrt{\frac{m}{N}}$. In all the cases, the concentration rate in Lemma 4 is faster than in Lemma 1, but one should note that the latter result holds for a larger class of error distributions.

5. Examples: matrix completion, multi-task learning. 1. MATRIX COMPLETION. As discussed in Section 3, for matrix completion problems the restricted isometry argument as in Theorem 2 is typically not applicable. We will therefore use Theorems 1 and 3. Note that the bound of Theorem 4 is too coarse in the context of matrix completion. Indeed, it can be replaced by the following corollary, which is an immediate consequence of Theorem 1 with $p = 1$ and Lemma 3.

Corollary 1. *Let the i.i.d. zero-mean random variables* ξ_i *satisfy the Bernstein condition (15). Assume that* $mT(m+T) > N$ *and that* X_i *are point masks, which are i.i.d. uniformly distributed on*

$$\Big\{ e_k(m) e_l'(T) : 1 \le k \le m,\ 1 \le l \le T \Big\} \tag{24}$$

and independent from $\xi_1, \dots, \xi_N$. *Let* τ *be given by (22) with some* $D \ge 2$. *Then the Schatten-*1 *estimator* $\hat{A}$ *defined with* $\lambda = 4\tau$ *satisfies:*

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \le 16 \bar{C} \|A^*\|_{S_1} \frac{m+T}{N} \tag{25}$$

with probability at least $1 - 4\exp\{-(2-\log 5)(m+T)\}$, *where* $\bar{C} = 4\sigma\sqrt{10D} + 8HD$.

We see that the rate of convergence in Corollary 1 is substantially faster than in Theorem 1. This is not surprising because the matrix completion masks (24) are the sparsest possible; they contain only one non-zero entry.

If we assume that the maximal singular value $\sigma_1(A^*)$ is bounded from above by some constant, then it is easy to see that for $m = T$ the bounds of Corollary 1 achieve the optimal rate "intrinsic dimension/sample size" $\sim rm/N$.

The next corollary follows immediately from Theorem 3 and the remarks before it.

Corollary 2. *Let* $\xi_1, \dots, \xi_N$ *be i.i.d.* $\mathcal{N}(0, \sigma^2)$ *random variables, and assume that* $m = T > 1$, $N > \mathrm{e}\, m$ *and that the* X_i *are point masks, which are i.i.d. uniformly distributed on*

$$\Big\{ e_k(m) e_l'(m) : 1 \le k \le m,\ 1 \le l \le m \Big\}$$

and independent from $\xi_1, \dots, \xi_N$. *Let* $A^* \in \mathbb{R}^{m \times m}$ *with* $\mathrm{rank}(A^*) \le r$ *and the maximal singular value* $\sigma_1(A^*) \le (N/m)^{C^*}$ *for some* $0 < C^* < \infty$. *Set* $p = (\log(N/m))^{-1}$, $c_\kappa = (2\kappa - 1)(2\kappa)\kappa^{-1/(2\kappa-1)}$ *where* $\kappa = (2-p)/(2-2p)$ *and*

$$\tau = c_\kappa (\vartheta/p)^{1-p/2} \left(\frac{m}{N} \right)^{1-p/2}$$

for some $\vartheta \ge C^2$ *with a universal constant* $C > 0$, *independent of* r, m *and* N. *Then the Schatten-*p *estimator* $\hat{A}$ *defined as a minimizer of (4) with* $\lambda = 4\tau$ *satisfies:*

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \le C_3 \vartheta \frac{rm}{N} \log\left(\frac{N}{m} \right) \tag{26}$$

with probability at least $1 - C\exp(-\vartheta m / C^2)$, *where the positive constant* C_3 *is also independent of* r, m *and* N.

Note that the bound of Corollary 2 achieves the optimal rate rm/N, up to logarithmic factor. This holds under weaker condition on A^* than the boundedness of $\sigma_1(A^*)$ needed to obtain the rate rm/N from Corollary 1 as discussed above.

2. Multi-task learning. For multi-task learning we can use both Theorem 2 and Theorem 3. Theorem 2 imposes stronger assumptions on the masks X_i, namely the restricted isometry (RI) condition. Nevertheless, the advantage is that Theorem 2 covers the computationally easy case $p = 1$.

Let us first discuss the form of the RI condition in the context of multi-task learning. Using the analog of (2) for a matrix $A = (a_1 \cdots a_T)$ we obtain

$$\begin{aligned} |\mathcal{L}(A)|_2^2 &= N^{-1} \sum_{i=1}^{N} \mathrm{tr}^2(X_i' A) \\ &= N^{-1} \sum_{t=1}^{T} \sum_{s=1}^{n} a_t' \mathbf{x}^{(t,s)} (\mathbf{x}^{(t,s)})' a_t \\ &= T^{-1} \sum_{t=1}^{T} a_t' \Psi_t a_t \end{aligned}$$

where $\Psi_t = n^{-1} \sum_{s=1}^{n} \mathbf{x}^{(t,s)} (\mathbf{x}^{(t,s)})'$ is the Gram matrix of predictors for the tth task. These matrices correspond to T separate regression models. The standard assumption is that they are normalized so that all the diagonal elements of each Ψ_t are equal to 1. This suggests that the natural RI scaling factor ν for such model is of the order $\sqrt{T}$. For example, in the simplest case when all the matrices Ψ_t are just equal to the $m \times m$ identity matrix, we find

$$|\mathcal{L}(A)|_2^2 = T^{-1} \sum_{t=1}^{T} a_t' \Psi_t a_t = T^{-1} \|A\|_{S_2}^2 .$$

Similarly, we get the RI condition with scaling factor $\nu \sim \sqrt{T}$ when the spectra of all the Gram matrices $\Psi_t,\ t = 1, \dots, T$, are included in a fixed interval $[a, b]$ with $0 < a < b < \infty$.

We now concretize the bounds for the stochastic term. The first approach is to use Lemma 1. Note that we can choose the best of the two bounds corresponding either to the ”row” conditions (16), (17) or to the ”column” conditions (18), (19). Interestingly, the two bounds are of different order of magnitude. To explain this, assume for simplicity that all the elements of vectors $\mathbf{x}^{(t,s)}$ are uniformly bounded in absolute value by a constant x_0. Then for the row norms we have $\max_{i,j} |X_{i(j,\cdot)}|_2 \le x_0$ because each row of X_i contains only one non-zero entry. Thus, we can take $S_{row} = x_0^2, H_{row} = x_0$, and the ”row” version of Lemma 1 gives $\tau \sim \sqrt{\frac{m \log m}{N}} = \sqrt{\frac{m \log m}{nT}}$. On the other hand, the ”column” version of Lemma 1 leads to a coarser bound. Indeed, the columns of X_i can contain m non-zero entries, so that their norms can be only bounded by $x_0\sqrt{m}$, and thus $S_{col} = x_0^2 m, H_{col} = x_0\sqrt{m}$. The ”column” version of Lemma 1 then gives $\tau \sim \sqrt{\frac{mT \log T}{N}} = \sqrt{\frac{m \log T}{n}}$. Discarding the logarithmic factors, this is larger in order than τ obtained from the ”row” version. With this best τ, we get the following immediate corollary of Theorem 2 with $p = 1$.

Corollary 3. *Let the i.i.d. zero-mean random variables ξ_i satisfy the Bernstein condition (15). Consider the multi-task learning problem with* $\mathrm{rank}(A^*) \le r$. *Assume that all the elements of vectors $\mathbf{x}^{(t,s)}$ are uniformly bounded in absolute value by a constant x_0. Assume also that the Restricted Isometry condition RI $(21r,\nu)$ holds with some $0 < \nu < \infty$, and with $0 < \delta_{21r} \le \delta_0$ for a sufficiently small δ_0. Set*

$$\lambda = 4x_0 \left(\sigma\sqrt{2D} + 2DH\sqrt{\frac{\log m}{nT}} \right) \sqrt{\frac{m \log m}{nT}}$$

for some $D > 1$. Let $\hat{A}$ be the Schatten-1 estimator with this parameter λ. Then with probability at least $1 - 2/m^{D-1}$ we have

$$\begin{aligned} \hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 &\le C_1' r\nu^2 \left(1 + \sqrt{\frac{\log m}{nT}}\right)^2 \frac{m \log m}{nT}, \\ \|\hat{A} - A^*\|_{S_q}^q &\le C_2' r\nu^{2q} \left(1 + \sqrt{\frac{\log m}{nT}}\right)^q \left(\frac{m \log m}{nT}\right)^{q/2}, \quad \forall\, q \in [1,2], \end{aligned}$$

where the constants C_1' and C_2' do not depend on m, T and n.

Some comments may be useful here. As discussed above, the right scaling factor ν for multi-task learning is of the order $\sqrt{T}$. This and the natural assumption $\log m \le nT$ transform the bounds of Corollary 3 (with $q = 2$) to:

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \le C_1'' \frac{rm \log m}{n}, \tag{27}$$

$$\frac{1}{T}\|\hat{A} - A^*\|_{S_2}^2 \le C_2'' \frac{rm \log m}{n}, \tag{28}$$

where the constants C_1'' and C_2'' do not depend on m, T and n. The bounds (27), (28) look similarly to those obtained for the Group Lasso estimator in multi-task setting by Lounici et al. (2009). The main difference is that the sparsity index s appearing in Lounici et al. (2009) is now replaced by rm. At first sight, this seems natural because s and rm are some quantities characterizing the dimension of the problem. In Lounici et al. (2009), the columns a_t^* of A^* were supposed to be sparse, with the sets of non-zero elements of cardinality not more than s. Here the sparsity is characterized by the rank r of A^*, and thus the intrinsic dimension of A^* is of order rm, as discussed in the introduction.

