



HAL
open science

Learning linear dynamical systems under convex constraints

Hemant Tyagi, Denis Efimov

► **To cite this version:**

Hemant Tyagi, Denis Efimov. Learning linear dynamical systems under convex constraints. 2024. <hal-04394297>

HAL Id: hal-04394297

<https://hal.science/hal-04394297v1>

Preprint submitted on 15 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Learning linear dynamical systems under convex constraints

Hemant Tyagi
hemant.tyagi@inria.fr

Denis Efimov
denis.efimov@inria.fr

Inria, Univ. Lille, CNRS, UMR 8524 - Laboratoire Paul Painlevé, F-59000

August 2, 2023

Abstract

We consider the problem of finite-time identification of linear dynamical systems from T samples of a single trajectory. Recent results have predominantly focused on the setup where no structural assumption is made on the system matrix $A^* \in \mathbb{R}^{n \times n}$, and have consequently analyzed the ordinary least squares (OLS) estimator in detail. We assume prior structural information on A^* is available, which can be captured in the form of a convex set \mathcal{K} containing A^* . For the solution of the ensuing constrained least squares estimator, we derive non-asymptotic error bounds in the Frobenius norm that depend on the local size of \mathcal{K} at A^* . To illustrate the usefulness of these results, we instantiate them for three examples, namely when (i) A^* is sparse and \mathcal{K} is a suitably scaled ℓ_1 ball; (ii) \mathcal{K} is a subspace; (iii) \mathcal{K} consists of matrices each of which is formed by sampling a bivariate convex function on a uniform $n \times n$ grid (convex regression). In all these situations, we show that A^* can be reliably estimated for values of T much smaller than what is needed for the unconstrained setting.

1 Introduction

We consider the problem of finite-time identification of a linear dynamical system (LDS) of the form

$$x_{t+1} = A^* x_t + \eta_{t+1} \quad \text{for } t = 0, 1, \dots, T \quad \text{and } x_0 = 0, \quad (1.1)$$

where $A^* \in \mathbb{R}^{n \times n}$ is the unknown system matrix to be estimated, $x_t \in \mathbb{R}^n$ is the observed state at time t , and $\eta_t \in \mathbb{R}^n$ is the unobserved (random) process noise. Such problems arise in many areas such as control theory, reinforcement learning and time-series analysis to name a few. An important line of research in recent years has focused on theoretically analyzing the performance of the ordinary least squares (OLS) estimator, by deriving non-asymptotic error bounds for the estimation of A^* (e.g., [33, 34, 31, 18]), holding with high probability provided T is sufficiently large. The analyses depend crucially on the spectrum of A^* – in particular on the spectral radius of A^* , namely $\rho(A^*)$.

The focus of this paper is the strictly stable setting where $\rho(A^*) < 1$. Denoting $\Gamma_s(A) = \sum_{k=0}^s A^k (A^k)^\top$ for $s \geq 0$ to be the controllability Grammian of the system, and $\lambda_{\min}(\cdot)$ to be the smallest eigenvalue of a symmetric matrix, it was shown recently [18] that the OLS estimate \hat{A} satisfies with probability at least $1 - \delta$

$$\|\hat{A} - A^*\|_2 \lesssim \sqrt{\frac{\log(1/\delta) + n}{\lambda_{\min}(\sum_{s=0}^{T-1} \Gamma_s(A^*))}},$$

provided $\lambda_{\min}(\sum_{s=0}^{T-1} \Gamma_s(A^*)) \gtrsim J^2(A^*)(\log(1/\delta) + n)$. Here $\|\cdot\|_2$ denotes the spectral norm and $(\eta_t)_{t \geq 1}$ are considered to be i.i.d subgaussian vectors – see Section 2 for a description of notations. The quantity $J(A^*)$ is defined in (2.1) and is finite when $\rho(A^*) < 1$; it is moreover bounded by a constant if $\|A^*\|_2 < c < 1$ for some constant c . Since $\lambda_{\min}(\sum_{s=0}^{T-1} \Gamma_s(A^*)) \geq T$, we can rewrite the above bound as

$$\|\hat{A} - A^*\|_2 \lesssim \sqrt{\frac{\log(1/\delta) + n}{T}} \quad \text{if} \quad T \gtrsim J^2(A^*)(\log(1/\delta) + n). \quad (1.2)$$

In other words, a meaningful error bound is ensured provided the length of the trajectory is at least of the order of the dimension n . Furthermore, this bound is also optimal in terms of dependence on δ, n and T [34].

It is natural to consider the scenario where additional structural information is available regarding A^* – in this case one would expect that incorporating such information in the estimation procedure should lead to an improved performance compared to the vanilla OLS estimator. In many cases of interest, A^* actually has an intrinsically low dimensional structure and it is possible to capture this structural information of A^* through a known convex set \mathcal{K} containing A^* . Computationally, the estimate \hat{A} is then obtained by the constrained least squares estimator

$$\hat{A} \in \operatorname{argmin}_{A \in \mathcal{K}} \sum_{t=0}^T \|x_{t+1} - Ax_t\|_2^2 \quad (1.3)$$

which is also a convex program that can typically be solved efficiently in practice. From a statistical perspective, one would expect to be able to improve the error bounds in (1.2) in terms of the dependence on the (extrinsic) dimension n . Some examples of such \mathcal{K} – which will also be used later for instantiating our more general results – are outlined below.

1. **(Example 1)** \mathcal{K} is a d -dimensional subspace of $\mathbb{R}^{n \times n}$ for some $d \leq n^2$. This was considered recently in [42] which cited applications of this setup in the time series analysis of spatio-temporal studies and social networks.
2. **(Example 2)** If A^* is k -sparse, i.e., has k non-zero entries, then one can choose \mathcal{K} to be a suitably scaled ℓ_1 ball such that $A^* \in \mathcal{K}$. Such assumptions exist in the literature for model (1.1) as will be discussed in Section 1.2. It is well known in the statistics and signal processing communities that the resulting estimator – known as the LASSO – promotes solutions which are sparse (see, e.g., [5, 25, 6]).
3. **(Example 3)** In this case, we consider A^* to be formed by sampling an unknown convex function $f_0 : [0, a]^2 \rightarrow \mathbb{R}$ on a uniform grid. Then, the set \mathcal{K} can be written as

$$\mathcal{K} = \left\{ A \in \mathbb{R}^{n \times n} : A_{ij} = f\left(\frac{i-1}{n-1}, \frac{j-1}{n-1}\right) \text{ for some convex } f : [0, a]^2 \rightarrow \mathbb{R} \right\}.$$

This problem is similar to the problem of convex-regression that has been studied extensively in the statistics community in the i.i.d setting, see, e.g., [17, 21, 11] and references therein. From a computational perspective, it is useful to know that \mathcal{K} can be characterized as a convex cone [32, Lemma 2.2] which thus leads to a convex quadratic program that can be solved efficiently through standard solvers in practice; there also exist efficient specialized solvers for convex regression (e.g., [10]).

In examples 1, 2, the intrinsic dimension of A^* is essentially captured by the quantities d or k , and we expect that the error bounds in (1.2) should improve in terms of exhibiting a milder dependence on n . In particular, when $d, k \ll n^2$, we expect the estimation error for A^* to be small for moderately large values of T . While this is less obvious for example 3, we will see later that a similar conclusion can be drawn if f_0 is piecewise affine with a ‘simple structure’.

1.1 Our contributions

For the setting where $\rho(A^*) < 1$ and $A^* \in \mathcal{K}$, we derive non-asymptotic bounds on the estimation error in the Frobenius norm $\|\hat{A} - A^*\|_F$ for the estimator (1.3), holding with high probability; see Theorems 1 and 2 for the full statements. Our bounds depend on the ‘local size’ of the set \mathcal{K} at A^* , captured via Talagrand’s γ_1, γ_2 functionals [37] (see Definitions 1 and 2 in Section 2). Upon instantiating our bounds for the aforementioned choices of \mathcal{K} , we obtain the following corollaries.

1. (**Example 1**) In this case, Theorem 1 states (see Corollary 1) that with probability at least $1 - \delta$,

$$\|\hat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + \sqrt{d}}{\sqrt{T}} \right) \quad \text{if } T \gtrsim J^4(A^*) \max \{d, \log^2(1/\delta)\}. \quad (1.4)$$

Suppose for simplicity that $\|A^*\|_2 < 1$ so that $J(A^*)$ is a constant. If $d = n^2$, we obtain the rate $\frac{n}{\sqrt{T}}$ which matches that obtained from (1.2) using the standard inequality $\|\hat{A} - A^*\|_F \leq \sqrt{n} \|\hat{A} - A^*\|_2$. Moreover, we would then also need $T \gtrsim n^2$ in (1.2) in order to drive $\|\hat{A} - A^*\|_F$ below a specified threshold. For general d , however, the sample complexity of estimating A^* is seen to be of order d which is relevant when $d \ll n^2$.

2. (**Example 2**) In this case we obtain Corollary 2 of Theorem 1 which is best interpreted for specific regimes of the sparsity level k . For instance, if k is of the order n , we show that

$$\|\hat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + \sqrt{n \log n}}{\sqrt{T}} + \frac{(n \log n)^{3/2}}{T} \right)$$

if $T \gtrsim J^4(A^*) \max \{n \log n, \log^2(1/\delta)\}$. Assuming $J(A^*)$ is constant, note that $T \gtrsim (n \log n)^{3/2}$ suffices to drive $\|\hat{A} - A^*\|_F$ below a specified threshold, however, this is still much milder than what we need in general.

3. (**Example 3**) Suppose f_0 is piecewise affine convex with K pieces, and where the i th piece lies on a convex subdomain Ω_i with $\Omega_1, \dots, \Omega_K$ a partition of the domain $[0, a]^2$. Suppose each Ω_i is formed by the intersection of at most s pairs of parallel halfspaces¹. Then Corollary 3 of Theorem 2 implies that if n is large enough and $T \gtrsim J^4(A^*) \max \{c^s K \log^s n, \log^2(\frac{1}{\delta})\}$ then with probability at least $1 - \delta$

$$\|\hat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{K c^s \log^{s+1} n}{T} + \frac{\log(1/\delta)}{\sqrt{T}} \right)$$

for some constant $c > 0$. So if $J(A^*)$, K and s are constants then it suffices that $T \gtrsim \text{polylog } n$ to drive the error below a specified threshold.

¹We use the terminology of [21] here. A pair of parallel halfspaces is the set $\{v \in \mathbb{R}^2 : a \leq w^\top v \leq b\}$ for scalars $a < b$ and a unit vector $w \in \mathbb{R}^2$.

1.2 Related work

Learning unstructured LDS. A recent line of work has focused on deriving non-asymptotic error bounds for learning linear systems of the form (1.1), without any explicit structural assumption on A^* . The majority of these works analyze the OLS under different assumptions on $\rho(A^*)$, namely: strict stability ($\rho(A^*) < 1$) [18, 33]; marginal stability ($\rho(A^*) \leq 1$) [34, 31]; purely explosive systems ($\rho(A^*) > 1$) [33, 31]. While \hat{A} is known in closed form, the main challenge in the analysis comes from handling the interaction between the matrix of covariates x_t , and that of noise terms η_t due to their dependencies. Common techniques used in the analysis involve concentration results for self normalized processes [27, 1], and Mendelson’s “small-ball” method [24], the latter of which was extended to dependent data in [34] leading to sharper error bounds. When $\rho(A^*) \leq 1$, the authors in [34] interpret the quantity $\lambda_{\min}(\Gamma_{T-1})$ as a measure of the signal-noise-ratio [34] – larger values lead to improved error bounds. As mentioned earlier, the results of [18] depend on a similar quantity, namely $\lambda_{\min}(\sum_{s=0}^{T-1} \Gamma_s)$, which plays a key role in their error bounds. These terms do not appear explicitly within our analysis and it is unclear (albeit interesting) how this can be done. The main tools that we employ involve concentration results for the suprema of second-order subgaussian chaos processes indexed by a set of matrices [20, 12]; see Section 3.1 for details.

Learning structured LDS. Existing works for learning structured LDS are typically focused on specific structures such as sparsity, group sparsity, or low-rank models.

- In [14], a more general version of (1.1) was considered where $x_{t+1} = A^*x_t + B^*u_t + \eta_t$, with $B^* \in \mathbb{R}^{n \times m}$ and $u_t \in \mathbb{R}^m$ denoting the inputs. Assuming the unknown A^*, B^* to be k -sparse, and $u_t = K_0x_t + v_t$ where v_t is random with a user specified distribution (K_0 is a feedback controller), a LASSO type estimator was analyzed. Assuming x_0 rests at its stationary distribution, uniform asymptotic stability of the closed-loop system, and certain technical assumptions involving A^*, B^* and K_0 , entry-wise error bounds were obtained for the estimation of A^*, B^* . It was shown that these bounds can sometimes be obtained with T of the order $k^2 \log(n + m)$. If k is of order n , this means that $T \gtrsim n^2 \log(n + m)$ samples are needed for recovering the *support* of A^*, B^* . This is larger than our sample complexity bound for controlling the Frobenius norm error.
- In [28], the model (1.1) was considered with A^* assumed to be k -sparse and strictly stable. Under certain assumptions on the problem parameters, it was shown for a LASSO-type estimator that the support of A^* is recovered exactly provided $T \gtrsim \text{poly}(k) \cdot \log n$.
- The model (1.1) is a vector autoregressive (VAR) model of order 1. In [23], a more general order d -VAR model² was considered involving matrices A_1^*, \dots, A_d^* . They analyzed a penalized least squares method where the penalty is imposed by a norm-based regularizer $R(\cdot)$; examples of R including: ℓ_1 norm, group LASSO, sparse group LASSO and OWL (ordered weighted ℓ_1 norm). The recent work [41] considered an order d VAR as well, but allowed for heavy-tail η_t ’s and model-misspecification. They proposed an estimator based on the Yule-Walker equation that deploys a regularizer (norm-based) R , this leads to near minimax optimal results for estimating the model parameters under different structural settings: sparsity, low-rankness, and linear constraints. Other existing results primarily make sparsity-type assumptions on the parameters [22, 3, 19, 35]. The recent work [3] considered the model (1.1) and assumed that A^* has a low-rank plus sparse decomposition, they analyzed a penalized least squares approach for recovering the low-rank and sparse components of A^* .

²Note that an order d -VAR can be written as an order-1 VAR in a higher dimensional space.

- The results in [42] are applicable to model (1.1), with additional linear information about A^* assumed to be available. This can be reformulated as saying that for a known d -dimensional basis $\{V_i\}_{i=1}^d \subset \mathbb{R}^{n \times n}$ and a known offset $\bar{V} \in \mathbb{R}^{n \times n}$, we have $A^* - \bar{V} \in \text{span}\{V_i\}_{i=1}^d$. This is identical to Example 1. If $\rho(A^*) < 1$ and $\|A^*\|_2 \leq C$ for some constant $C > 0$, they show that $\|A^* - \hat{A}\|_F \lesssim \sqrt{\frac{d \log(d/\delta)}{T}}$ provided the smallest singular value of A^* is sufficiently smaller than 1. This is similar to our bound in (1.4). They also cover the setting $\rho(A^*) \leq 1$ where the analysis uses the small ball method [24].

Learning structured signals from random linear measurements. Consider the relatively easier setting of linear regression with independent covariates and noise, i.e.,

$$y = X\beta^* + \eta \tag{1.5}$$

where $X \in \mathbb{R}^{m \times n}$ is the matrix of covariates, $\beta^* \in \mathbb{R}^n$ is the unknown signal, η is noise, and X is *independent* of η . Suppose for simplicity that the entries of X and η are i.i.d centered Gaussian's. The problem of recovering β^* – assuming it belongs to a convex set $\mathcal{K} \subseteq \mathbb{R}^n$ – has received significant interest over the past decade from the statistical and signal processing communities. It is now known that the efficient recovery of β^* is possible via convex programs (e.g., via solving least squares with constraint \mathcal{K} , or by penalizing the objective) with the sample complexity m depending on the Gaussian width of the ‘local size’ of \mathcal{K} at β^* ; see for e.g., [26, 30, 7, 38, 29] and also [2] who introduced a related notion of ‘statistical dimension’. For some sets \mathcal{K} (such as the ℓ_1 ball), sharp estimates for the Gaussian width are available through tools such as Gordon’s escape through the mesh theorem [15], that leads to tight sample complexity bounds. While our proof technique is similar in spirit to these papers (in particular [26, 29]), the model in (1.1) leads to additional technical difficulties. For instance, we cannot use Gordon’s theorem anymore and require other concentration tools for the underlying second order subgaussian chaos. To our knowledge, existing works for finite-time identification of (1.1) do not provide bounds for general closed convex sets \mathcal{K} ; our main goal is to fill this gap (to an extent) by drawing ideas from the above literature.

Shape constrained regression. There is a rich line of work for the Gaussian sequence model

$$y_i = \beta_i^* + \eta_i; \quad i = 1, \dots, n,$$

where $\beta_i^* = f(x_i)$ (with x_i ’s the design points) for some unknown function $f \in \mathcal{F}$, with the function class \mathcal{F} known. This is known as ‘shape constrained’ nonparametric regression since the shape of f (depending on \mathcal{F}) imparts structure to β^* of the form $\beta^* \in \mathcal{K} \subset \mathbb{R}^n$ where

$$\mathcal{K} := \{(f(x_1), \dots, f(x_n)) \in \mathbb{R}^n : f \in \mathcal{F}\}.$$

When η_i ’s are i.i.d centered Gaussians, then the maximum likelihood estimator (MLE) of β^* would be the orthogonal projection of y on \mathcal{K} . There exist many interesting examples of \mathcal{F} which impart an intrinsically low-dimensional structure to \mathcal{K} , thus leading to better error bounds (in terms of n) for the MLE as compared to the unconstrained case. Some examples include: convex functions (e.g., [8, 21]), isotonic functions (e.g., [8, 4]) and Lipschitz functions [26]; see [16] for a nice overview of this area and also the seminal work of Chatterjee [9]. In this respect, our work can be interpreted as shape constrained regression for the vector autoregressive model in (1.1).

2 Problem setup and results

Notation. For any vector $x \in \mathbb{R}^n$, $\|x\|_p$ denotes the usual ℓ_p norm of x . For $X \in \mathbb{R}^{n \times m}$, we denote $\|X\|_2$, $\|X\|_F$ to be respectively the spectral and Frobenius norms of X , while $\langle X, Y \rangle = \text{Tr}(X^\top Y)$ denotes the inner product between X and Y . Also, $\text{vec}(X) \in \mathbb{R}^{nm}$ is formed by stacking the columns of X and $\|X\|_{1,1} = \|\text{vec}(X)\|_1$ denotes the entry-wise ℓ_1 norm of X . The symbol \otimes denotes the Kronecker product between matrices. Sets will be usually denoted by calligraphic letters, and their cardinalities by $|\cdot|$. For any set $\mathcal{S} \subset \mathbb{R}^{n \times m}$ and $X \in \mathbb{R}^{n \times m}$, we will denote

$$\mathcal{S} - X = \{A - X : A \in \mathcal{S}\}.$$

For $a, b > 0$, we say $a \lesssim b$ if there exists a constant $C > 0$ such that $a \leq Cb$. If $a \lesssim b$ and $a \gtrsim b$, then we write $a \asymp b$. For $n \times n$ matrices, we denote the unit Frobenius sphere by \mathbb{S}_n , the unit Frobenius ball by \mathcal{B}_n , and the identity matrix by I_n . The values of symbols used for denoting constants (e.g., c, C, c_1 etc.) may change from line to line. Finally, recall that the subgaussian norm of a random variable X is given by $\|X\|_{\psi_2} := \sup_{p \geq 1} p^{-1/2} (\mathbb{E} |X|^p)^{1/p}$, see for e.g. [39]. We say X is L -subgaussian if $\|X\|_{\psi_2} \leq L$.

2.1 Setup

Consider the autonomous linear dynamical system in (1.1) where $(\eta_t)_{t \geq 1}$ are assumed to be zero-mean, independent and identically distributed (i.i.d) random variables for $t = 0, \dots, T$. Specifically, η_t is assumed to have i.i.d L -subgaussian entries (for some constant L), each of unit variance. Given $(x_t)_{t=0}^{T+1}$, our goal is to estimate A^* under the constraint that $A^* \in \mathcal{K}$ for a closed convex set $\mathcal{K} \subseteq \mathbb{R}^{n \times n}$. We focus on the estimator (1.3) which is a convex program that can usually be efficiently solved by interior-point methods, and in many cases by more efficient methods (e.g., projected gradient descent) specialized to the structure of \mathcal{K} . Also, if η_t were i.i.d Gaussian's, then (1.3) would simply correspond to the maximum likelihood estimator (MLE) of A^* .

In our analysis, we will assume A^* is strictly stable, i.e., its spectral radius $\rho(A^*) < 1$. The quantity $J(A^*)$ defined as

$$J(A^*) := \sum_{i=0}^{\infty} \|(A^*)^i\|_2 \quad (2.1)$$

was introduced in [18] for the analysis of the OLS estimator for strictly stable linear dynamical systems, and will also appear in our results. It is not difficult to verify that $J(A^*)$ is bounded if $\rho(A^*) < 1$, although it could grow with n . If $\|A^*\|_2 < 1$ then $J(A^*) \leq \frac{\|A^*\|_2}{1 - \|A^*\|_2}$.

Before stating our results, we need to present some definitions which will be used later on.

2.2 Preliminaries

We begin by recalling Talagrand's γ_α functionals [37] which can be interpreted as a measure of the complexity of a (not necessarily convex) set.

Definition 1 ([37]). *Let (\mathcal{S}, d) be a metric space. We say that a sequence of subsets of \mathcal{S} , namely $(\mathcal{S}_r)_{r \geq 0}$ is an admissible sequence if $|\mathcal{S}_0| = 1$ and $|\mathcal{S}_r| \leq 2^{2^r}$ for every $r \geq 1$. Then for any $0 < \alpha < \infty$, the γ_α functional of (\mathcal{S}, d) is defined as*

$$\gamma_\alpha(\mathcal{S}, d) := \inf \sup_{s \in \mathcal{S}} \sum_{r=0}^{\infty} 2^{r/\alpha} d(s, \mathcal{S}_r)$$

with the infimum being taken over all admissible sequences of \mathcal{S} .

The following properties of the γ_α functional are useful to note.

1. It can be verified that for any two metrics d_1, d_2 such that $d_1 \leq ad_2$ for some $a > 0$, it holds that $\gamma_\alpha(\mathcal{S}, d_1) \leq a\gamma_\alpha(\mathcal{S}, d_2)$.
2. For $\mathcal{S}' \subset \mathcal{S}$, we have³ that $\gamma_\alpha(\mathcal{S}', d) \leq C_\alpha \gamma_\alpha(\mathcal{S}, d)$ for $C_\alpha > 0$ depending only on α .