However, comparing parameter s of Lounici et al. (2009) to our rm is not fair. Indeed, in our case rm characterizes the number of unknown parameters to be estimated. In Lounici et al. (2009) the number of unknown parameters is of the order sT (s parameters in each of T columns), not s. This suggests that the bounds (27), (28) are, in fact, too coarse, as compared to their analogs in Lounici et al. (2009). Indeed, the next results show that they can be improved by the factor of T under somewhat more restrictive condition on $X_i's$.

COROLLARY 4. *Let $\xi_1, \dots, \xi_N$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables. Consider the multi-task learning problem with $\operatorname{rank}(A^*) \le r$. Assume that the spectra of the Gram matrices Ψ_t are uniformly in t bounded from above by a constant $c_1 < \infty$. Assume also that the Restricted Isometry condition RI ($21r$,ν) holds with some $0 < \nu < \infty$ and with $0 < \delta_{21r} \le \delta_0$ for a sufficiently small δ_0. Set*

$$\lambda = 32\sigma \sqrt{\frac{c_1(m+T)}{nT^2}}\,.$$

Let $\hat{A}$ be the Schatten-1 estimator with this parameter λ. Then with probability at least $1 - 2\exp\{-(2 - \log 5)(m+T)\}$ we have

$$\begin{aligned} \hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 &\le \bar{C}_1 c_1 \sigma^2 r\nu^2 \left(\frac{m+T}{nT^2}\right), \\ \|\hat{A} - A^*\|_{S_q}^q &\le \bar{C}_2 c_1^{q/2} \sigma^q r\nu^{2q} \left(\frac{m+T}{nT^2}\right)^{q/2}, \quad \forall\, q \in [1,2], \end{aligned}$$

where $\bar{C}_1$ is an absolute constant and $\bar{C}_2$ depends only on q.

Proof of Corollary 4 is straightforward in view of Theorem 2, Lemma 4, and the fact that, under the premises of Corollary 4, we have

$$|\mathcal{L}(A)|_2^2 \; = T^{-1}\sum_{t=1}^{T} a_t'\Psi_t a_t \le \; (c_1/T)\|A\|_{S_2}^2$$

for all matrices $A \in \mathbb{R}^{m\times T}$, so that Assumption 1 holds with $c_0 = c_1/T$.

Taking in the bounds of Corollary 4 the natural scaling factor $\nu \sim \sqrt{T}$ we obtain the following analogs of (27) and (28)

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \;\le\; \tilde{C}_1\,\frac{r(m+T)}{nT}, \tag{29}$$

$$\frac{1}{T}\|\hat{A} - A^*\|_{S_2}^2 \;\le\; \tilde{C}_2\,\frac{r(m+T)}{nT}, \tag{30}$$

where the constants $\tilde{C}_1$ and $\tilde{C}_2$ do not depend on m, T and n. We see that the rates in (29) and (30) are strictly better than in (27) and (28). Moreover, due to Corollary 4, the probability for (29) and (30) to hold converges exponentially to 1 if either m or T tends to ∞. In the case of (27) and (28), Corollary 3 assures only polynomial in m convergence of the probabilities. Note, however, that the assumptions of Corollary 4 are more restrictive.

A remarkable fact is that the rates in Corollary 4 are free of the logarithmic inflation factor. This is one of the differences between the matrix estimation problems and vector estimation ones, where the logarithmic risk inflation is inevitable, as first noticed by Foster and George (1994).

Note also that for $m = T$ the bounds (29) and (30) achieve the optimal rate "intrinsic dimension/sample size" $\sim rm/N$ ($N = nT$ in the multi-task learning).

Finally, we give the following corollary based on application of Lemma 2.

Corollary 5. *Let $\xi_1, \dots, \xi_N$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables, and assume that $m = T > 1$, $n > \mathrm{e}$. Consider the multi-task learning problem with $A^* \in \mathbb{R}^{m\times m}$, $\operatorname{rank}(A^*) \le r$ such that the maximal singular value $\sigma_1(A^*) \le n^{C^*}$ for some $0 < C^* < \infty$. Assume that the spectra of the Gram matrices Ψ_t are uniformly in t bounded from above by $c_0 T$ where $c_0 < \infty$ is a constant. Set $p = (\log n)^{-1}$, $c_\kappa = (2\kappa-1)(2\kappa)\kappa^{-1/(2\kappa-1)}$ where $\kappa = (2-p)/(2-2p)$ and*

$$\lambda = 4c_\kappa(\vartheta/p)^{1-p/2}\Big(\frac{1}{n}\Big)^{1-p/2}$$

for some $\vartheta \ge C^2$ and a universal constant $C > 0$, independent of r, m and n. Then the Schatten-p estimator $\hat{A}$ with this parameter λ satisfies

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \le C_3\vartheta\,\frac{r}{n}\log n \tag{31}$$

with probability at least $1 - C\,\exp(-\vartheta m/C^2)$ where the positive constant C_3 is independent of r, m and n.

Corollary 5 follows from Theorem 3. Indeed, it suffices to remark that, under the premises of Corollary 5, we have

$$|\mathcal{L}(A)|_2^2 \ = T^{-1}\sum_{t=1}^{T} a_t' \Psi_t a_t \leq \ c_0 \|A\|_{S_2}^2$$

for all matrices $A \in \mathbb{R}^{m\times m}$, so that Assumption 1 holds.

Since $m = T$, we can write (31) in the form

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \leq C_3' \frac{rm}{nT} \log n \tag{32}$$

Clearly, this bound is better than (27) for large T. Moreover, it achieves the optimal rate ”intrinsic dimension/sample size”$\sim rm/N$, up to logarithms (recall that $N = nT$ in the multi-task learning).

Another remark concerns the possible range of m. In (27) we need at least $rm < n$ to get a reasonable bound. So, the ”dimension larger than the sample size” framework is not covered by Corollary 3. In contrast to this, the bounds of Corollary 4 (for $m \sim T$) and Corollary 5 make sense when the dimension m is larger than the task sample size n; we only need the condition $m \ll \exp(n)$ for Corollary 5 to be meaningful, and no condition on the dimension m in Corollary 4.

Note also that, on the difference from Corollaries 3 and 4, Corollary 5 holds when the RI assumption is violated and under a weaker condition on the masks X_i. The price to pay is to assume that the singular values of A^* do not grow exponentially fast. Also, the estimator of Corollary 5 corresponds to $p < 1$, so that no efficient computational algorithms are actually available. Therefore, Corollary 5 remains mainly of theoretical interest.

6. Proof of Theorem 2.

We first give two lemmas on matrix decomposition needed in our proof, which are essentially provided by Recht, Fazel and Parrilo (2008) (subsequently, RFP(08) for short).

LEMMA 5. *Let A and B be matrices of the same dimension. If $AB' = 0$, $A'B = 0$, then*

$$\|A+B\|_{S_p}^p \ = \ \|A\|_{S_p}^p \ + \ \|B\|_{S_p}^p, \quad \forall\ p > 0.$$

PROOF. For $p = 1$ the result is Lemma 2.3 in RFP(08). The argument obviously extends to any $p > 0$ since RFP(08) show that the singular values of $A+B$ are equal to the union (with repetition) of the singular values of A and B. □

LEMMA 6. *Let $A \in \mathbb{R}^{m\times T}$ with $\mathrm{rank}(A) = r$ and singular value decomposition $A = U\Lambda V'$. Let $B \in \mathbb{R}^{m\times T}$ be arbitrary. Then there exists a decompositon $B = B_1 + B_2$ with the following properties:*

(i) $\mathrm{rank}(B_1) \leq 2\mathrm{rank}(A) = 2r$,

(ii) $AB_2' = 0$, $A'B_2 = 0$,

(iii) $\mathrm{tr}\big(B_1'B_2\big) = 0$.

(iv) B_1 and B_2 are of the form

$$B_1 \ = \ U\begin{pmatrix} \tilde{B}_{11} & \tilde{B}_{12} \\ \tilde{B}_{21} & 0 \end{pmatrix} V' \quad \text{and} \quad B_2 \ = \ U\begin{pmatrix} 0 & 0 \\ 0 & \tilde{B}_{22} \end{pmatrix} V'.$$

The points (i)-(iii) are the statement of Lemma 3.4 in RFP(08), the representation (iv) is provided in its proof.

Proof of Theorem 2. First note that there exists a decomposition $\hat{A} = \hat{A}^{(1)} + \hat{A}^{(2)}$ with the following properties:

(i) $\mathrm{rank}(\hat{A}^{(1)} - A^*) \le 2\,\mathrm{rank}(A^*) = 2r$,

(ii) $A^*(\hat{A}^{(2)})' = 0$, $(A^*)'\hat{A}^{(2)} = 0$,

(iii) $\mathrm{tr}\big((\hat{A}^{(1)} - A^*)'\hat{A}^{(2)}\big) = 0$.

This follows from Lemma 6 with $A = A^*$ and $B = \hat{A} - A^*$. In the notation of Lemma 6 we have $B_1 = \hat{A}^{(1)} - A^*$ and $B_2 = \hat{A}^{(2)}$.

From the basic inequality (5) and (6) with $\delta = 1/2$ we find

$$(1 - I_{\{0<p<1\}}/2)\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \ \le\ 2^{2-p}\tau\|\hat{A} - A^*\|_{S_p}^p \ +\ 4\tau\Big(\|A\|_{S_p}^p - \|\hat{A}\|_{S_p}^p\Big). \tag{33}$$

In particular, for the case $p = 1$

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \ \le\ 2\tau\|\hat{A} - A^*\|_{S_p}^p \ +\ 4\tau\Big(\|A\|_{S_p}^p - \|\hat{A}\|_{S_p}^p\Big). \tag{34}$$

For brevity, we will conduct the proof with the numerical constants given in (34), i.e., with those for $p = 1$. The proof for general p differs only in the values of the constants, but their expressions become cumbersome.

Using (3), we get

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \ \le\ 2\tau\|\hat{A} - A^*\|_{S_p}^p \ +\ 4\tau\|A^*\|_{S_p}^p + 2\tau\|\hat{A}^{(2)}\|_{S_p}^p \ -\ 4\tau\|\hat{A}\|_{S_p}^p. \tag{35}$$

By (3) again and by Lemma 5,

$$\begin{aligned}\|\hat{A}\|_{S_p}^p \ &\ge\ \|A^* + \hat{A}^{(2)}\|_{S_p}^p - \|\hat{A}^{(1)} - A^*\|_{S_p}^p \\ &=\ \|A^*\|_{S_p}^p + \|\hat{A}^{(2)}\|_{S_p}^p - \|\hat{A}^{(1)} - A^*\|_{S_p}^p,\end{aligned}$$

since $(A^*)'\hat{A}^{(2)} = 0$ and $A^*(\hat{A}^{(2)})' = 0$ by construction. Together with (35) this yields

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \ \le\ 2\tau\|\hat{A}^{(1)} - A\|_{S_p}^p - 2\tau\|\hat{A}^{(2)}\|_{S_p}^p + 4\tau\|\hat{A}^{(1)} - A^*\|_{S_p}^p, \tag{36}$$

from which one may deduce

$$\hat{d}_{2,N}(\hat{A}, A^*)^2 \ \le\ 6\tau\|\hat{A}^{(1)} - A^*\|_{S_p}^p \ \text{ and} \tag{37}$$

$$\|\hat{A}^{(2)}\|_{S_p}^p \ \le\ 3\|\hat{A}^{(1)} - A^*\|_{S_p}^p. \tag{38}$$

Consider now the following decomposition of the matrix $\hat{A}^{(2)}$. First recall that $\hat{A}^{(2)}$ is of the form

$$\hat{A}^{(2)} \ =\ U\begin{pmatrix} 0 & 0 \\ 0 & \tilde{B}_{22}\end{pmatrix}V'.$$

Write $\tilde{B}_{22} = W_1 \Lambda(\tilde{B}_{22}) W_2'$ with diagonal matrix $\Lambda(\tilde{B}_{22})$ of dimension r' and $W_1'W_1 = W_2'W_2 = I_{r' \times r'}$ for some $r' \leq \min(m,T)$. In the next step, W_1 and W_2 are complemented to orthogonal matrices $\bar{W}_1$ and $\bar{W}_2$ of dimension $\min(m,T) \times \min(m,T)$. For instance, set

$$\bar{W}_2' \;=\; \begin{pmatrix} & 0 \\ * & \\ & W_2' \end{pmatrix} \in \mathbb{R}^{\min(m,T)\times\min(m,T)},$$

where $*$ complements the columns of the matrix $\binom{0}{W_2'}$ to an orthonormal basis in $\mathbb{R}^{m\times T}$, and proceed analogously with W_1. In particular, $\bar{W}_1'\bar{W}_1 = \bar{W}_2'\bar{W}_2 = I_{\min(m,T)\times\min(m,T)}$. Also

$$\hat{A}^{(2)} \;=\; U \begin{pmatrix} 0 & 0 \\ 0 & W_1\Lambda(\tilde{B}_{22})W_2' \end{pmatrix} V' \;=\; U\bar{W}_1 \begin{pmatrix} 0 & 0 \\ 0 & \Lambda(\tilde{B}_{22}) \end{pmatrix} \bar{W}_2'V' \;=:\; U\bar{W}_1 D \bar{W}_2' V'.$$

We now represent $\hat{A}^{(2)}$ as a finite sum of matrices $\hat{A}^{(2)} = \sum_{j=1}^{R'} \hat{A}_j^{(2)}$ with

$$\hat{A}_i^{(2)} \;=\; U\bar{W}_1 D_i \bar{W}_2' V'$$

and

$$D_i \;=\; \begin{pmatrix} 0 & 0 \\ 0 & \Lambda_i \end{pmatrix}$$

where the $r' \times r'$ diagonal matrix Λ_i has the form $\Lambda_i = \mathrm{diag}(\lambda_j I_{\{j\in I_j\}})$. We denote here by I_1 the set of ar indices from $\{1,\dots,\min(m,T)\}$ corresponding to the ar largest in absolute value diagonal entries of Λ, by I_2 the set of indices corresponding to the next ar largest in absolute value diagonal entries λ_j, etc. Clearly, the matrices $\hat{A}_k^{(2)}$ are mutually orthogonal: $\mathrm{tr}\big((\hat{A}_j^{(2)})'\hat{A}_k^{(2)}\big) = 0$ for $j \neq k$ and $\mathrm{rank}(\hat{A}_j^{(2)}) \leq ar$. Moreover, $\hat{A}_i^{(2)}$ is orthogonal to $\hat{A}^{(1)} - A^*$.

Let $\sigma_1 \geq \sigma_2 \geq \dots$ be the singular values of $\hat{A}^{(2)}$, then $\sigma_1 \geq \dots \geq \sigma_{ar}$ are the singular values of $\hat{A}_1^{(2)}$, $\sigma_{ar+1} \geq \dots \geq \sigma_{2ar}$ those of $\hat{A}_2^{(2)}$, etc. By construction, we have $\mathrm{Card}(I_i) = ar$ for all i, and for all $k \in I_{i+1}$:

$$\sigma_k \;\leq\; \min_{j\in I_i} \sigma_j \;\leq\; \Big(\frac{1}{ar}\sum_{j\in I_i} \sigma_j^p\Big)^{1/p}.$$

Thus,

$$\sum_{k\in I_{i+1}} \sigma_k^2 \;\leq\; ar\Big(\frac{1}{ar}\sum_{j\in I_i}\sigma_j^p\Big)^{2/p}$$

from which one can deduce for all $j \geq 2$:

$$\|\hat{A}_j^{(2)}\|_{S_2} \;=\; \Big(\sum_{k\in I_{i+1}} \sigma_k^2\Big)^{1/2} \;\leq\; (ar)^{\frac{1}{2}-\frac{1}{p}}\Big(\sum_{j\in I_i}\sigma_j^p\Big)^{1/p} \;=\; (ar)^{\frac{1}{2}-\frac{1}{p}} \|\hat{A}_j^{(2)}\|_{S_p}$$

and consequently

$$\sum_{j\geq 2} \|\hat{A}_j^{(2)}\|_{S_2} \;\leq\; (ar)^{\frac{1}{2}-\frac{1}{p}} \sum_{j\geq 1} \|\hat{A}_j^{(2)}\|_{S_p}$$

Because of the elementary inequality $x^{1/p}+y^{1/p}\leq(x+y)^{1/p}$ for any non-negative x,y and $0<p\leq 1$,

$$\sum_{j\geq 2}\|\hat{A}^{(2)}_j\|_{S_p} = \sum_{j\geq 2}\Big(\sum_{k\in I_j}\sigma_k^p\Big)^{1/p} \leq \Big(\sum_{j\geq 2}\sum_{k\in I_j}\sigma_k^p\Big)^{1/p} \leq \Big(\sum_k \sigma_k^p\Big)^{1/p} = \|\hat{A}^{(2)}\|_{S_p}.$$

Therefore,

$$\begin{aligned}\sum_{j\geq 2}\|\hat{A}^{(2)}_j\|_{S_2} &\leq (ar)^{\frac{1}{2}-\frac{1}{p}}\|\hat{A}^{(2)}\|_{S_p}\\ &\leq 3^{1/p}(ar)^{\frac{1}{2}-\frac{1}{p}}\|\hat{A}^{(1)}-A^*\|_{S_p} \quad \text{(using inequality (38))}\\ &\leq 3^{1/p}(ar)^{\frac{1}{2}-\frac{1}{p}}(2r)^{\frac{1}{p}-\frac{1}{2}}\|\hat{A}^{(1)}-A^*\|_{S_2},\end{aligned}$$

whereby the last inequality results from $\operatorname{rank}\big(\hat{A}^{(1)}-A^*\big)\leq 2r$ and

$$\Big(\frac{1}{2r}\sum_{k\leq 2r}\sigma_k^p\Big)^{1/p} \leq \Big(\frac{1}{2r}\sum_{k\leq 2r}\sigma_k^2\Big)^{1/2}.$$

Finally,

$$\sum_{j\geq 2}\|\hat{A}^{(2)}_j\|_{S_2} \leq 3^{1/p}\Big(\frac{a}{2}\Big)^{\frac{1}{2}-\frac{1}{p}}\|\hat{A}^{(1)}-A^*\|_{S_2}. \tag{39}$$