The γ_α functionals can be upper bounded in terms of the covering numbers of the set \mathcal{S} . For any $\epsilon > 0$, denote $\mathcal{N}(\mathcal{S}, d, \epsilon)$ to be the minimum number of balls of radius ϵ (with centers in \mathcal{S}) which are needed to cover \mathcal{S} . Then, one can show⁴ that

$$\gamma_\alpha(\mathcal{S}, d) \leq c_\alpha \int_0^{\text{diam}(\mathcal{S})} \log^{1/\alpha} \mathcal{N}(\mathcal{S}, d, \epsilon) d\epsilon \quad (2.2)$$

where $c_\alpha > 0$ depends only on α , and $\text{diam}(\mathcal{S})$ is the diameter of \mathcal{S} . For $\alpha = 2$, the right-hand side (RHS) of (2.2) is the well-known Dudley entropy integral [13]. In fact, by Talagrand's majorizing measure theorem [37], $\gamma_2(\mathcal{S}, d)$ characterizes the expected suprema of centered Gaussian processes $(X_s)_{s \in \mathcal{S}}$ as

$$c\gamma_2(\mathcal{S}, d) \leq \mathbb{E} \sup_{s \in \mathcal{S}} X_s \leq C\gamma_2(\mathcal{S}, d) \quad (2.3)$$

for some universal constants $c, C > 0$, with the canonical distance $d(s, s') := (\mathbb{E}[X_s - X_{s'}]^2)^{1/2}$. For example, if $\mathcal{S} \subset \mathbb{R}^{n \times m}$, and $X_s = \langle G, s \rangle$ for a $n \times m$ matrix G with iid standard Gaussian entries, we have $d(s, s') = \|s - s'\|_F^2$. Then (2.3) implies $\mathbb{E} \sup_{s \in \mathcal{S}} \langle G, s \rangle \asymp \gamma_2(\mathcal{S}, \|\cdot\|_F)$ where $\mathbb{E} \sup_{s \in \mathcal{S}} \langle G, s \rangle$ is known as the *Gaussian width* of the set \mathcal{S} , denoted as $w(\mathcal{S})$.

Tangent cone. One of our sample complexity bounds for estimating A^* will depend on the local size of the *tangent cone* of the set \mathcal{K} at A^* .

Definition 2 (Tangent cone). *For a convex set \mathcal{K} and $A \in \mathcal{K}$, the tangent cone at A is defined as*

$$\mathcal{T}_{\mathcal{K}, A} := \text{cl} \{t(B - A) : t \geq 0, B \in \mathcal{K}\}$$

where $\text{cl}(\cdot)$ denotes the closure of a set.

As we will see shortly, our first main result, namely Theorem 1 below, will involve the gamma functionals $\gamma_1(\mathcal{T}_{\mathcal{K}, A^*} \cap \mathbb{S}_n, \|\cdot\|)$ and $\gamma_2(\mathcal{T}_{\mathcal{K}, A^*} \cap \mathbb{S}_n, \|\cdot\|)$ with $\|\cdot\|$ corresponding to either the spectral or Frobenius norm. Small values of these terms will translate to weaker requirements on the sample size T for accurately estimating A^* . By virtue of the earlier discussion, note that $\gamma_2(\mathcal{T}_{\mathcal{K}, A^*} \cap \mathbb{S}_n, \|\cdot\|_F) \asymp w(\mathcal{T}_{\mathcal{K}, A^*} \cap \mathbb{S}_n)$.

2.3 Main results

Our first main result is the following theorem which bounds the estimation error $\|\widehat{A} - A^*\|_F$. The proofs of our main results are detailed in Section 3.

³This holds for the γ_α functional defined in [37, Definition 2.2.19] with $C_\alpha = 1$. However, γ_α as in Definition 1 is equivalent to that in [37, Definition 2.2.19] up to a constant depending only on α ; see [37, Section 2.3].

⁴This can be deduced using [37, Corollary 2.3.2], and by replicating the arguments after the proof of [37, Lemma 2.2.11] to general $\alpha \geq 1$.

Theorem 1 (Tangent cone structure). *There exist constants $C_1, C_2, C_3, C_4 > 0$ depending only on L such that for any $\delta \in (0, 1)$ and $B \in \mathcal{K}$, if*

$$T \geq C_1 J^4(A^*) \max \left\{ \gamma_2^2(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2), \log^2(C_2/\delta) \right\},$$

then it holds with probability at least $1 - \delta$ that

$$\begin{aligned} \|\hat{A} - A^*\|_F &\leq C_3 J(A^*) \left(\frac{\log(C_2/\delta) + \gamma_2(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_F)}{\sqrt{T}} + \frac{\gamma_1(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2)}{T} \right) \\ &\quad + C_4 J^2(A^*) \|A^* - B\|_F. \end{aligned}$$

In the formulation of this theorem a generic matrix $B \in \mathcal{K}$ is introduced, and a natural choice for it is $B = A^*$, which minimizes the last term in RHS above. However, the shape of $\mathcal{T}_{\mathcal{K}, B}$, which becomes important in evaluation of $\gamma_\alpha(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2)$ for $\alpha = 1, 2$, may be more suitable for calculation if $B \neq A^*$ (but sufficiently ‘close’ to A^*). We will illustrate this on Examples 1 and 2 in the next section.

Our second main result is useful in situations where \mathcal{K} does not have a small tangent cone structure at A^* , or at points in \mathcal{K} sufficiently close to A^* .

Theorem 2 (without tangent cone structure). *There exist constants $C_1, C_2, C_3, C_4 > 0$ depending only on L such that the following is true. For any $\delta \in (0, 1)$, $x > 0$ and $B \in \mathcal{K}$, suppose that*

$$T \geq C_1 J^4(A^*) \max \left\{ \frac{\gamma_2^2((\mathcal{K} - B) \cap x\mathcal{B}_n, \|\cdot\|_2)}{x^2}, \log^2(C_2/\delta) \right\}.$$

Then with probability at least $1 - \delta$, the estimate \hat{A} satisfies

$$\begin{aligned} \|\hat{A} - A^*\|_F &\leq C_3 J(A^*) \left[\frac{\log(C_2/\delta)}{\sqrt{T}} + \frac{\gamma_1((\mathcal{K} - B) \cap x\mathcal{B}_n, \|\cdot\|_2)}{Tx} + \frac{\gamma_2((\mathcal{K} - B) \cap x\mathcal{B}_n, \|\cdot\|_F)}{\sqrt{T}x} \right] \\ &\quad + C_4 J^2(A^*) \|A^* - B\|_F + x. \end{aligned}$$

Suppose we take $B = A^*$ and say A^* lies in the interior of \mathcal{K} . Then $\mathcal{T}_{\mathcal{K}, A^*} = \mathbb{R}^{n \times n}$ which renders Theorem 1 to be unhelpful. In such a situation, Theorem 2 can be meaningful in terms of capturing the local size of \mathcal{K} at A^* at a suitably small scale $x = o(1)$ (as $T \rightarrow \infty$). The optimal choice of x will depend on how the terms involving the γ_1 and γ_2 functionals scale with x . We will apply Theorem 2 on Example 3 in the next section. Interestingly, we will see that both the γ_1 and γ_2 terms are bounded by quantities depending *linearly* on x . In this case, the limit $x = 0$ is the optimal choice.

Remark 1. *Both Theorems 1 and 2 have analogues in the literature for recovering (in the ℓ_2 norm) a structured signal β^* in the linear model (1.5), where the design matrix X and the noise η are typically independent; see for instance [26, 29]. Indeed, our proofs are structured along the same lines as that in [26] (which in turn follows ideas from [29]). However, our linear model in (1.1) is different, especially due to the temporal correlation between x_t and $\eta_t, \eta_{t-1}, \dots, \eta_1$. This leads to additional technical difficulties in the analysis – as opposed to [26, 29] we cannot use Gordon’s mesh theorem [15] and need to instead use a different set of concentration tools (see Section 3.1) for controlling the terms in Lemma 1 (see Section 3.2). One implication of this is the appearance of the γ_1 term in our bounds which can often be larger than the γ_2 term.*

2.4 Instantiating our results on examples

We now recall the examples from Section 1 for which the bounds in Theorems 1 and 2 can be made explicit in terms of n . The proofs are detailed in Section 4.

Example 1 (d -dimensional subspace). In this case, observe from the definition of $\mathcal{T}_{\mathcal{K},B}$ that $\mathcal{T}_{\mathcal{K},B} = \mathcal{K}$ for any choice of $B \in \mathcal{K}$. One can then use standard covering number bounds to bound the γ_α -functionals in Theorem 1. This leads to the following corollary of Theorem 1 where we take $B = A^*$.

Corollary 1. *Let $\mathcal{K} \subset \mathbb{R}^{n \times n}$ be a d -dimensional subspace. If $T \gtrsim J^4(A^*) \max\{d, \log^2(1/\delta)\}$ for $\delta \in (0, 1)$, then it holds with probability at least $1 - \delta$ that*

$$\|\widehat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + \sqrt{d}}{\sqrt{T}} \right).$$

In the unconstrained case where $\mathcal{K} = \mathbb{R}^{n \times n}$, so that $d = n^2$, (1.3) is the OLS estimator. Then Corollary 1 states that if $T \gtrsim J^4(A^*) \max\{n^2, \log^2(1/\delta)\}$, we have

$$\|\widehat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + n}{\sqrt{T}} \right). \quad (2.4)$$

Existing error bounds in the literature for the OLS estimator are typically in the spectral norm, but can of course be converted to the Frobenius norm with an extra factor of \sqrt{n} . Indeed, the result of [18] in (1.2) implies

$$\|\widehat{A} - A^*\|_F \leq \sqrt{n} \|\widehat{A} - A^*\|_2 \lesssim \sqrt{\frac{n \log(1/\delta) + n^2}{T}} \quad (2.5)$$

provided $T \gtrsim J^2(A^*)(\log(1/\delta) + n)$. The bound in (2.5) is of the same order as in (2.4), barring the extra $J(A^*)$ term in our bound. Moreover, when $J(A^*)$ is a constant, note that T needs to be at least of the order n^2 – in both (2.4) and (2.5) – in order to drive the error below a specified threshold. Of course, in case $d \ll n^2$, then the requirement $T \gtrsim d$ in Corollary 1 is relatively mild, as one would expect, given that \mathcal{K} has an intrinsic dimension d .

Example 2 (ℓ_1 ball). We now consider the setting where A^* is k -sparse, i.e., has at most k non-zero entries for some $1 \leq k \leq n^2$. A standard strategy for recovering sparse matrices (resp. vectors) is to take \mathcal{K} to be a suitably scaled ball in the $\|\cdot\|_{1,1}$ (resp. ℓ_1) norm. We take $\mathcal{K} := \|A^*\|_{1,1} \mathcal{B}_{1,n}$ where

$$\mathcal{B}_{1,n} := \{A \in \mathbb{R}^{n \times n} : \|A\|_{1,1} \leq 1\} \quad (2.6)$$

so that A^* lies on the boundary of \mathcal{K} . Then the tangent cone $\mathcal{T}_{\mathcal{K},A^*}$ has the form

$$\mathcal{T}_{\mathcal{K},A^*} := \text{cl}\{tU : \|A^* + U\|_{1,1} \leq \|A^*\|_{1,1}, t \geq 0\}.$$

One can bound $w(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n)$ by using existing results in the literature [30] (see also [36, 2, 7, 38]) – these results apply for vectors but can be directly invoked in our setting by treating $n \times n$ matrices as vectors in \mathbb{R}^{n^2} . This leads to the same order-wise bound on $\gamma_2(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F)$ due to (2.3) with $\mathcal{S} = \mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n$. Furthermore, the γ_1 functional term can be bounded in terms of $w(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n)$ using Sudakov’s minoration inequality [40, Theorem 7.4.1]. These considerations lead to the following corollary of Theorem 1.

Corollary 2. *Suppose A^* is k -sparse and $\mathcal{K} = \|A^*\|_{1,1} \mathcal{B}_{1,n}$ where $\mathcal{B}_{1,n}$ is the unit ball defined in (2.6), and denote $\beta(n, k) := \sqrt{k \log(n^2/k) + k}$. For any $\delta \in (0, 1)$, if*

$$T \gtrsim J^4(A^*) \max\{\beta^2(n, k), \log^2(1/\delta)\}$$

then with probability at least $1 - \delta$ it holds that

$$\|\widehat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + \beta(n, k)}{\sqrt{T}} + \frac{n\beta(n, k)[1 + \log(\frac{n}{\beta(n, k)})]}{T} \right).$$

As a sanity check, note that when $k = n^2$ then $\beta(n, k) = n$ and we recover the statement of Corollary 1 with $d = n^2$. This is expected since A^* does not possess any additional structure, and hence the constraint \mathcal{K} – the purpose of which is to promote sparse solutions – does not provide any benefit. The non-trivial sparsity regime is when $k = o(n^2)$. Consider for instance the case where $k \asymp n$. Then, we have $\beta(n, k) \asymp \sqrt{n \log n}$ and Corollary 2 gives the error bound

$$\|\widehat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{\log(1/\delta) + \sqrt{n \log n}}{\sqrt{T}} + \frac{(n \log n)^{3/2}}{T} \right) \quad (2.7)$$

provided $T \gtrsim J^4(A^*) \max\{n \log n, \log^2(1/\delta)\}$. Notice that while this condition on T implies the error bound in (2.7), we actually need T to be at least of the order $J(A^*)(n \log n)^{3/2}$ to drive the error in (2.7) below a specified threshold. This “gap” is of course due to the term $(n \log n)^{3/2}/T$ in (2.7) – this term arises from the bound on $\gamma_1(\mathcal{K}, A^* \cap \mathbb{S}_n, \|\cdot\|_F)$ within the proof of the corollary and is likely suboptimal. Nevertheless, the bound in (2.7) clearly has a milder dependence on n as compared to that obtained for the OLS in (2.4).

Example 3 (convex regression). We will now illustrate the use of Theorem 2 in the context of convex regression. Consider the case where \mathcal{K} is the collection of matrices formed by sampling bivariate convex functions on a square grid of $\Omega = [0, 1/\sqrt{2}]^2$, i.e.,

$$\mathcal{K} := \left\{ A \in \mathbb{R}^{n \times n} : A_{ij} = f \left(\frac{i-1}{(n-1)\sqrt{2}}, \frac{j-1}{(n-1)\sqrt{2}} \right) \text{ for some } f : \Omega \rightarrow \mathbb{R} \text{ convex} \right\}. \quad (2.8)$$

Then the following corollary of Theorem 2 is obtained taking $B = A^*$ with A^* formed by sampling a convex function with a simple (i.e., piecewise affine) structure.

Corollary 3. *Let $\Omega_1, \dots, \Omega_K$ be convex subsets of $\Omega = [0, 1/\sqrt{2}]^2$ which (a) form a partition of Ω , and (b) are such that each Ω_i is an intersection of at most s pairs of parallel halfspaces. Suppose $A^* \in \mathcal{K}$ for some convex f_0 which is affine on each Ω_i . Then there exists a constant $c > 0$ such that the following is true. For any $\delta \in (0, 1)$, if*

$$\frac{n^2}{\log^s n} \geq Kc^s \quad \text{and} \quad T \gtrsim J^4(A^*) \max \left\{ c^s K \log^s n, \log^2 \left(\frac{1}{\delta} \right) \right\},$$

then with probability at least $1 - \delta$,

$$\|\widehat{A} - A^*\|_F \lesssim J(A^*) \left(\frac{Kc^s \log^s n \log(\frac{n^2}{c^s K \log^s n})}{T} + \frac{\log(1/\delta) + c^{s/2} K^{1/2} \log^{s/2} n}{\sqrt{T}} \right).$$

The proof uses metric entropy estimates from [21, Section 4] for a collection of bivariate convex functions lying within a ball of radius x around f_0 . This then leads to bounds for the γ_1 and γ_2 functionals appearing in Theorem 2. Interestingly, we find for the specified A^* that

$$\begin{aligned} \gamma_1((\mathcal{K} - A^*) \cap x\mathcal{B}_n, \|\cdot\|_2) &\lesssim xc^s K \log^s n \log \left(\frac{n^2}{c^s K \log^s n} \right), \\ \gamma_2((\mathcal{K} - A^*) \cap x\mathcal{B}_n, \|\cdot\|_2) &\lesssim xc^{s/2} K^{1/2} \log^{s/2} n. \end{aligned}$$

Hence the only dependence on x in the error bound of Theorem 2 is through the last term, and since the bound holds for any $x > 0$, we can take the limit $x = 0$.

The scaling factor of $\sqrt{2}$ in the definition of Ω essentially ensures that Ω lies within the unit ℓ_2 ball in \mathbb{R}^2 , which is needed in order to use the results of [21]. Note that in order to drive the error below a threshold, it suffices to have

$$T \gtrsim J^4(A^*) \max \left\{ c^s K \log^s n \log \left(\frac{n^2}{c^s K \log^s n} \right), \log^2 \left(\frac{1}{\delta} \right) \right\}.$$

If, e.g., s and K are constants (i.e., f_0 has a simple structure), then

$$T \gtrsim J^4(A^*) \max \left\{ \text{polylog}(n), \log^2 \left(\frac{1}{\delta} \right) \right\}$$

is sufficient, which is much milder than the requirement for OLS.

3 Proof of main results

We begin by outlining in Section 3.1 some important concentration tools used in the proof of our main results. Section 3.2 contains the proof of Theorem 1 while Section 3.3 contains the proof of Theorem 2.

3.1 Technical tools

For a set of matrices \mathcal{A} , let us define the terms

$$d_F(\mathcal{A}) = \sup_{A \in \mathcal{A}} \|A\|_F, \quad d_2(\mathcal{A}) = \sup_{A \in \mathcal{A}} \|A\|_2, \quad (3.1)$$

which can be thought of as other types of complexity measures of the set \mathcal{A} (the ‘‘radius’’ of \mathcal{A}).

The following result from [20] provides a concentration bound for the suprema of second order subgaussian chaos processes involving positive semidefinite (p.s.d) matrices.

Theorem 3 ([20]). *Let \mathcal{A} be a set of matrices and ξ be a vector whose entries are independent, zero-mean, variance 1, and are L -subgaussian random variables. Denote*

$$P = \gamma_2(\mathcal{A}, \|\cdot\|_2) [\gamma_2(\mathcal{A}, \|\cdot\|_2) + d_F(\mathcal{A})] + d_F(\mathcal{A}) d_2(\mathcal{A}) \\ V = d_2(\mathcal{A}) [\gamma_2(\mathcal{A}, \|\cdot\|_2) + d_F(\mathcal{A})], \quad \text{and} \quad U = d_2^2(\mathcal{A})$$

where d_2, d_F are as in (3.1). Then there exist constants $c_1, c_2 > 0$ depending only on L such that for any $t > 0$ it holds that

$$\mathbb{P} \left(\sup_{A \in \mathcal{A}} \left| \|A\xi\|_2^2 - \mathbb{E}[\|A\xi\|_2^2] \right| \geq c_1 P + t \right) \leq 2 \exp \left(-c_2 \min \left\{ \frac{t^2}{V^2}, \frac{t}{U} \right\} \right).$$

We will also use the following result from [12] for bounding the suprema of general second order subgaussian chaos processes, where the matrices are not necessarily p.s.d. The bound on moments in part 1 is stated on page 15 in [12]; part 2 follows by passing from moment bounds to tail bounds in a standard manner via Markov’s inequality, see for example [12, Lemma A.1].

Theorem 4 ([12]). *Let \mathcal{A} be a set of matrices and ξ be a vector whose entries are independent, zero-mean, 1-subgaussian random variables. For $A \in \mathcal{A}$, denote $C_A(\xi) := \xi^\top A \xi - \mathbb{E}[\xi^\top A \xi]$. With d_2, d_F as in (3.1), the following is true.*

1. There exists a universal constant $c > 0$ such that for any $p \geq 1$,

$$\begin{aligned} \left(\mathbb{E} \sup_{A \in \mathcal{A}} |C_A(\xi)|^p \right)^{1/p} &\leq c(\gamma_1(\mathcal{A}, \|\cdot\|_2) + \gamma_2(\mathcal{A}, \|\cdot\|_F) + \sqrt{p}d_F(\mathcal{A}) + pd_2(\mathcal{A})) \\ &= c\rho(\mathcal{A}, p). \end{aligned}$$

2. There exists a universal constant $c' > 0$ (depending on c) such that for any $u \geq 1$,

$$\mathbb{P} \left(\sup_{A \in \mathcal{A}} |C_A(\xi)| \geq c' \rho(\mathcal{A}, u) \right) \leq e^{-u}.$$

3.2 Proof of Theorem 1

Let us define the matrices (each of size $n \times T$)

$$\tilde{X} = [x_2 \cdots x_{T+1}], \quad X = [x_1 \cdots x_T], \quad \text{and } E = [\eta_2 \cdots \eta_{T+1}],$$

so that $\tilde{X} = A^*X + E$ with $x_0 = 0$ and $x_1 = \eta_1$. Then, (1.3) can be rewritten as

$$\hat{A} \in \underset{A \in \mathcal{K}}{\operatorname{argmin}} \|\tilde{X} - AX\|_F^2. \quad (3.2)$$

Step 1. Our starting point is the following inequality which follows from first-order optimality conditions for constrained convex programs.

Lemma 1. For any $B \in \mathcal{K}$ the solution \hat{A} of (3.2) satisfies

$$\|(\hat{A} - B)X\|_F^2 \leq \langle (\hat{A} - B)X, E \rangle + \|(A^* - B)X\|_F \|(\hat{A} - B)X\|_F.$$

Proof. We first expand $\|\tilde{X} - \hat{A}X\|_F^2$ as

$$\begin{aligned} \|\tilde{X} - \hat{A}X\|_F^2 &= \|\tilde{X} - BX\|_F^2 + \|(B - \hat{A})X\|_F^2 + 2 \langle \tilde{X} - BX, (B - \hat{A})X \rangle \\ &= \|\tilde{X} - BX\|_F^2 - \|(B - \hat{A})X\|_F^2 + 2 \langle \tilde{X} - \hat{A}X, (B - \hat{A})X \rangle \\ &\leq \|\tilde{X} - BX\|_F^2 - \|(B - \hat{A})X\|_F^2 \end{aligned} \quad (3.3)$$

where the last inequality follows due to $\langle \tilde{X} - \hat{A}X, (B - \hat{A})X \rangle \leq 0$ by the first order optimality condition for \hat{A} . Using $\tilde{X} = A^*X + E$, we obtain

$$\begin{aligned} \|\tilde{X} - \hat{A}X\|_F^2 &= \|(\hat{A} - A^*)X\|_F^2 + \|E\|_F^2 + 2 \langle (A^* - \hat{A})X, E \rangle, \\ \|\tilde{X} - BX\|_F^2 - \|(B - \hat{A})X\|_F^2 &= \|(A^* - B)X\|_F^2 + \|E\|_F^2 \\ &\quad + 2 \langle (A^* - B)X, E \rangle - \|(B - \hat{A})X\|_F^2. \end{aligned}$$

Plugging these expressions in (3.3) leads to the inequality

$$\|(A^* - \hat{A})X\|_F^2 \leq \|(A^* - B)X\|_F^2 + 2 \langle (\hat{A} - B)X, E \rangle - \|(B - \hat{A})X\|_F^2. \quad (3.4)$$

Expanding the LHS of (3.4) leads to the lower bound

$$\begin{aligned} \|(A^* - \hat{A})X\|_F^2 &= \|(A^* - B)X\|_F^2 + \|(B - \hat{A})X\|_F^2 + 2 \langle (A^* - B)X, (B - \hat{A})X \rangle \\ &\geq \|(A^* - B)X\|_F^2 + \|(B - \hat{A})X\|_F^2 - 2\|(A^* - B)X\|_F \|(B - \hat{A})X\|_F, \end{aligned}$$

and plugging this in (3.4) readily leads to the stated inequality in the lemma. \square

Our next goal is to bound the terms appearing in Lemma 1.