We now proceed with the final argument. First note that $\operatorname{rank}\big((\hat{A}^{(1)}-A^*)+\hat{A}^{(2)}_1\big)\leq(2+a)r$. Next, by the triangular inequality, the restricted isometry condition and the orthogonality of $\hat{A}^{(2)}_j$ and $\hat{A}^{(1)}-A^*$ we obtain

$$\begin{aligned}\nu\hat{d}_{2,N}(\hat{A},A^*) &= \nu\,|\,\mathcal{L}(\hat{A}-A^*)\,|_2\\ &\geq \nu\,|\,\mathcal{L}(\hat{A}^{(1)}-A^*+\hat{A}^{(2)}_1)\,|_2-\nu\sum_{j\geq 2}|\,\mathcal{L}(\hat{A}^{(2)}_j)\,|_2\\ &\geq \big(1-\delta_{(2+a)r}\big)\|\hat{A}^{(1)}-A^*+\hat{A}^{(2)}_1\|_{S_2} - (1+\delta_{ar})\sum_{j\geq 2}\|\hat{A}^{(2)}_j\|_{S_2}\\ &\geq \|\hat{A}^{(1)}-A^*\|_{S_2}\Big(\big(1-\delta_{(2+a)r}\big)-\big(1+\delta_{ar}\big)3^{1/p}\Big(\frac{a}{2}\Big)^{\frac{1}{2}-\frac{1}{p}}\Big). \qquad (40)\end{aligned}$$

Recall that

$$a = \min\Big\{k\in\mathbb{N}:\ k>\big(6^{1/p}/\sqrt{2}\big)^{\frac{2p}{2-p}}\Big\}.$$

Then $1-3^{1/p}\big(a/2\big)^{\frac{1}{2}-\frac{1}{p}}>0$. Now, $\delta_{(2+a)r}\geq\delta_{ar}$, and thus

$$\big(1-\delta_{(2+a)r}\big)-\big(1+\delta_{ar}\big)3^{1/p}\Big(\frac{a}{2}\Big)^{\frac{1}{2}-\frac{1}{p}} \geq \Big(1-3^{1/p}\Big(\frac{a}{2}\Big)^{\frac{1}{2}-\frac{1}{p}}\Big)-2\delta_{(2+a)r} > 0$$

whenever

$$\delta_{(2+a)r} < \frac{1}{2}\Big(1-3^{1/p}\Big(\frac{a}{2}\Big)^{\frac{1}{2}-\frac{1}{p}}\Big).$$

Thus, there exists a universal constant $\kappa=\kappa(p)$ such that

$$\nu^2\hat{d}_{2,N}\big(\hat{A},A^*\big)^2 \geq \kappa\|\hat{A}^{(1)}-A^*\|^2_{S_2}. \tag{41}$$

Now, the inequalities (37) and (41) yield

$$\kappa \| \hat{A}^{(1)} - A^* \|_{S_2}^2 \;\le\; 6\tau\nu^2 \| \hat{A}^{(1)} - A^* \|_{S_p}^p \;\le\; 6\tau\nu^2 (2r)^{1-p/2} \| \hat{A}^{(1)} - A^* \|_{S_2}^p , \tag{42}$$

where the latter inequality results from the fact that we have $\operatorname{rank}\big(\hat{A}^{(1)} - A^*\big) \le 2r$, which implies

$$\| \hat{A}^{(1)} - A^* \|_{S_p} \;\le\; (2r)^{1-p/2} \| \hat{A}^{(1)} - A^* \|_{S_2} . \tag{43}$$

From (41) and (42) we obtain

$$\kappa \| \hat{A}^{(1)} - A^* \|_{S_2}^{2-p} \;\le\; 6\tau\nu^2 (2r)^{1-p/2} . \tag{44}$$

Now, from (37), (43) and (44) we find

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^2 \;\le\; 6\tau (2r)^{1-p/2} \| \hat{A}^{(1)} - A^* \|_{S_2}^p \le 2r (6\tau)^{\frac{2}{2-p}} \kappa^{-\frac{p}{2-p}} \nu^{\frac{2p}{2-p}} . \tag{45}$$

This proves (8). It remains to prove (9). We first demonstrate (9) for $q = 2$, then for $q = p$, and finally obtain (9) for all $q \in [p, 2]$ by Schatten norm interpolation.

Using (39), (40), (45), we find

$$\begin{aligned} \big(1 - \delta_{(2+a)r}\big) \| \hat{A}^{(1)} - A^* + \hat{A}_1^{(2)} \|_{S_2} &\le \nu \hat{d}_{2,N}\big(\hat{A}, A^*\big) + \big(1 + \delta_{ar}\big) \sum_{j \ge 2} \| \hat{A}_j^{(2)} \|_{S_2} \\ &\le C \sqrt{r}\, \tau^{\frac{1}{2-p}} \nu^{\frac{2}{2-p}} \end{aligned}$$

for some constant $C = C(p) > 0$. This and again (39) yield

$$\| \hat{A} - A^* \|_{S_2} \le \| \hat{A}^{(1)} - A^* + \hat{A}_1^{(2)} \|_{S_2} + \sum_{j \ge 2} \| \hat{A}_j^{(2)} \|_{S_2} \le C' \sqrt{r}\, \tau^{\frac{1}{2-p}} \nu^{\frac{2}{2-p}}$$

for some constant $C' = C'(p) > 0$. Thus, we have proved (9) for $q = 2$. Next, using inequalities (3) and (38) we obtain

$$\| \hat{A} - A^* \|_{S_p}^p \le \| \hat{A}^{(1)} - A^* \|_{S_p}^p + \| \hat{A}^{(2)} \|_{S_p}^p \le 4 \| \hat{A}^{(1)} - A^* \|_{S_p}^p .$$

Combining this with (43) and (44) we get (9) for $q = p$. Finally, (9) for arbitrary $q \in [p, 2]$ follows from the norm interpolation formula

$$\| A \|_{S_q}^q \le \| A \|_{S_p}^{\frac{p(2-q)}{2-p}} \| A \|_{S_2}^{\frac{2(q-p)}{2-p}} ,$$

cf. Lemma 10 of the Appendix with $\theta = \frac{p(2-p)}{q(2-q)}$.

7. Proofs of the Lemmas.

Proof of Lemma 1.. First observe that, by the trace duality,

$$\Big| \frac{1}{N} \sum_{i=1}^{N} \xi_i \operatorname{tr}\big(X_i'(\hat{A} - A^*)\big) \Big| \;\le\; \| \hat{A} - A^* \|_{S_1} \| \mathbf{M} \|_{S_\infty} \tag{46}$$

where $\mathbf{M} = \frac{1}{N}\sum_{i=1}^N \xi_i X_i$. Furthermore,

$$\|\mathbf{M}\|_{S_\infty} \;=\; \sup_{\substack{u\in\mathbb{R}^T:\\ |u|_2=1}} |\mathbf{M}u|_2 \;\le\; \sqrt{m}\max_{1\le j\le m}\sup_{\substack{u\in\mathbb{R}^T:\\ |u|_2=1}} \big| u'\bar\eta_j|,$$

with vectors $\bar\eta_j = N^{-1}\sum_{i=1}^N \xi_i X_{i(j,\cdot)}$. Consequently, for any $t>0$,

$$\begin{aligned}\mathbb{P}\Big(\|\mathbf{M}\|_{S_\infty} \;\ge\; t\sqrt{\frac{m\log m}{N}}\Big) \;&\le\; \mathbb{P}\Big(\sqrt{m}\max_{1\le j\le m}|\bar\eta_j|_2 \ge t\sqrt{\frac{m\log m}{N}}\Big)\\ &\le\; m\max_{1\le j\le m}\mathbb{P}\Big(|\bar\eta_j|_2 \ge t\sqrt{\frac{\log m}{N}}\Big).\end{aligned}$$

To proceed with the evaluation of the latter probability we use the following concentration bound (Pinelis and Sakhanenko, 1985).

Lemma 7. *Let $\zeta_1,\dots,\zeta_N$ be independent zero mean random variables in a separable Hilbert space $\mathcal{H}$ such that*

$$\sum_{i=1}^N \mathbb{E}\|\zeta_i\|_{\mathcal{H}}^l \;\le\; \frac{1}{2}l!\,B^2L^{l-2},\quad l=2,3\dots, \tag{47}$$

with some finite constants $B,L>0$. Then

$$\mathbb{P}\Big(\Big\|\sum_{i=1}^N\zeta_i\Big\|_{\mathcal{H}} \;\ge\; x\Big) \;\le\; 2\,\exp\Big(-\frac{x^2}{2B^2+2xL}\Big),\quad \forall\, x>0.$$

Setting $\zeta_i = \xi_i X_{i(j,\cdot)}$, $\mathcal{H}=\mathbb{R}^T$, note first that, by the Bernstein condition (15),

$$\begin{aligned}\sum_{i=1}^N \mathbb{E}\|\zeta_i\|_{\mathcal{H}}^l \;&=\; \mathbb{E}|\xi_i|^l\sum_{i=1}^N |X_{i(j,\cdot)}|_2^l\\ &\le\; \frac{1}{2}l!\,\sigma^2H^{l-2}\Big(\max_j\sum_{i=1}^N|X_{i(j,\cdot)}|_2^2\Big)\max_{i,j}|X_{i(j,\cdot)}|_2^{l-2}\\ &\le\; \frac{1}{2}l!\,B^2L^{l-2},\end{aligned}$$

where $B^2 = \sigma^2S_{row}^2N$ and $L = H_{row}H$, i.e., condition (47) is satisfied. Now an application of Lemma 7 yields for any $t>0$

$$\begin{aligned}\mathbb{P}\Big(|\bar\eta_j|_2 \;\ge\; t\sqrt{\frac{\log m}{N}}\Big) \;&=\; \mathbb{P}\Big(\Big|\frac{1}{\sqrt{N}}\sum_{i=1}^N\xi_i\Big|_2 \;>\; t\sqrt{\log m}\Big)\\ &\le\; 2\exp\Big(-\frac{N(\log m)t^2}{2B^2+2tL\sqrt{N\log m}}\Big)\\ &=\; 2\exp\Big(-\frac{N(\log m)t^2}{2\sigma^2S_{row}^2N+2tL\sqrt{N\log m}}\Big).\end{aligned}$$

Define $t = \sqrt{2D\sigma^2 S_{row}^2} + 2DL\sqrt{\frac{\log m}{N}}$ for some $D > 1$. Then

$$\frac{t^2}{\bar{B} + \bar{L}t} \;\geq\; D, \quad \text{where } \bar{B} = 2\sigma^2 S_{row}^2, \;\; \bar{L} = 2L\sqrt{\frac{\log m}{N}}.$$

With this choice of t,

$$\mathbb{P}\left(|\bar{\eta}_j|_2 \;\geq\; t\sqrt{\frac{\log m}{N}} \right) \;\leq\; 2\exp\big(-D\log m\big) \;=\; 2m^{-D}$$

and therefore $\mathbb{P}\big(\|\mathbf{M}\|_{S_\infty} \geq \tau_{row}\big) \leq 2m^{1-D}$, where

$$\tau_{row} \;=\; \left(\sqrt{2D\sigma^2 S_{row}^2} + 2DH_{row}H\sqrt{\frac{\log m}{N}}\right)\sqrt{\frac{m\log m}{N}}\,.$$

Similarly, using $\|\mathbf{M}\|_{S_\infty} = \sup_{|v|_2=1}|v'\mathbf{M}|_2$, and assuming (18) and (19), we get $\mathbb{P}\big(\|\mathbf{M}\|_{S_\infty} \geq \tau_{col}\big) \leq 2T^{1-D}$, where

$$\tau_{col} \;=\; \left(\sqrt{2D\sigma^2 S_{col}^2} + 2DH_{col}H\sqrt{\frac{\log T}{N}}\right)\sqrt{\frac{T\log T}{N}}\,.$$

□

For the proof of Lemma 2 we will need some notation. The pth Schatten class of $M \times M$-matrices is denoted by S_p^M, and we write

$$\mathcal{B}(S_p^M) \;=\; \big\{A \in \mathbb{R}^{M\times M} : \|A\|_{S_p} \leq 1\big\}$$

for the corresponding closed Schatten-p unit ball in $\mathbb{R}^{M\times M}$. For any pseudo-metric space $(\mathcal{T}, d)$ and any $\varepsilon > 0$ we define the covering number

$$\mathcal{N}\big(\mathcal{T}, d, \varepsilon\big) \;=\; \min\Big\{\mathrm{Card}(\mathcal{T}_0) : \mathcal{T}_0 \subset \mathcal{T} \text{ and } \inf_{s\in\mathcal{T}_0} d(t,s) \leq \varepsilon \text{ for all } t \in \mathcal{T}\Big\}.$$

In other words, $\mathcal{N}\big(\mathcal{T}, d, \varepsilon\big)$ is the smallest number of closed balls of radius ε in the metric d needed to cover the set $\mathcal{T}$. We will sometimes write $\mathcal{N}\big(\mathcal{T}, \|\cdot\|, \varepsilon\big)$ instead of $\mathcal{N}\big(\mathcal{T}, d, \varepsilon\big)$ if the metric d is associated with the norm $\|\cdot\|$. The empirical norm $\|\cdot\|_{2,N}$ corresponds to $\hat{d}_{2,N}$, i.e., for all $A \in \mathbb{R}^{M\times M}$,

$$\|A\|_{2,N}^2 \;=\; \frac{1}{N}\sum_{j=1}^N \mathrm{tr}(A'X_j)^2.$$

Proof of Lemma 2. First note that we cannot rely the proof directly on the trace duality and norm interpolation (cf. Lemma 10), i.e., on the inequalities

$$\begin{aligned} \left|\frac{1}{N}\sum_{i=1}^N \xi_i \mathrm{tr}\big(X_i'(\hat{A} - A^*)\big)\right| \;&\leq\; \big\|\hat{A} - A^*\big\|_{S_1}\big\|\mathbf{M}\big\|_{S_\infty} \\ &\leq\; \big\|\hat{A} - A^*\big\|_{S_2}^{1-\frac{p}{2-p}}\big\|\hat{A} - A^*\big\|_{S_p}^{\frac{p}{2-p}}\big\|\mathbf{M}\big\|_{S_\infty}. \qquad (48) \end{aligned}$$

Indeed, one may think that we could have bounded here the S_∞-norm of $\mathbf{M}$ in the same way as above, and then the proof would be complete if we were able to bound from above $\|\hat{A}-A^*\|_{S_2}^2$ by $\hat{d}_{2,N}(\hat{A},A^*)^2$ times a constant factor. Unfortunately, this is not possible. Even the Restricted Isometry condition cannot help here because $\hat{A}-A^*$ is not necessarily of small rank. Nevertheless, we will show now that by other techniques it is possible to derive an inequality similar to (48) with $\hat{d}_{2,N}(\hat{A},A^*)$ instead of $\|\hat{A}-A^*\|_{S_2}$.

Since

$$\sup_{B\in\mathbb{R}^{M\times M}}\left|\frac{\frac{1}{\sqrt{N}}\sum_{i=1}^N\xi_i\mathrm{tr}(B'X_i)}{\|B\|_{2,N}^{1-\frac{p}{2-p}}\|B\|_{S_p}^{\frac{p}{2-p}}}\right| = \sup_{B\in\mathcal{B}(S_p^M)}\left|\frac{\frac{1}{\sqrt{N}}\sum_{i=1}^N\xi_i\mathrm{tr}(B'X_i)}{\|B\|_{2,N}^{1-\frac{p}{2-p}}}\right|,$$

the expression on the LHS of (21) is not greater than

$$\frac{\sqrt{M}}{\sqrt{p}\sqrt{N}}\,\hat{d}_{2,N}\big(\hat{A},A^*\big)^{1-\frac{p}{2-p}}\,\big\|\,\hat{A}-A^*\big\|_{S_p}^{\frac{p}{2-p}}\sup_{B\in\mathcal{B}(S_p^M)}\left|\frac{(M/p)^{\frac{p-2}{2p}}N^{-1/2}\sum_{i=1}^N\xi_i\mathrm{tr}(B'X_i)}{\big((M/p)^{\frac{p-2}{2p}}\|B\|_{2,N}\big)^{1-\frac{p}{2-p}}}\right|.$$

Due to the linear dependence in M of the ε-entropies of the quasi-convex Schatten class embeddings $S_p^M\hookrightarrow S_2^M$ (cf. Corollary 6) and the fact that the required bound should be uniform in M and in p for $p\searrow 0$, we introduced an additional weighting by $(M/p)^{\frac{p-2}{2p}}$. Now define

$$\mathcal{G}_{M,p} = \Big\{A\in\mathbb{R}^{M\times M}:(M/p)^{\frac{2-p}{2p}}A\in\mathcal{B}(S_p^M)\Big\}.$$

By the entropy bound of Corollary 6 and Assumption 1,

$$\log\mathcal{N}\big(\mathcal{G}_{M,p},\hat{d}_{2,N},\varepsilon\big)\le\log\mathcal{N}\big(\mathcal{G}_{M,p},\sqrt{c_0}\,\|\cdot\|_{S_2},\varepsilon\big)\le p\,\alpha_0(p)\big(\varepsilon/\sqrt{c_0}\big)^{-\frac{2p}{2-p}},$$

whence

$$\int_0^\delta\sqrt{\log\mathcal{N}\big(\mathcal{G}_{M,p},\hat{d}_{2,N},\varepsilon\big)}\,d\varepsilon\le c_0^{\frac{p}{2(2-p)}}p\,\alpha_0(p)\frac{2-p}{2-2p}\,\delta^{1-\frac{p}{2-p}}. \tag{49}$$

We remark that due to the order specification of α_0 in Corollary 6, the expression

$$c_0^{\frac{p}{2(2-p)}}p\,\alpha_0(p)\frac{2-p}{2-2p} \tag{50}$$

is uniformly bounded as long as p stays uniformly bounded away from 1. Note that for $p=1$ the entropy integral on the LHS in (49) does not converge.

CLAIM 1. *For any $q\in(0,1)$, there exist constants $c(q)$ and $c'(q)$, such that for all $0<p\le q$, all $0<\delta\le\sqrt{c_0}$ and uniformly in M and N:*

$$\mathbb{P}\Bigg(\sup_{\substack{B\in\mathcal{G}_{M,p}:\\ \|B\|_{2,N}\le\delta}}\Big|\frac{1}{\sqrt{N}}\sum_{j=1}^N\xi_j\mathrm{tr}(X_j'B)\Big|\ge T\Bigg)\le c(q)\exp\Big(-\frac{T^2}{c(q)^2\delta^2}\Big) \tag{51}$$

for all $T\ge c'(q)\delta^{1-\frac{p}{2-p}}$.