Step 2. Consider first the term $\|(\widehat{A} - B)X\|_F^2$ which can be written as

$$\|(\widehat{A} - B)X\|_F^2 = \sum_{t=0}^T \|(\widehat{A} - B)x_t\|_2^2 = \|(I_T \otimes (\widehat{A} - B)) \text{vec}(X)\|_2^2.$$

One can verify that $\text{vec}(X) = \Gamma\xi$ where

$$\Gamma = \begin{bmatrix} I_n & 0 & \dots & 0 \\ A^* & I_n & \dots & 0 \\ \vdots & & \ddots & \vdots \\ (A^*)^{T-1} & \dots & A^* & I_n \end{bmatrix} \in \mathbb{R}^{Tn \times Tn}, \quad \xi = \begin{bmatrix} \eta_1 \\ \eta_2 \\ \vdots \\ \eta_T \end{bmatrix} \in \mathbb{R}^{Tn},$$

and so, we can write $\|(\widehat{A} - B)X\|_F^2 = \|(I_T \otimes (\widehat{A} - B))\Gamma\xi\|_2^2$. As shown in [18], we can bound $\|\Gamma\|_2 \leq J(A^*)$; recall its definition from (2.1). Now let us define the set

$$\widetilde{\mathcal{T}}_{\mathcal{K},B} := \{(I_T \otimes A)\Gamma : A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n\};$$

we then clearly have

$$\frac{(I_T \otimes (\widehat{A} - B))\Gamma}{\|\widehat{A} - B\|_F} = \left(I_T \otimes \left(\frac{\widehat{A} - B}{\|\widehat{A} - B\|_F} \right) \right) \Gamma \in \widetilde{\mathcal{T}}_{\mathcal{K},B}$$

since $\widehat{A} - B \in \mathcal{T}_{\mathcal{K},B}$. Therefore we can bound the term $\|(\widehat{A} - B)X\|_F^2$ as

$$\left(\inf_{W \in \widetilde{\mathcal{T}}_{\mathcal{K},B}} \|W\xi\|_2^2 \right) \|\widehat{A} - B\|_F^2 \leq \|(\widehat{A} - B)X\|_F^2 \leq \|\widehat{A} - B\|_F^2 \left(\sup_{W \in \widetilde{\mathcal{T}}_{\mathcal{K},B}} \|W\xi\|_2^2 \right). \quad (3.5)$$

The term $\|W\xi\|_2^2$ is a second order subgaussian chaos involving p.s.d matrices, and we wish to control its infimum and supremum over the set $\widetilde{\mathcal{T}}_{\mathcal{K},B}$ in (3.5). This is done using Theorem 3 and leads to the following lemma.

Lemma 2. Denote $\overline{U}, \overline{V}, \overline{P}$ as

$$\begin{aligned} \overline{U} &= J^2(A^*), \quad \overline{V} = J(A^*) \left[\gamma_2(\widetilde{\mathcal{T}}_{\mathcal{K},B}, \|\cdot\|_2) + \sqrt{T}J(A^*) \right], \\ \overline{P} &= \gamma_2(\widetilde{\mathcal{T}}_{\mathcal{K},B}, \|\cdot\|_2) \left[\gamma_2(\widetilde{\mathcal{T}}_{\mathcal{K},B}, \|\cdot\|_2) + \sqrt{T}J(A^*) \right] + \sqrt{T}J^2(A^*). \end{aligned}$$

Then there exist constants $c_1, c_2 > 0$ depending only on L such that for any $t > 0$, it holds with probability at least $1 - 2 \exp(-c_2 \min\{\frac{t^2}{\overline{V}}, \frac{t}{\overline{U}}\})$ that

1. $\inf_{W \in \widetilde{\mathcal{T}}_{\mathcal{K},B}} \|W\xi\|_2^2 \geq T - c_1\overline{P} - t$, and
2. $\sup_{W \in \widetilde{\mathcal{T}}_{\mathcal{K},B}} \|W\xi\|_2^2 \leq TJ^2(A^*) + c_1\overline{P} + t$.

Proof. We will use Theorem 3 to obtain the stated bounds. Let us first bound the terms $d_F(\widetilde{\mathcal{T}}_{\mathcal{K},B})$, $d_2(\widetilde{\mathcal{T}}_{\mathcal{K},B})$ as follows.

$$d_F(\widetilde{\mathcal{T}}_{\mathcal{K},B}) = \sup_{W \in \widetilde{\mathcal{T}}_{\mathcal{K},B}} \|W\|_F = \sup_{A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n} \|(I_T \otimes A)\Gamma\|_F \leq J(A^*)\sqrt{T} \quad (3.6)$$

since $\|(I_T \otimes A)\Gamma\|_F \leq \|\Gamma\|_2 \|I_T \otimes A\|_F \leq J(A^*)\sqrt{T}$ for any $A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n$. Furthermore,

$$d_2(\tilde{\mathcal{T}}_{\mathcal{K},B}) = \sup_{W \in \tilde{\mathcal{T}}_{\mathcal{K},B}} \|W\|_2 = \sup_{A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n} \|(I_T \otimes A)\Gamma\|_2 \leq J(A^*) \quad (3.7)$$

since $\|(I_T \otimes A)\Gamma\|_2 \leq \|A\|_2 \|\Gamma\|_2 \leq \|A\|_F \|\Gamma\|_2 \leq J(A^*)$ for all $A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n$.

Next, we use almost matching bounds on $\mathbb{E}[\|(I_T \otimes A)\Gamma\xi\|_2^2]$ holding uniformly over $A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n$ as

$$\mathbb{E}[\|(I_T \otimes A)\Gamma\xi\|_2^2] = \|(I_T \otimes A)\Gamma\|_F^2 = \begin{cases} \leq TJ^2(A^*) \\ \geq T\|A\|_F^2 = T \end{cases}$$

where for the lower bound we used the fact that A appears T times within $(I_T \otimes A)\Gamma$. Finally, one can readily verify that the terms $\bar{P}, \bar{U}, \bar{V}$ are (resp.) upper bounds for P, U and V , the latter terms defined in Theorem 3. This concludes the proof. \square

Hence the event in Lemma 2 implies the bounds

$$\|(\hat{A} - B)X\|_F^2 \geq \|\hat{A} - B\|_F^2 (T - c_1\bar{P} - t), \quad (3.8)$$

$$\|(\hat{A} - B)X\|_F \leq \|\hat{A} - B\|_F (TJ^2(A^*) + c_1\bar{P} + t)^{1/2}, \quad (3.9)$$

$$\text{and } \|(A^* - B)X\|_F \leq \|A^* - B\|_F (TJ^2(A^*) + c_1\bar{P} + t)^{1/2}. \quad (3.10)$$

Step 3. Our goal now is to control the term $\langle (\hat{A} - B)X, E \rangle$. We can first bound it as

$$\langle (\hat{A} - B)X, E \rangle = \|\hat{A} - B\|_F \left\langle \left(\frac{\hat{A} - B}{\|\hat{A} - B\|_F} \right) X, E \right\rangle \leq \|\hat{A} - B\|_F \left(\sup_{A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n} \langle AX, E \rangle \right).$$

It remains to control the supremum term, which in fact is the supremum of a second order sub-gaussian chaos involving matrices that are not necessarily p.s.d (as will be seen in the proof below). This is achieved via Theorem 4 and leads to the following lemma.

Lemma 3. *For any $u \geq 1$, we have with probability at least $1 - \exp(-u)$,*

$$\sup_{A \in \mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n} |\langle AX, E \rangle| \leq c_3 \left(uJ(A^*) + \sqrt{uT}J(A^*) + \gamma_1(\tilde{\mathcal{T}}_{\mathcal{K},B}, \|\cdot\|_2) + \gamma_2(\tilde{\mathcal{T}}_{\mathcal{K},B}, \|\cdot\|_F) \right)$$

for some constant $c_3 > 0$ depending only on L .

Proof. We begin by rewriting $\langle AX, E \rangle$ as

$$\langle AX, E \rangle = \langle A^\top E, X \rangle = \langle (I_T \otimes A^\top) \tilde{\xi}, \Gamma \xi \rangle = \tilde{\xi}^\top ((I_T \otimes A)\Gamma) \xi \quad \text{where } \tilde{\xi} = \text{vec}(E) = \begin{bmatrix} \eta_2 \\ \eta_3 \\ \vdots \\ \eta_{T+1} \end{bmatrix} \in \mathbb{R}^{Tn}.$$

Denoting $\eta \in \mathbb{R}^{(T+1)n}$ to be the vector formed by stacking $\eta_1, \dots, \eta_{T+1}$, we can further simplify $\tilde{\xi}^\top ((I_T \otimes A)\Gamma) \xi$ as

$$\tilde{\xi}^\top ((I_T \otimes A)\Gamma) \xi = \eta^\top \underbrace{\begin{bmatrix} 0 & 0 \\ (I_T \otimes A)\Gamma & 0 \end{bmatrix}}_{M_A} \eta = \eta^\top M_A \eta.$$

This in particular implies $\mathbb{E}[\langle AX, E \rangle] = \mathbb{E}[\eta^\top M_A \eta] = 0$. Denoting the set

$$\mathcal{M} = \{M_A : A \in \mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n\}$$

we obtain

$$\sup_{A \in \mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n} \langle AX, E \rangle = \sup_{M_A \in \mathcal{M}} \eta^\top M_A \eta.$$

We now use Theorem 4 to control this suprema. To this end, we observe that

$$d_F(\mathcal{M}) = \sup_{M_A \in \mathcal{M}} \|M_A\|_F = \sup_{A \in \mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n} \|(I_T \otimes A)\Gamma\|_F \leq \sqrt{T}J(A^*), \quad (\text{using (3.6)})$$

$$d_2(\mathcal{M}) = \sup_{M_A \in \mathcal{M}} \|M_A\|_2 = \sup_{A \in \mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n} \|(I_T \otimes A)\Gamma\|_2 \leq J(A^*). \quad (\text{using (3.7)})$$

It is easy to see that $\gamma_\alpha(\mathcal{M}, \|\cdot\|) = \gamma_\alpha(\tilde{\mathcal{T}}_{\mathcal{K}, B}, \|\cdot\|)$ where $\|\cdot\|$ denotes any matrix norm. Then by using Theorem 4 with ξ therein corresponding to $\frac{u}{T}$, we readily arrive at the statement of the lemma. \square

The event in Lemma 3 implies the bound

$$\begin{aligned} \langle (\hat{A} - B)X, E \rangle &\leq c_3 \|\hat{A} - B\|_F \left(uJ(A^*) + \sqrt{uT}J(A^*) + \gamma_1(\tilde{\mathcal{T}}_{\mathcal{K}, B}, \|\cdot\|_2) + \gamma_2(\tilde{\mathcal{T}}_{\mathcal{K}, B}, \|\cdot\|_F) \right) \\ &= c_3 \varrho(\tilde{\mathcal{T}}_{\mathcal{K}, B}, u) \|\hat{A} - B\|_F \end{aligned} \quad (3.11)$$

with $\varrho(\cdot, \cdot)$ as defined in Theorem 4.