Proof of Claim 1. The bound is essentially stated in van de Geer (2000) as Lemma 3.2 (further referred to as VG(00)). The constant in VG(00) depends neither on the $\|\cdot\|_{2,N}$-diameter of the function class nor on the function class itself and is valid, in particular, for $\varepsilon = 0$, in the notation of VG(00). The uniformity in $0 < p \le q$ follows from the uniform boundedness of (50) over $p \in (0, q]$. The required case corresponds to $K = \infty$ in the notation of VG(00). Its proof follows by taking $\varepsilon = 0$ and applying the theorem of monotone convergence as $K \to \infty$, since the RHS of the inequality is independent of K.

CLAIM 2. *For any $q \in (0,1)$, there exists a constant $C(q)$ such that for any $0 < p \le q$*

$$\mathbb{P}\left(\sup_{B\in\mathcal{G}_{M,p}} \left| \frac{\frac{1}{\sqrt{N}}\sum_{j=1}^{N} \xi_j \operatorname{tr}(B'X_j)}{\|B\|_{2,N}^{1-\frac{p}{2-p}}} \right| \;\ge\; T\right) \;\le\; C(q)\exp\Big(-T^2 M/C(q)^2\Big) \tag{52}$$

for all $T \ge C(q)$.

Proof of Claim 2. First observe that

$$\sup_{A\in\mathcal{G}_{M,p}} \|A\|_{2,N} \le \sqrt{c_0} \sup_{A\in\mathcal{G}_{M,p}} \|A\|_{S_2} \;\le\; \sqrt{c_0}(M/p)^{\frac{p-2}{2p}} \sup_{A\in\mathcal{B}(S_2^M)} \|A\|_{S_2} = \sqrt{c_0}(M/p)^{\frac{p-2}{2p}},$$

where the last inequality follows from $\mathcal{B}(S_p^M) \subset \mathcal{B}(S_2^M)$. Define the decomposition of $\mathcal{G}_{M,p}$

$$\mathcal{G}_{M,p}^{(k)} \;=\; \Big\{A \in \mathcal{G}_{M,p} :\, 2^{k\frac{p-2}{2p}}\sqrt{c_0}(M/p)^{\frac{p-2}{2p}} \;\le\; \|A\|_{2,N} \;\le\; 2^{(k-1)\frac{p-2}{2p}}\sqrt{c_0}(M/p)^{\frac{p-2}{2p}}\Big\},\;\; k \in \mathbb{N}.$$

Then by straightforward peeling-off the class $\mathcal{G}_{M,p}$, we obtain for all $T \ge c'(q)$

$$\begin{aligned}
&\mathbb{P}\left(\sup_{B\in\mathcal{G}_{M,p}} \left| \frac{\frac{1}{\sqrt{N}}\sum_{j=1}^{N} \xi_j \operatorname{tr}(B'X_j)}{\|B\|_{2,N}^{1-\frac{p}{2-p}}} \right| \;\ge\; T\right)\\
&\qquad\le\; \sum_{k=1}^{\infty} \mathbb{P}\left(\sup_{B\in\mathcal{G}_{M,p}^{(k)}} \left| \frac{1}{\sqrt{N}}\sum_{j=1}^{N} \xi_j \operatorname{tr}(B'X_j)\right| \ge T\big(2^{k\frac{p-2}{2p}}\sqrt{c_0}(M/p)^{\frac{p-2}{2p}}\big)^{1-\frac{p}{2-p}}\right)\\
&\qquad\le\; \sum_{k=1}^{\infty} c(q)\exp\left(-\frac{T^2\big(2^{k\frac{p-2}{2p}}\sqrt{c_0}(M/p)^{\frac{p-2}{2p}}\big)^{-\frac{2p}{2-p}}}{c(q)^2}\right)\\
&\qquad\le\; \sum_{k=1}^{\infty} c(q)\exp\left(-\frac{T^2 M 2^k C_0(q)}{qc(q)^2}\right) \qquad (53)
\end{aligned}$$

with the definition

$$C_0(q) \;=\; \inf_{0<p\le q} c_0^{-\frac{p}{2-p}}.$$

It remains to note that the last sum in (53) is bounded by $C(q)\exp\Big(-\frac{T^2 M}{C(q)^2}\Big)$ whenever $T \ge C(q)$ for some suitable constant $C(q)$.

In particular, the result reveals that the LHS of (21) is bounded by

$$\hat{d}_{2,N}\big(\hat{A}, A^*\big)^{1-\frac{p}{2-p}} \,\|\,\hat{A} - A^*\,\|_{S_p}^{\frac{p}{2-p}} \sqrt{\vartheta/p}\Big(\frac{M}{N}\Big)^{1/2} \tag{54}$$

with probability at least $1 - C\,\exp(-\vartheta M/C^2)$ for any $\sqrt{\vartheta} \geq C(q)$.

We now use the following simple consequence of the concavity of the logarithm which is stated, for instance, in Tsybakov and van de Geer (2005) (Lemma 5).

Lemma 8. *For any positive v, t and any $\kappa \geq 1$, $\delta > 0$ we have*

$$vt^{1/(2\kappa)} \;\leq\; (\delta/2)t + c_\kappa \delta^{-1/(2\kappa-1)} v^{2\kappa/(2\kappa-1)},$$

where $c_\kappa = (2\kappa-1)(2\kappa)\kappa^{-1/(2\kappa-1)}$.

Taking in Lemma 8

$$t = \hat{d}_{2,N}(\hat{A}, A^*)^2, \quad v = \big\|\hat{A} - A^*\big\|_{S_p}^{\frac{p}{2-p}} \sqrt{\vartheta/p}\Big(\frac{M}{N}\Big)^{1/2},$$

and $\kappa = (2-p)/(2-2p)$ shows that for any $\delta > 0$

$$(54) \;\leq\; (\delta/2)\,\hat{d}_{2,N}(\hat{A}, A^*)^2 \;+\; \tau\delta^{p-1}\big\|\hat{A} - A^*\big\|_{S_p}^{p}$$

with probability at least $1 - C\,\exp(-\vartheta M/C^2)$. □

Proof of Lemma 3. In view of (46), it suffices to evaluate the norm $\big\|\mathbf{M}\big\|_{S_\infty}$ of the matrix $\mathbf{M} = \frac{1}{N}\sum_{i=1}^N \xi_i X_i$. This is a random matrix with non-i.i.d. entries, so the standard results on maximal eigenvalues of subgaussian matrices (cf. Mendelson et al. (2007), Vershynin (2007)) do not apply. However, due to the specific structure of the matrix, we can adapt the argument to our case. Note first that

$$\big\|\mathbf{M}\big\|_{S_\infty} = \max_{u\in\mathcal{S}^{T-1}} |\mathbf{M}u|_2 = \max_{v\in\mathcal{S}^{m-1}, u\in\mathcal{S}^{T-1}} v'\mathbf{M}u,$$

where $\mathcal{S}^{m-1}$ is the unit sphere in $\mathbb{R}^m$. Therefore, denoting by $\mathcal{M}_m$ and $\mathcal{M}_T$ the minimal 1/2-nets in Euclidean metric on $\mathcal{S}^{m-1}$ and $\mathcal{S}^{T-1}$ respectively, we easily get

$$\big\|\mathbf{M}\big\|_{S_\infty} \leq 2 \max_{u\in\mathcal{M}_T} |\mathbf{M}u|_2 \leq 4 \max_{v\in\mathcal{M}_m, u\in\mathcal{M}_T} |v'\mathbf{M}u|.$$

Now, $\mathrm{Card}(\mathcal{M}_m) \leq 5^m$, cf. Kolmogorov and Tikhomirov (1959), so that by the union bound, for any $t > 0$,

$$\mathbb{P}\Big(\|\mathbf{M}\|_{S_\infty} \geq \tau\Big) \;\leq\; 5^{m+T} \max_{v\in\mathcal{M}_m, u\in\mathcal{M}_T} \mathbb{P}\Big(\big|v'\mathbf{M}u\big| \geq \tau/4\Big). \tag{55}$$

It remains to bound the last probability in (55) for fixed u, v. Let us fix some $v \in \mathcal{S}^{m-1}, u \in \mathcal{S}^{T-1}$ and introduce the random event

$$\mathcal{A} = \left\{\frac{1}{N}\sum_{i=1}^N (u'X_i v)^2 \leq \frac{5(m+T)}{N}\right\}.$$

Note that $\mathbb{E}(u'X_i v)^2 = \sum_{k=1}^m \sum_{l=1}^T u_k^2 v_l^2 \mathbb{P}(X_1 = e_k(m)e_l'(T)) = (mT)^{-1}|u|_2^2|v|_2^2 = (mT)^{-1}$, and consider the zero-mean random variables $\eta_i = (u'X_i v)^2 - \mathbb{E}(u'X_i v)^2 =$

$(u'X_iv)^2 - (mT)^{-1}$. We have $|\eta_i| \le 2\max_i (u'X_iv)^2 \le 2|u|_2^2|v|_2^2 = 2$ (a.s.). Furthermore,

$$\mathbb{E}(\eta_i^2) \le \mathbb{E}(u'X_iv)^4 \le \sum_{k=1}^{m}\sum_{l=1}^{T} u_k^4 v_l^4 \mathbb{P}(X_1 = e_k(m)e_l'(T))$$
$$= (mT)^{-1}\sum_{k=1}^{m} u_k^4 \sum_{l=1}^{T} v_l^4 \le (mT)^{-1}.$$

Therefore, using Bernstein's inequality and the condition $(m+T)/N > (mT)^{-1}$ we get

$$\mathbb{P}(\mathcal{A}^c) \le 2\exp\left(-\frac{N(4(m+T)/N)^2}{2(mT)^{-1} + (4/3)(4(m+T)/N)}\right) \le 2\exp(-2(m+T)), \tag{56}$$

where $\mathcal{A}^c$ is the complement of $\mathcal{A}$. We now bound the conditional probability

$$\mathbb{P}\left(\left|v'\mathbf{M}u\right| \ge \tau/4 \,\middle|\, X_1,\dots,X_N\right) = \mathbb{P}\left(\left|\frac{1}{N}\sum_{i=1}^{N}\xi_i(u'X_iv)\right| \ge \tau/4 \,\middle|\, X_1,\dots,X_N\right).$$

Note that conditionally on $X_1,\dots,X_N$, the $\xi_i(u'X_iv)$ are independent zero-mean random variables with

$$\sum_{i=1}^{N}\mathbb{E}\left(|\xi_i(u'X_iv)|^l \,\middle|\, X_1,\dots,X_N\right) \le \mathbb{E}|\xi_1|^l \sum_{i=1}^{N}|u'X_iv|^2, \quad \forall\, l \ge 2,$$

where we used the fact that $|u'X_iv|^{l-2} \le (|u|_2|v|_2)^{l-2} = 1$ (a.s.) for $l \ge 2$. This and the Bernstein condition (15) yield that, for $(X_1,\dots,X_N) \in \mathcal{A}$,

$$\sum_{i=1}^{N}\mathbb{E}\left(|\xi_i(u'X_iv)|^l \,\middle|\, X_1,\dots,X_N\right) \le \frac{l!}{2}B^2H^{l-2}.$$

with $B^2 = 5(m+T)\sigma^2$. Therefore, by Lemma 7, for $(X_1,\dots,X_N) \in \mathcal{A}$ we have

$$\mathbb{P}\left(\left|v'\mathbf{M}u\right| \ge \tau/4 \,\middle|\, X_1,\dots,X_N\right) \le 2\exp\left(-\frac{N^2\tau^2/16}{10\sigma^2(m+T) + N\tau H/2}\right). \tag{57}$$

For τ defined in (22) the last expression does not exceed $2\exp(-D(m+T))$. Together with (55) and (56), this proves the lemma. □

PROOF OF LEMMA 4. We act as in the proof of Lemma 3 but since the matrices X_i are now deterministic, we do not need to introduce the event $\mathcal{A}$. By Assumption 1,

$$\frac{1}{N}\sum_{i=1}^{N}(u'X_iv)^2 \equiv |\mathcal{L}(vu')|_2^2 \le c_0||vu'||_{S_2}^2 = c_0$$

for all $v \in \mathcal{S}^{m-1}, u \in \mathcal{S}^{T-1}$. Hence, $\frac{1}{N}\sum_{i=1}^{N}\xi_i(u'X_iv)$ is a zero-mean Gaussian random variable with variance not larger than $c_0\sigma^2/N$. Therefore,

$$\mathbb{P}\left(\left|v'\mathbf{M}u\right| \ge \tau/4\right) \le 2\exp\left(-\frac{N\tau^2}{32c_0\sigma^2}\right).$$

For τ as in (23) the last expression does not exceed $2\exp(-2(m+T))$. Combining this with (55) we get the lemma. □

Appendix. Here we derive bounds for the kth entropy numbers of the embeddings $S_p^M \hookrightarrow S_2^M$ for $0 < p < 1$, where S_p^M denotes the pth Schatten class of real $M \times M$-matrices. Corresponding results for the $l_p^M \hookrightarrow l_2^M$-embeddings are given first by Edmunds and Triebel (1989) but their proof does not carry over to the Schatten spaces. Pajor (1998) provides bounds for the $S_p^M \hookrightarrow S_2^M$ embeddings in the convex case, $p \geq 1$. His approach is based on the trace duality (Hölder inequality for $p^{-1} + q^{-1} = 1$) and the geometric formulation of Sudakov's minoration

$$\varepsilon \sqrt{\log \mathcal{N}(A, |\cdot|_2, \varepsilon)} \;\leq\; c\, \mathbb{E} \sup_{t \in A} \langle G, t \rangle$$

with a d-dimensional standard Gaussian vector and an arbitrary subset A of $\mathbb{R}^d$. Here $|\cdot|_2$ is the Euclidean norm in $\mathbb{R}^d$ and $\langle \cdot, \cdot \rangle$ is the corresponding scalar product. Guédon and Litvak (2000) derive a slightly sharper bound for the $l_p \hookrightarrow l_q$-embeddings than Edmunds and Triebel (1989) with a different technique. In addition, they prove lower bounds. We adjust their ideas concerning finite ℓ_p spaces to the non-convex Schatten spaces.

We denote by $e_k\big(id_{p,r}^M\big)$ the kth entropy number of the embedding $S_p^M \hookrightarrow S_r^M$ for $0 < p < r < \infty$, i.e., the infimum of all $\varepsilon > 0$ such that there exist 2^{k-1} balls in S_r^M of radius ε that cover $\mathcal{B}(S_p^M)$. For the general definition of kth entropy numbers $e_k(T : F \to E)$ of bounded linear operators T between quasi-Banach spaces F and E we refer to Edmunds and Triebel (1996).

Recall that a homogeneous non-negative functional $\|\cdot\|$ is called C-quasi-norm, if it satisfies for all x, y the inequality $\|x+y\| \leq C \max\big(\|x\|, \|y\|\big)$. Finally, any p-norm is a C-quasi-norm with $C = 2^{1/p}$ (cf., e.g., Edmunds and Triebel 1996, page 2). We will use the following lemma.

Lemma 9 (Guédon and Litvak, 2000). *Assume that* $\|\cdot\|_i$ *are symmetric* C_i-*quasi-norms on* $\mathbb{R}^n$ *for* $i = 0, 1$, *and for some* $\theta \in (0,1)$, $\|\cdot\|_\theta$ *is a quasi-norm on* $\mathbb{R}^n$ *such that* $\|x\|_\theta \leq \|x\|_0^\theta \|x\|_1^{1-\theta}$ *for all* $x \in \mathbb{R}^n$. *Then for any quasi-normed space* F, *any linear operator* $T : F \to \mathbb{R}^n$, *and all integers* k *and* m, *we have*

$$e_{m+k-1}\big(T : F \to E_\theta\big) \;\leq\; \Big(C_0 e_m\big(T : F \to E_0\big)\Big)^\theta \Big(C_1 e_k\big(T : F \to E_1\big)\Big)^{1-\theta}.$$

where E_t *stands for* $\mathbb{R}^n$ *equipped with quasi-norm* $\|\cdot\|_t$, $t \in \{0, \theta, 1\}$.

Recall also the following simple fact.

Lemma 10 (Interpolation inequality). *For* $0 < p < q < r < \infty$ *let* $\theta \in [0,1]$ *be such that*

$$\frac{\theta}{p} + \frac{1-\theta}{r} \;= \frac{1}{q}.$$

Then, for all $A \in \mathbb{R}^{m \times T}$,

$$\big\| A \big\|_{S_q} \;\leq\; \big\| A \big\|_{S_p}^{\theta} \big\| A \big\|_{S_r}^{1-\theta}.$$

Proof is immediate in view of the inequalities

$$\sum_j a_j^q = \sum_j a_j^{\theta q} a_j^{(1-\theta)q} \leq \Big(\sum_j a_j^p\Big)^{\frac{\theta q}{p}} \Big(\sum_j a_j^r\Big)^{\frac{(1-\theta)q}{r}}$$

valid for any non-negative a_j's.

PROPOSITION 1 (Entropy numbers). *Let $0 < p < 1$, $p < r \le \infty$. Then there exists an absolute constant β independent of p and r, such that for all integers k and M we have*

$$e_k\big(id^M_{p,r}\big) \;\le\; \min\left\{1,\, \alpha(\beta,p,r)\Big(\frac{M}{k}\Big)^{1/p-1/r}\right\}$$

with

$$\alpha(\beta,p,r) \;\le\; 2^{1+1/r}\Big(\frac{\beta}{p}\Big)^{1/p-1/r}\Big(\frac{1}{1-p}\Big)^{(1/p-1)(1/p-1/r)}.$$

PROOF. The fact that $e_k\big(id^M_{p,r}\big)$ is bounded by 1 is obvious, since $\mathcal{B}(S^M_p) \subset \mathcal{B}(S^M_r)$. Consider the other case. We start with $r = \infty$ and then extend the result to $r < \infty$ by interpolation. Fix some number $L > M$ and let $D = D(M, L, p)$ be the smallest constant which satisfies, for all $1 \le k \le L$,

$$e_k\big(id^M_{p,\infty}\big) \;\le\; D\bigg(\frac{M}{k}\bigg)^{1/p}. \tag{58}$$

Let us show that $\alpha = \sup_{M,L} D(M, L, p)$ is finite. Since $\|\cdot\|_{S_p}$, $p < 1$, can be viewed as a quasi-norm on $\mathbb{R}^{M^2}$ (isomorphic to $\mathbb{R}^{M\times M}$), Lemma 9 applies with $F = E_0 = S^M_p$, $E_1 = S^M_\infty$, $\theta = p$, $E_\theta = S^M_1$ and $m = 1$. This gives

$$e_k\big(id^M_{p,1}\big) \;\le\; 4\,\Big(e_k\big(id^M_{p,\infty}\big)\Big)^{1-p}. \tag{59}$$