Step 4: Putting it together. Using the results from (3.8), (3.9), (3.10), (3.11) in Lemma 1, we have with probability at least $1 - 2 \exp(-c_2 \min\{\frac{t^2}{\bar{V}^2}, \frac{t}{\bar{V}}\}) - \exp(-u)$ that

$$\begin{aligned} \|\hat{A} - B\|_F^2 (T - c_1 \bar{P} - t) &\leq c_3 \varrho(\tilde{\mathcal{T}}_{\mathcal{K}, B}, u) \|\hat{A} - B\|_F + \|A^* - B\|_F \|\hat{A} - B\|_F (TJ^2(A^*) + c_1 \bar{P} + t) \\ \iff \|\hat{A} - B\|_F &\leq c_3 \frac{\varrho(\tilde{\mathcal{T}}_{\mathcal{K}, B}, u)}{T - c_1 \bar{P} - t} + \|A^* - B\|_F \left(\frac{TJ^2(A^*) + c_1 \bar{P} + t}{T - c_1 \bar{P} - t} \right). \end{aligned} \quad (3.12)$$

Our aim is to simplify the above bounds and also showcase the dependency on the local tangent cone $\mathcal{T}_{\mathcal{K}, B}$. To this end, the following claim is useful.

Claim 1. For any $\alpha > 0$,

$$\begin{aligned} \gamma_\alpha(\tilde{\mathcal{T}}_{\mathcal{K}, B}, \|\cdot\|_2) &\leq J(A^*) \gamma_\alpha(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2), \\ \gamma_\alpha(\tilde{\mathcal{T}}_{\mathcal{K}, B}, \|\cdot\|_F) &\leq \sqrt{T}J(A^*) \gamma_\alpha(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_F). \end{aligned}$$

Proof. For any $X, Y \in \mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n$, consider the matrices $A_1 = (I_n \otimes X)\Gamma$ and $A_2 = (I_n \otimes Y)\Gamma$. Clearly, $\|A_1 - A_2\|_2 \leq J(A^*)\|X - Y\|_2$ and $\|A_1 - A_2\|_F \leq \sqrt{T}J(A^*)\|X - Y\|_F$. This then readily implies the stated bounds using the definition of γ_α functionals. \square

Claim 1 leads to the following bounds on the terms \bar{P}, \bar{V} .

$$\begin{aligned} \bar{P} &\leq \bar{P}_1 := J^2(A^*) \left(\gamma_2^2(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T} \right) \\ \bar{V} &\leq \bar{V}_1 := J^2(A^*) \left(\gamma_2(\mathcal{T}_{\mathcal{K}, B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T} \right) \end{aligned}$$

Furthermore, for $1 \leq u \leq T$ we can bound $\varrho(\tilde{\mathcal{T}}_{\mathcal{K},B}, u)$ as

$$\varrho(\tilde{\mathcal{T}}_{\mathcal{K},B}, u) \leq J(A^*) \left(2\sqrt{uT} + \gamma_1(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_F) \right).$$

The above considerations lead to the following simplification of (3.12),

$$\begin{aligned} \|\hat{A} - B\|_F &\leq c_3 \frac{J(A^*) \left(2\sqrt{uT} + \gamma_1(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_F) \right)}{T - c_1\bar{P}_1 - t} \\ &\quad + \|A^* - B\|_F \frac{TJ^2(A^*) + c_1\bar{P}_1 + t}{T - c_1\bar{P}_1 - t} \end{aligned} \quad (3.13)$$

which holds with probability at least $1 - 2 \exp(-c_2 \min\{\frac{t^2}{\bar{V}_1^2}, \frac{t}{\bar{U}_1}\}) - \exp(-u)$.

Now choosing $t = \bar{V}_1\sqrt{u}$, note that

$$\min \left\{ \frac{t^2}{\bar{V}_1^2}, \frac{t}{\bar{U}_1} \right\} = \min \left\{ u, \sqrt{u} \frac{\bar{V}_1}{\bar{U}} \right\} \geq \min(u, \sqrt{u}) = \sqrt{u}$$

where the inequality holds since $\bar{V}_1 \geq \bar{U}$, and the final equality uses the condition $u \geq 1$. Also note that (3.13) holds provided $c_1\bar{P}_1 + t < T$. Using the condition $u \geq 1$, it is then easily verified that

$$c_1\bar{P}_1 + t \leq c'_1 J^2(A^*) \left(\gamma_2^2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2) + \sqrt{uT} \right)$$

for some constant $c'_1 > 0$ depending on c_1 . So it suffices to ensure that $c_1\bar{P}_1 + t \leq T/2$ for which a sufficient condition is

$$J^2(A^*) \left(\frac{\gamma_2^2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2)}{T} + \frac{\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2)}{\sqrt{T}} + \sqrt{\frac{u}{T}} \right) \leq \frac{1}{2c'_1}. \quad (3.14)$$

Since $J(A^*) \geq 1$, hence (3.14) holds provided

$$T \geq c''_1 J^4(A^*) \max \{ \gamma_2^2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2), u \}$$

for a suitably large constant $c''_1 > 1$ (depending on c'_1). Observe that the above condition implies $u \leq T$.

The above considerations are summarized in the form of the following theorem.

Theorem 5. *There exist constants $C_1, C_2, C_3, C_4 > 0$ depending only on L such that the following is true. For any $u \geq 1$ and $B \in \mathcal{K}$, suppose that*

$$T \geq C_1 J^4(A^*) \max \{ \gamma_2^2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2), u \}.$$

Then with probability at least $1 - C_2 \exp(-C_3\sqrt{u})$, the estimate \hat{A} in (3.2) satisfies

$$\|\hat{A} - B\|_F \leq C_4 J(A^*) \left(\sqrt{\frac{u}{T}} + \frac{\gamma_1(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2)}{T} + \frac{\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_F)}{\sqrt{T}} \right) + 3J^2(A^*) \|A^* - B\|_F.$$

The statement of Theorem 1 is obtained by choosing $u = \frac{1}{C_3^2} \log^2(\frac{C_2}{\delta})$ for any $\delta \in (0, 1)$, and by the triangle inequality

$$\begin{aligned} \|\hat{A} - A^*\|_F &\leq \|\hat{A} - B\|_F + \|A^* - B\|_F \\ &\leq C_4 J(A^*) \left(\sqrt{\frac{u}{T}} + \frac{\gamma_1(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_2)}{T} + \frac{\gamma_2(\mathcal{T}_{\mathcal{K},B} \cap \mathbb{S}_n, \|\cdot\|_F)}{\sqrt{T}} \right) + 4J^2(A^*) \|A^* - B\|_F. \end{aligned}$$

This completes the proof.

3.3 Proof of Theorem 2

We will follow similar notation as in the proof of Theorem 1. Our starting point is inequality in Lemma 1 which we recall below.

$$\|(\widehat{A} - B)X\|_F^2 \leq \langle (\widehat{A} - B)X, E \rangle + \|(A^* - B)X\|_F \|(\widehat{A} - B)X\|_F.$$

Suppose $\|\widehat{A} - B\|_F \geq x$ since otherwise the statement of the theorem holds trivially. Denote $a = \frac{x}{\|\widehat{A} - B\|_F} \leq 1$. Since \mathcal{K} is convex, hence $U = a\widehat{A} + (1-a)B \in \mathcal{K}$ which implies $a(\widehat{A} - B) = U - B$.

Denoting $\widetilde{\mathcal{K}} = \mathcal{K} - B$ from now on, note that $\widetilde{\mathcal{K}}$ is a star-shaped set⁵ since \mathcal{K} is convex and $\widetilde{\mathcal{K}}$ contains the origin. We would like to obtain an analogue of [29, Lemma 4.4] to our setting where we lower bound

$$\inf_{M \in \widetilde{\mathcal{K}}, \|M\|_F \geq x} \|MX\|_F = \inf_{M \in \widetilde{\mathcal{K}}, \|M\|_F \geq x} \|(I_T \otimes M)\Gamma\xi\|_2$$

Denoting $\widetilde{M} = x \frac{M}{\|M\|_F}$ for any $M \in \widetilde{\mathcal{K}}$ such that $\|M\|_F \geq x$, clearly $\widetilde{M} \in \widetilde{\mathcal{K}}$ since $\widetilde{\mathcal{K}}$ is star-shaped. This implies

$$\inf_{M \in \widetilde{\mathcal{K}}, \|M\|_F \geq x} \frac{\|(I_T \otimes M)\Gamma\xi\|_2}{\|M\|_F} = \inf_{\widetilde{M} \in \widetilde{\mathcal{K}} \cap x\mathcal{S}_n} \frac{\|(I_T \otimes \widetilde{M})\Gamma\xi\|_2}{x}. \quad (3.15)$$

Now we can bound $\mathbb{E}[\|(I_T \otimes \widetilde{M})\Gamma\xi\|_2^2]$ uniformly over $\widetilde{M} \in \widetilde{\mathcal{K}} \cap x\mathcal{S}_n$ as

$$\mathbb{E}[\|(I_T \otimes \widetilde{M})\Gamma\xi\|_2^2] = \|(I_T \otimes \widetilde{M})\Gamma\|_F^2 \begin{cases} \leq x^2 T J^2(A^*) \\ \geq T \|\widetilde{M}\|_F^2 = T x^2. \end{cases}$$

Furthermore, denoting $\mathcal{S}_x := \{(I_T \otimes \widetilde{M})\Gamma : \widetilde{M} \in \widetilde{\mathcal{K}} \cap x\mathcal{S}_n\}$, it is easy to verify that

$$d_F(\mathcal{S}_x) \leq J(A^*)x\sqrt{T}, \quad d_2(\mathcal{S}_x) \leq J(A^*)x.$$

This implies the following bounds on the terms P, U and V .

$$\begin{aligned} P &\leq \gamma_2(\mathcal{S}_x, \|\cdot\|_2) \left(\gamma_2(\mathcal{S}_x, \|\cdot\|_2) + xJ(A^*)\sqrt{T} \right) + x^2 J^2(A^*)\sqrt{T} =: \overline{P}, \\ V &\leq xJ(A^*) \left(\gamma_2(\mathcal{S}_x, \|\cdot\|_2) + xJ(A^*)\sqrt{T} \right) =: \overline{V}, \\ U &= d_2^2(\mathcal{S}_x) \leq x^2 J^2(A^*) =: \overline{U}. \end{aligned}$$

Note that (3.15) also holds with inf replaced by sup on both sides. Hence using Theorem 3, it follows that with probability at least $1 - 2 \exp(-c_2 \min\{\frac{t^2}{\overline{V}^2}, \frac{t}{\overline{U}}\})$,

1. $\inf_{M \in \widetilde{\mathcal{K}}, \|M\|_F \geq x} \frac{\|(I_T \otimes M)\Gamma\xi\|_2}{\|M\|_F} \geq \frac{(Tx^2 - c_1 \overline{P} - t)^{1/2}}{x}$, and
2. $\sup_{M \in \widetilde{\mathcal{K}}, \|M\|_F \geq x} \frac{\|(I_T \otimes M)\Gamma\xi\|_2}{\|M\|_F} \leq \frac{(Tx^2 J^2(A^*) + c_1 \overline{P} + t)^{1/2}}{x}$,

⁵A set \mathcal{X} is star shaped if $\lambda\mathcal{X} \subset \mathcal{X}$ for any $\lambda \in [0, 1]$. Any convex set containing the origin is star-shaped.

which in turn implies the bounds

$$\|(\widehat{A} - B)X\|_F^2 \geq \frac{\|\widehat{A} - B\|_F^2}{x^2}(Tx^2 - c_1\overline{P} - t), \quad (3.16)$$

$$\|(\widehat{A} - B)X\|_F \leq \frac{\|\widehat{A} - B\|_F}{x}(Tx^2 J^2(A^*) + c_1\overline{P} + t)^{1/2}, \quad (3.17)$$

$$\text{and } \|(A^* - B)X\|_F \leq \frac{\|A^* - B\|_F}{x}(Tx^2 J^2(A^*) + c_1\overline{P} + t)^{1/2}. \quad (3.18)$$

We now have to bound $\langle (\widehat{A} - B)X, E \rangle$. Denoting $\delta = \|\widehat{A} - B\|_F$, the fact that $\widetilde{\mathcal{K}}$ is star-shaped implies $\delta^{-1}\widetilde{\mathcal{K}} \subset x^{-1}\widetilde{\mathcal{K}}$ (since $\delta^{-1} \leq x^{-1}$). This then implies

$$\left\langle \frac{(\widehat{A} - B)X}{\|\widehat{A} - B\|_F}, E \right\rangle \leq \sup_{U \in \delta^{-1}\widetilde{\mathcal{K}} \cap x\mathcal{B}_n} \langle UX, E \rangle \leq \sup_{U \in x^{-1}\widetilde{\mathcal{K}} \cap x\mathcal{B}_n} \langle UX, E \rangle = \frac{1}{x} \sup_{U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n} \langle UX, E \rangle$$

which means that

$$\left\langle (\widehat{A} - B)X, E \right\rangle \leq \frac{\|\widehat{A} - B\|_F}{x} \sup_{U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n} \langle UX, E \rangle.$$

Now as in the proof of Theorem 1, we can write

$$\langle UX, E \rangle = \eta^\top \underbrace{\begin{bmatrix} 0 & 0 \\ (I_T \otimes U)\Gamma & 0 \end{bmatrix}}_{M_U} \eta = \eta^\top M_U \eta$$

with $\mathbb{E}[\langle UX, E \rangle] = 0$. Denoting $\mathcal{U}_x := \{M_U : U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n\}$, we then obtain

$$\sup_{U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n} \langle UX, E \rangle = \sup_{M_U \in \mathcal{U}_x} \eta^\top M_U \eta.$$

It can be checked that $d_F(\mathcal{U}_x) \leq xJ(A^*)\sqrt{T}$ and $d_2(\mathcal{U}_x) \leq J(A^*)x$. Moreover, denoting $\widetilde{\mathcal{S}}_x := \{(I_T \otimes U)\Gamma : U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n\}$, we have that $\gamma_\alpha(\mathcal{U}_x, \|\cdot\|) = \gamma_\alpha(\widetilde{\mathcal{S}}_x, \|\cdot\|)$ for $\alpha > 0$ and any norm $\|\cdot\|$.

Then by invoking Theorem 4, it follows that for $u \geq 1$, it holds with probability at least $1 - \exp(-u)$ that

$$\begin{aligned} \sup_{U \in \widetilde{\mathcal{K}} \cap x\mathcal{B}_n} |\langle UX, E \rangle| &\leq c_3 \left(\gamma_1(\widetilde{\mathcal{S}}_x, \|\cdot\|_2) + \gamma_2(\widetilde{\mathcal{S}}_x, \|\cdot\|_F) + \sqrt{uT}xJ(A^*) + uxJ(A^*) \right) \\ &=: \bar{\varrho}(\widetilde{\mathcal{S}}_x, x, u), \\ \implies \left\langle (\widehat{A} - B)X, E \right\rangle &\leq c_3 \frac{\|\widehat{A} - B\|_F}{x} \bar{\varrho}(\widetilde{\mathcal{S}}_x, x, u). \end{aligned} \quad (3.19)$$

Hence using (3.16), (3.17), (3.18) and (3.19), it holds w.p at least $1 - 2 \exp(-c_2 \min\{\frac{t^2}{\overline{V}}, \frac{t}{\overline{U}}\}) - \exp(-u)$ (for $u \geq 1$) that

$$\|\widehat{A} - B\|_F \leq c_3 \frac{x\bar{\varrho}(\widetilde{\mathcal{S}}_x, x, u)}{Tx^2 - c_1\overline{P} - t} + \|A^* - B\|_F \left(\frac{Tx^2 J^2(A^*) + c_1\overline{P} + t}{Tx^2 - c_1\overline{P} - t} \right). \quad (3.20)$$

It remains to simplify the expressions of the terms involved in our bounds. To this end, we obtain analogous to Claim 1 that for any $\alpha > 0$,

$$\begin{aligned}\gamma_\alpha(\mathcal{S}_x, \|\cdot\|_2) &\leq \gamma_\alpha(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_2)J(A^*), \\ \gamma_\alpha(\mathcal{S}_x, \|\cdot\|_F) &\leq \sqrt{T}J(A^*)\gamma_\alpha(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_F)J(A^*), \\ \gamma_\alpha(\tilde{\mathcal{S}}_x, \|\cdot\|_2) &\leq \gamma_\alpha(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2)J(A^*), \\ \gamma_\alpha(\tilde{\mathcal{S}}_x, \|\cdot\|_F) &\leq \sqrt{T}J(A^*)\gamma_\alpha(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F)J(A^*).\end{aligned}$$

Furthermore, we can bound \bar{P}, \bar{V} as

$$\begin{aligned}\bar{P} &\leq \gamma_2(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_2)J(A^*) \left(\gamma_2(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_2)J(A^*) + xJ(A^*)\sqrt{T} \right) + x^2J^2(A^*)\sqrt{T} =: \bar{P}_1, \\ \bar{V} &\leq xJ^2(A^*) \left(\gamma_2(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_2) + x\sqrt{T} \right) =: \bar{V}_1,\end{aligned}$$

and since $1 \leq u \leq T$, we also obtain

$$\begin{aligned}\bar{\varrho}(\tilde{\mathcal{S}}_x, x, u) &\leq \gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2)J(A^*) + \gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F)\sqrt{T}J(A^*) + 2\sqrt{uT}xJ(A^*) \\ &= J(A^*) \left(\gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F) + 2x\sqrt{uT} \right).\end{aligned}$$

Hence when $1 \leq u \leq T$, we can simplify (3.20) as

$$\begin{aligned}\|\hat{A} - B\|_F &\leq c_3 \frac{xJ(A^*) \left(2x\sqrt{uT} + \gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2) + \sqrt{T}\gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F) \right)}{Tx^2 - c_1\bar{P} - t} \\ &\quad + \|A^* - B\|_F \left(\frac{Tx^2J^2(A^*) + c_1\bar{P} + t}{Tx^2 - c_1\bar{P} - t} \right),\end{aligned}$$

which is analogous to (3.13). The remaining steps of the proof follow the same calculations as in the proof of Theorem 1 (details omitted) which leads us to the statement of the theorem. We only remark that the condition on T follows from the fact $\gamma_2(\tilde{\mathcal{K}} \cap x\mathbb{S}_n, \|\cdot\|_2) \lesssim \gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2)$ since $\mathbb{S}_n \subset \mathcal{B}_n$.

4 Proof of corollaries for examples

4.1 Proof of Corollary 1

Since $\mathcal{T}_{\mathcal{K}, B} = \mathcal{K}$ for any $B \in \mathcal{K}$, we choose $B = A^*$. We will bound the terms $\gamma_2(\mathcal{K} \cap \mathbb{S}_n, \|\cdot\|_F)$, $\gamma_1(\mathcal{K} \cap \mathbb{S}_n, \|\cdot\|_2)$ using (2.2). We first employ the simplified bound

$$\gamma_1(\mathcal{K} \cap \mathbb{S}_n, \|\cdot\|_2) \leq \gamma_1(\mathcal{K} \cap \mathbb{S}_n, \|\cdot\|_F)$$

since $\|\cdot\|_2 \leq \|\cdot\|_F$. Now the set $\mathcal{K} \cap \mathcal{B}_n \subset \mathbb{R}^{n \times n}$ is a d -dimensional unit ball w.r.t the $\|\cdot\|_F$ norm, hence it follows⁶ from standard volumetric estimates (see for e.g. [40, Corollary 4.2.13]) that

$$\mathcal{N}(\mathcal{K} \cap \mathcal{B}_n, \|\cdot\|_F, \epsilon) \leq \left(\frac{3}{\epsilon} \right)^d \quad \text{for } \epsilon \leq 1$$

and $\mathcal{N}(\mathcal{K} \cap \mathcal{B}_n, \|\cdot\|_F, \epsilon) = 1$ for $\epsilon > 1$. Since $\text{diam}(\mathcal{K} \cap \mathcal{B}_n) = 2$, it follows from (2.2) that for $\alpha = 1, 2$,

$$\gamma_\alpha(\mathcal{K} \cap \mathbb{S}_n, \|\cdot\|_F) \lesssim \gamma_\alpha(\mathcal{K} \cap \mathcal{B}_n, \|\cdot\|_F) \lesssim \int_0^1 d^{1/\alpha} \log^{1/\alpha} \left(\frac{3}{\epsilon} \right) d\epsilon \lesssim d^{1/\alpha},$$

where the first inequality holds since $\mathcal{K} \cap \mathbb{S}_n \subset \mathcal{K} \cap \mathcal{B}_n$. This completes the proof.

⁶We simply treat a matrix $X \in \mathbb{R}^{n \times n}$ by its vector form $\text{vec}(X) \in \mathbb{R}^{n^2}$.

4.2 Proof of Corollary 2

Since A^* is k -sparse, we can bound $w(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n)$ as

$$w(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n) \lesssim \sqrt{k \log(n^2/k) + k} =: \beta(n, k);$$

see for instance [38, Section 4.3] which can be directly applied to our setting. Then using (2.3) and the fact $\|\cdot\|_2 \leq \|\cdot\|_F$, we obtain

$$\gamma_2(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_2) \leq \gamma_2(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F) \lesssim w(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n) \lesssim \beta(n, k).$$

Now starting with the inequality $\gamma_1(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_2) \leq \gamma_1(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F)$, we seek to bound the term $\gamma_1(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F)$ using (2.2). To this end, we can use Sudakov's minoration inequality⁷ which yields for any $\epsilon > 0$,

$$\log \mathcal{N}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F, \epsilon) \lesssim \frac{w^2(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n)}{\epsilon^2} \lesssim \frac{\beta^2(n, k)}{\epsilon^2}. \quad (4.1)$$

We can also obtain an alternate bound on $\mathcal{N}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F, \epsilon)$ by using the following useful fact [40, Exercise 4.2.10]: for a metric space (\mathcal{S}, d) and $\mathcal{S}' \subset \mathcal{S}$, it holds for $\epsilon > 0$ that $\mathcal{N}(\mathcal{S}', d, \epsilon) \leq \mathcal{N}(\mathcal{S}, d, \epsilon/2)$. Applied to our setting, this yields the bound

$$\mathcal{N}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F, \epsilon) \leq \mathcal{N}(\mathbb{S}_n, \|\cdot\|_F, \epsilon/2) \leq \mathcal{N}(\mathcal{B}_n, \|\cdot\|_F, \epsilon/4) \leq \left(\frac{12}{\epsilon}\right)^{n^2}, \quad (4.2)$$

where the final bound follows directly from [40, Corollary 4.2.13]. Hence from (4.1), (4.2), we have for all $\epsilon > 0$ that

$$\log \mathcal{N}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F, \epsilon) \lesssim \min \left\{ \frac{\beta^2(n, k)}{\epsilon^2}, n^2 \log \left(\frac{12}{\epsilon} \right) \right\}.$$

Since $\text{diam}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n) = 2$, we can employ (2.2) to obtain the bound

$$\gamma_1(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F) \leq \int_0^2 \log \mathcal{N}(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F, \epsilon) d\epsilon \lesssim \underbrace{\int_0^{\epsilon_1} n^2 \log(12/\epsilon) d\epsilon}_{u_1} + \underbrace{\int_{\epsilon_1}^2 \frac{\beta^2(n, k)}{\epsilon^2} d\epsilon}_{u_2}, \quad (4.3)$$

for any $\epsilon_1 \in (0, 2)$. It is easy to verify that

$$u_1 = n^2 \epsilon_1 \left(\log \left(\frac{12}{\epsilon_1} \right) + 1 \right) \leq 2n^2 \epsilon_1 \log \left(\frac{12}{\epsilon_1} \right),$$

and $u_2 = \left(\frac{2 - \epsilon_1}{2\epsilon_1} \right) \beta^2(n, k)$.