Here the factor 4 follows from the relations $C_1 = 2$ and $C_0^p \le 2$. Now, (59) and the factorization theorem for entropy numbers of bounded linear operators between quasi-Banach spaces (see, e.g., Edmunds and Triebel 1996, page 8), with factorization via S^M_1, leads to the bound

$$\begin{aligned} e_k\big(id^M_{p,\infty}\big) \;&\le\; e_{[(1-p)k]}\big(id^M_{p,1}\big)e_{[pk]}\big(id^M_{1,\infty}\big) \\ &\le\; 4\Big(e_{[(1-p)k]}\big(id^M_{p,\infty}\big)\Big)^{1-p} e_{[pk]}\big(id^M_{1,\infty}\big), \end{aligned} \tag{60}$$

where for any $x \in (0, \infty)$, $[x]$ denotes the smallest integer which is larger or equal to x. Proposition 5 of Pajor (1998) entails

$$\log\mathcal{N}\Big(\mathcal{B}(S^M_1), \|\cdot\|_{S_\infty}, \varepsilon\Big) \;\le\; c\,M/\varepsilon, \quad \forall\,\varepsilon > 0,$$

and hence

$$e_k\big(id^M_{1,\infty}\big) \;\le\; c'\,M/k \tag{61}$$

with constants c and c' independent of M, ε and k. Note that, in contrast to the $l^M_1 \hookrightarrow l^M_\infty$-embedding, for which the k'th entropy numbers are bounded by $c''k^{-1}\log\big(1 + M/k\big)$ with some $c'' > 0$ and $\log_2 M \le k \le M$ (see, e.g., Edmunds and Triebel 1996, page 98), we have in (61) not a logarithmic but linear dependence of M in the upper bound. Plugging (58) and (61) into (60) yields

$$\begin{aligned} e_k\big(id^M_{p,\infty}\big) \;&\le\; 4\bigg(D\Big(\frac{M}{(1-p)k}\Big)^{1/p}\bigg)^{1-p}\frac{c'M}{pk} \\ &=\; \frac{4\,c'}{p}\Big(\frac{1}{1-p}\Big)^{(1-p)/p}D^{1-p}\Big(\frac{M}{k}\Big)^{1/p}. \end{aligned}$$

Thus, by definition of D,

$$D^p \;\leq\; \frac{4\,c'}{p}\Big(\frac{1}{1-p}\Big)^{(1-p)/p},$$

which shows that D is uniformly bounded in M and L. This proves the proposition for $r = \infty$.

Consider now the case $r < \infty$. In view of Lemma 10 with $\theta = p/r$, we can apply Lemma 9 with $F = E_0 = S_p^M$, $E_1 = S_\infty^M$, $\theta = p/r$, $E_\theta = S_r^M$ and $m = 1$. This yields

$$\begin{aligned} e_k\big(id^M_{p,r}\big) \;&\leq\; 2^{1+1/r}\Big(e_k\big(id^M_{p,\infty}\big)\Big)^{1-p/r} \\ &\leq\; 2^{1+1/r} D^{1-p/r}\Big(\frac{M}{k}\Big)^{1/p-1/r}. \end{aligned}$$

□

Corollary 6. *For any $p \in (0,1)$, there exists a positive constant $\alpha_0(p)$ such that for all integers $M \geq 1$ and any $\varepsilon > 0$,*

$$\log \mathcal{N}\big(\mathcal{B}(S_p^M), \|\cdot\|_{S_2}, \varepsilon\big) \;\leq\; \alpha_0(p)\, M\, \varepsilon^{-\frac{2p}{2-p}}.$$

Moreover, $\alpha_0(p) = O(1/p)$ for $p \searrow 0$.

Proof. The result follows by transforming the entropy number bound of Proposition 1 into an entropy bound. Specification of the constant in Proposition 1 yields

$$\alpha_0(p) \;=\; O\bigg(\frac{\beta}{p}\Big(1 + \frac{p}{1-p}\Big)^{(1-p)/p}\bigg) = O(1/p)$$

as $p \searrow 0$. □

References.

[1] Abernethy,J., Bach,F., Evgeniou,T. and Vert, J.-P. (2009). A New Approach to Collaborative Filtering: Operator Estimation with Spectral Regularization. *Journal of Machine Learning Research,* ***10*** 803–826.

[2] Amini, A. and Wainwright, M. (2009) High-dimensional analysis of semidefinite relaxations for sparse principal components. *Annals of Statistics,* ***37***, 2877–2921.

[3] Argyriou, A., Evgeniou, T. and Pontil, M. (2008) Convex multi-task feature learning. *Machine Learning,* ***73***, 243–272.

[4] Argyriou, A., Micchelli, C.A., and Pontil, M. (2008) On spectral learning. *Journal of Machine Learning Research*, to appear.

[5] Argyriou, A., Micchelli, C.A., Pontil, M., and Ying, Y. (2007) A Spectral Regularization Framework for Multi-Task Structure Learning. *Proceedings of NIPS-2007.*

[6] Bach, F.R. (2008). Consistency of trace norm minimization. *Journal of Machine Learning Research,* ***9***, 1019–1048.

[7] Bickel, P.J., Levina, E. (2008) Regularized estimation of large covariance matrices. *Annals of Statistics,* ***36***, 199–227.

[8] Bickel, P.J., Ritov, Y. and Tsybakov, A.B. (2009) Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics,* ***37***, 1705–1732.

[9] Bunea, F., Tsybakov, A.B. and Wegkamp, M.H. (2007) Aggregation for Gaussian regression. *Annals of Statistics* ***35***, 1674–1697.

[10] Candès, E. J. and Plan, Y. (2009) Matrix completion with noise. Submitted.

[11] Candès, E. J. and Recht, B. (2008) Exact matrix completion via convex optimization. *Found. of Comput. Math.*, to appear.

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

[13] Candès, E. J. and Tao,T. (2009) The power of convex relaxation: Near-optimal matrix completion. Submitted.

[14] Dalalyan, A. and Tsybakov, A. (2008). Aggregation by exponential weighting, sharp oracle inequalities and sparsity. *Machine Learning* ***72*** 39–61.

[15] Edmunds, D.E. and Triebel, H. (1996) *Function spaces, Entropy numbers, Differential operators.* Cambridge University Press, Cambridge.

[16] Edmunds, D.E. and Triebel, H. (1989) Entropy numbers and approximation numbers in function spaces. *Proc. London Math. Soc.* ***58***, 137–152.

[17] Foster, D. P. and George, E. I. (1994) The risk inflation criterion for multiple regression *Annals of Statistics* ***22*** 1947–1975.

[18] Gohberg, I.C. and Krein, M.G. (1969) *Introduction to the Theory of Linear Nonselfadjoint Operators.* Amer. Math. Soc., Providence, RI.

[19] Guédon, O. and Litvak, A. E. (2000). Euclidean projections of a p-convex body. Geometric aspects of functional analysis, Israel Seminar (GAFA) 1996–2000, V.D. Milman and G. Schechtman, eds., *Lecutre Notes in Mathematics* ***1745***, 95–108, Springer 2000.

[20] Kolmogorov, A.N. and Tikhomirov,V.M. (1959) The ϵ-entropy and ϵ-capacity of sets in function spaces *Uspekhi Matem. Nauk,* ***14***, 3–86.

[21] Koltchinskii, V. (2008) Oracle inequalities in empirical risk minimization and sparse recovery problems. Saint-Flour Lectures Notes.

[22] Lounici, K., Pontil, M., Tsybakov, A.B. and van de Geer, S. (2009) Taking advantage of sparsity in multi-task learning. *Proceedings of COLT-2009.*

[23] Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics* **34** 1436–1462.

[24] Mendelson, S., Pajor, A. and Tomczak-Jaegermann, N. (2007) Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis* ***17***, 1248–1282.

[25] Pajor, A. (1998). Metric entropy of the Grassmann manifold. *Convex Geometric Analysis* ***34***, 181–188.

[26] Pinelis, I.F. and Sakhanenko, A.I. (1985) Remarks on inequalities for the probabilities of large deviations. *Theory of Prob. Appl.* ***30***, 143–148.

[27] Ravikumar, P., Wainwright, M., Raskutti, G. and Yu, B. (2008) High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. Submitted.

[28] Recht, B, Fazel, M. and Parrilo,P.A. (2007) Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization. `arxiv:0706.4138`

[29] Tsybakov, A. and van de Geer, S. (2005). Square root penalty: adaptation to the margin in classification and in edge estimation *Ann. Statist.* ***33***, 1203–1224.

[30] van de Geer, S. (2000) *Empirical Processes in M-estimation.* Cambridge University Press, Cambridge.

[31] Vershynin, R. (2007) Some problems in asymptotic convex geometry and random matrices motivated by numerical algorithms. *Banach spaces and their applications in analysis*, 209–218, Walter de Gruyter, Berlin.

[32] Zhao, P. and Yu, B. (2006). On model selection consistency of Lasso. *Journal of Machine Learning Research* **7** 2541–2563.

Universität Hamburg
Department Mathematik
Bundesstrasse 55
D-20146 Hamburg
Germany
E-mail: angelika.rohde@math.uni-hamburg.de

Laboratoire de Statistique, CREST
3, avenue Pierre Larousse
F-92240 Malakoff Cedex
France
E-mail: alexandre.tsybakov@upmc.fr