Using this in (4.3), and choosing $\epsilon_1 = \beta(n, k)/n$, we obtain the bound

$$\gamma_1(\mathcal{T}_{\mathcal{K},A^*} \cap \mathbb{S}_n, \|\cdot\|_F) \lesssim n\beta(n, k) \left[\log \left(\frac{n}{\beta(n, k)} \right) + 1 \right].$$

The statement of the corollary now follows directly. Note that for the above choice of ϵ_1 to be in the interval $(0, 2)$, we need $n \geq \beta(n, k)/2$ to hold but this is always satisfied.

⁷see for instance [40, Theorem 7.4.1]; the canonical metric therein is $\|\cdot\|_F$.

4.3 Proof of Corollary 3

The proof relies on using metric entropy bounds developed in [21, Section 4] for the collection of convex functions lying within a ball around f_0 . In what follows, we will borrow some notations and definitions from [21].

Let \mathcal{S} denote an equispaced $n \times n$ grid of $\Omega = [0, 1/\sqrt{2}]^2$ with spacing $\delta = \frac{1}{(n-1)\sqrt{2}}$ along each axis. Note that Ω lies within the unit ℓ_2 ball of \mathbb{R}^2 (centered at origin). This is a requirement for the results in [21] to hold, but can be ensured by a suitable scaling. For any function defined on Ω , let us define the quasi-norm

$$\ell_{\mathcal{S}}(f, \Omega, p) = \left(\frac{1}{|\Omega \cap \mathcal{S}|} \sum_{b \in \Omega \cap \mathcal{S}} |f(b)|^p \right)^{1/p}.$$

Then for any convex f_0 defined on Ω , and $x > 0$, define a ball centered at f_0 ,

$$\mathcal{B}_{\mathcal{S}}^p(f_0; x; \Omega) = \{f : \Omega \rightarrow \mathbb{R} \text{ s.t } f \text{ is convex and } \ell_{\mathcal{S}}(f - f_0, \Omega, p) \leq x\}.$$

Now denoting $\tilde{\mathcal{K}} = \mathcal{K} - A^*$, note that

$$\tilde{\mathcal{K}} \cap x\mathcal{B}_n = \{A - A^* : A \in \mathcal{K}, \|A - A^*\|_F \leq x\}.$$

We need to bound the metric entropy $\log \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon)$. To this end, it is not difficult to verify that

$$\mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) = \mathcal{N}\left(\mathcal{B}_{\mathcal{S}}^2\left(f_0; \frac{x}{n}; \Omega\right), \ell_{\mathcal{S}}(\cdot, \Omega, 2), \frac{\epsilon}{n}\right). \quad (4.4)$$

The RHS of (4.4) is bounded by [21, Corollary 4.2], which in our setup implies

$$\log \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) = \log \mathcal{N}\left(\mathcal{B}_{\mathcal{S}}^2\left(f_0; \frac{x}{n}; \Omega\right), \ell_{\mathcal{S}}(\cdot, \Omega, 2), \frac{\epsilon}{n}\right) \leq c^s K \left(\frac{x}{\epsilon}\right) \log^s n$$

for some constant $c > 0$. Moreover, using standard volumetric estimates as earlier (e.g., [40, Proposition 4.2.12, Corollary 4.2.13]), we also have $\log \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) \leq n^2 \log(6x/\epsilon)$, thus implying

$$\log \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) \leq \min \left\{ c^s K \left(\frac{x}{\epsilon}\right) \log^s n, n^2 \log \left(\frac{6x}{\epsilon}\right) \right\}.$$

Now it remains to bound the gamma functionals by deploying (2.2).

1. **(Bounding $\gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2)$).** Since $\text{diam}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n) = 2x$, we have using (2.2), for any $0 < \epsilon_1 < 2x$, that

$$\begin{aligned} \gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2) &\leq \gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F) \lesssim \int_0^{2x} \log \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) d\epsilon \\ &\lesssim n^2 \int_0^{\epsilon_1} \log \left(\frac{6x}{\epsilon}\right) d\epsilon + c^s K x \log^s n \int_{\epsilon_1}^{2x} \frac{1}{\epsilon} d\epsilon \\ &\lesssim n^2 \epsilon_1 \log \left(\frac{6x}{\epsilon_1}\right) + c^s K x \log^s n \log \left(\frac{2x}{\epsilon_1}\right). \end{aligned}$$

In the final inequality above, we used similar calculations as within the proof of Corollary 2. Choosing $\epsilon_1 = \frac{c^s K x \log^s n}{n^2}$ then implies

$$\gamma_1(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F) \lesssim x c^s K \log^s n \log \left(\frac{n^2}{c^s K \log^s n}\right)$$

where we note that $\epsilon_1 < 2x$ is ensured provided $n^2 \geq K c^s \log^s n$.

2. (**Bounding** $\gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2)$). Again using (2.2), we obtain

$$\begin{aligned}
\gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_2) &\leq \gamma_2(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F) \\
&\lesssim \int_0^{2x} \log^{1/2} \mathcal{N}(\tilde{\mathcal{K}} \cap x\mathcal{B}_n, \|\cdot\|_F, \epsilon) d\epsilon \\
&\lesssim \int_0^{2x} c^{s/2} \frac{(Kx)^{1/2}}{\sqrt{\epsilon}} \log^{s/2} n d\epsilon \\
&\lesssim xc^{s/2} K^{1/2} \log^{s/2} n
\end{aligned}$$

Plugging these bounds in Theorem 2, we see that the dependence on x cancels out from all terms barring the additive term at the end. Since Theorem 2 is true for any $x > 0$, we can now take the limit $x = 0$ to obtain the corollary.

Acknowledgment

H.T would like to thank Yassir Jedra and Sjoerd Dirksen for helpful clarifications regarding their results in their works [18], [12] respectively.

References

- [1] Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24, 2011.
- [2] Dennis Amelunxen, Martin Lotz, Michael B. McCoy, and Joel A. Tropp. Living on the edge: phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 3(3):224–294, 2014.
- [3] Sumanta Basu and George Michailidis. Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics*, 43(4):1535–1567, 2015.
- [4] Pierre C. Bellec. Sharp oracle inequalities for Least Squares estimators in shape restricted regression. *The Annals of Statistics*, 46(2):745 – 780, 2018.
- [5] Emmanuel J. Candès and Yaniv Plan. Near-ideal model selection by ℓ_1 minimization. *The Annals of Statistics*, 37(5A):2145 – 2177, 2009.
- [6] Emmanuel J. Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006.
- [7] Venkat Chandrasekaran, Benjamin Recht, Pablo A. Parrilo, and Alan S. Willsky. The convex geometry of linear inverse problems. *Found. Comput. Math.*, 12(6):805–849, 2012.
- [8] Sabyasachi Chatterjee, Adityanand Guntuboyina, and Bodhisattva Sen. On risk bounds in isotonic and other shape restricted regression problems. *The Annals of Statistics*, 43(4):1774 – 1800, 2015.
- [9] Sourav Chatterjee. A new perspective on least squares under convex constraint. *The Annals of Statistics*, 42(6):2340–2381, 2014.
- [10] Wenyu Chen and Rahul Mazumder. Multivariate convex regression at scale. *arXiv: 2005.11588*, 2021.
- [11] I.C. Demetriou and P. Tzitziris. Infant mortality and economic growth: Modeling by increasing returns and least squares. In *Proceedings of the World Congress on Engineering*, 2017.
- [12] Sjoerd Dirksen. Tail bounds via generic chaining. *Electronic Journal of Probability*, 20(none):1 – 29, 2015.
- [13] R.M Dudley. The sizes of compact subsets of Hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- [14] Salar Fattahi, Nikolai Matni, and Somayeh Sojoudi. Learning sparse dynamical systems from a single sample trajectory. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2682–2689, 2019.
- [15] Y. Gordon. On Milman’s inequality and random subspaces which escape through a mesh in R^n . In *Geometric Aspects of Functional Analysis*, pages 84–106, 1988.
- [16] Adityanand Guntuboyina and Bodhisattva Sen. Nonparametric Shape-Restricted Regression. *Statistical Science*, 33(4):568 – 594, 2018.

- [17] Clifford Hildreth. Point estimates of ordinates of concave functions. *Journal of the American Statistical Association*, 49(267):598–619, 1954.
- [18] Yassir Jedra and Alexandre Proutiere. Finite-time identification of stable linear systems optimality of the least-squares estimator. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 996–1001, 2020.
- [19] Anders Bredahl Kock and Laurent Callot. Oracle inequalities for high dimensional vector autoregressions. *Journal of Econometrics*, 186(2):325–344, 2015.
- [20] Felix Krahmer, Shahar Mendelson, and Holger Rauhut. Suprema of chaos processes and the restricted isometry property. *Communications on Pure and Applied Mathematics*, 67(11):1877–1904, 2014.
- [21] Gil Kur, Fuchang Gao, Adityanand Guntuboyina, and Bodhisattva Sen. Convex regression in multidimensions: Suboptimality of least squares estimators. *arXiv: 2006.02044*, 2020.
- [22] Po-Ling Loh and Martin J. Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *The Annals of Statistics*, 40(3):1637 – 1664, 2012.
- [23] Igor Melnyk and Arindam Banerjee. Estimating structured vector autoregressive models. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 830–839, 2016.
- [24] Shahar Mendelson. Learning without concentration. In *Proceedings of The 27th Conference on Learning Theory*, volume 35, pages 25–39, 2014.
- [25] Sahand N. Negahban, Pradeep Ravikumar, Martin J. Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [26] Matey Neykov. Gaussian regression with convex constraints. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89, pages 31–38, 2019.
- [27] Shao Qi-Man Pena Victor H., Lai Tze L. *Self-Normalized Processes Limit Theory and Statistical Applications*. Probability and Its Applications. Springer Berlin Heidelberg, 2009.
- [28] José Pereira, Morteza Ibrahimi, and Andrea Montanari. Learning networks of stochastic differential equations. In *Advances in Neural Information Processing Systems*, volume 23, 2010.
- [29] Yaniv Plan and Roman Vershynin. The generalized Lasso with non-linear observations. *IEEE Trans. Inf. Theor.*, 62(3):1528–1537, 2016.
- [30] Mark Rudelson and Roman Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008.
- [31] Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, volume 97, pages 5610–5618, 2019.
- [32] Emilio Seijo and Bodhisattva Sen. Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics*, 39(3):1633 – 1657, 2011.

- [33] Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite time identification in unstable linear systems. *Automatica*, 96:342–353, 2018.
- [34] Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Proceedings of the 31st Conference On Learning Theory*, volume 75, pages 439–473, 2018.
- [35] Song Song and Peter J. Bickel. Large vector auto regressions. *arXiv:1106.3915*, 2011.
- [36] Mihailo Stojnic. Various thresholds for ℓ_1 -optimization in compressed sensing. *arXiv:0907.3666*, 2009.
- [37] Michel Talagrand. *Upper and lower bounds for stochastic processes*, volume 60. Springer, 2014.
- [38] Joel A. Tropp. *Convex Recovery of a Structured Signal from Independent Random Linear Measurements*, pages 67–101. 2015.
- [39] Roman Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, page 210–268. Cambridge University Press, 2012.
- [40] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Number 47 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- [41] Di Wang and Ruey S. Tsay. Rate-optimal robust estimation of high-dimensional vector autoregressive models. *The Annals of Statistics*, 51(2):846 – 877, 2023.
- [42] Yao Zheng and Guang Cheng. Finite-time analysis of vector autoregressive models under linear restrictions. *Biometrika*, 108(2):469–489, 2021.