

Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification

Max Simchowitz

MSIMCHOW@BERKELEY.EDU

Horia Mania

HMANIA@BERKELEY.EDU

Stephen Tu

STEPHENT@BERKELEY.EDU

Michael I. Jordan

JORDAN@BERKELEY.EDU

Benjamin Recht

BRECHT@BERKELEY.EDU

Department of Electrical Engineering and Computer Science, University of California, Berkeley

Editors: Sebastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

We prove that the ordinary least-squares (OLS) estimator attains nearly minimax optimal performance for the identification of linear dynamical systems from a single observed trajectory. Our upper bound relies on a generalization of Mendelson’s small-ball method to dependent data, eschewing the use of standard mixing-time arguments. Our lower bounds reveal that these upper bounds match up to logarithmic factors. In particular, we capture the correct signal-to-noise behavior of the problem, showing that *more unstable* linear systems are *easier* to estimate. This behavior is qualitatively different from arguments which rely on mixing-time calculations that suggest that unstable systems are more difficult to estimate. We generalize our technique to provide bounds for a more general class of linear response time-series.

Keywords: Linear dynamical systems, autoregressive processes, time series, system identification, empirical process theory

1. Introduction

System identification—the problem of estimating the parameters of a dynamical system given a time series of its trajectories—is a fundamental problem in time-series analysis, control theory, robotics, and reinforcement learning. Despite its importance, sharp, non-asymptotic analyses for the sample complexity of system identification are rare. In particular, it is not known how many trajectories are required to identify the parameters of an unknown *linear* system. Properly characterizing this sample complexity would have profound implications, since accurate error bounds are indispensable for designing robust and high-performing control systems. It is important that the bounds be sharp, in the sense that they do not drastically *overestimate* the number of required measurements from system trajectories, which are often time-consuming and prohibitively expensive to collect. More broadly, a deeper understanding of system identification would inform other statistical problems where one wishes to learn from non-i.i.d. or time-correlated data.

We focus on the problem of identifying a discrete-time *linear dynamical system* from an observed trajectory. Such systems are described by two parameter matrices A_* and B_* , and the dynamics evolve according to the law $X_{t+1} = A_*X_t + B_*u_t + \eta_t$, where $X_t \in \mathbb{R}^d$ is the state of the system, u_t is the input of the system, and $\eta_t \in \mathbb{R}^d$ denotes unobserved process noise. Linear systems are fundamental in control theory, since they are able to capture the behavior of many natural

systems and also able to accurately describe the evolution of an even broader class of systems near their equilibria. Despite the importance of understanding the statistical properties of system identification, the relationship between the matrix A_* and the statistical rate for estimating this matrix remains poorly understood. We note that the larger the state vectors X_t are in comparison to the process noise, the larger the *signal-to-noise ratio* for estimating A_* is. As a result, larger matrices A_* (larger in an appropriate sense, discussed later) lead to states X_t of larger norm, which in turn should make the estimation of A_* easier. However, it is difficult to theoretically formalize this intuition because the sequence of measurements X_0, X_1, \dots, X_{T-1} used for estimation is not i.i.d. and it is dependent on the noise $\eta_0, \eta_1, \dots, \eta_{T-2}$. Even the computationally straightforward ordinary least-squares (OLS) estimator is difficult to analyze. Standard analyses for OLS on random design linear regression (Hsu et al., 2014) cannot be used due to the dependency between the covariates X_t and the process noise η_t .

In the statistics and machine learning literature, correlated data is usually dealt with using mixing-time arguments (Yu, 1994), which relies on fast convergence to a stationary distribution that allows correlated samples to be treated roughly as if they were independent. While this approach has been successfully used to develop generalization bounds for time-series data (Mohri and Rostamizadeh, 2007a), a fundamental limitation of mixing-time arguments is that the bounds deteriorate when the underlying process is slower to mix. In the case of linear systems, this behavior is qualitatively incorrect. For linear systems, the rate of mixing is intimately tied to the eigenvalues of the matrix A_* , specifically the *spectral radius* $\rho(A_*)$. When $\rho(A_*) < 1$ (i.e. when the system is *stable*), the process mixes to a stationary distribution at a rate that deteriorates as $\rho(A_*)$ approaches the boundary of one. However, as discussed above, as $\rho(A_*)$ increases we expect estimation to become easier due to better signal-to-noise ratio, and not harder as mixing-time arguments suggest. We note that recent work by Faradonbeh et al. (2017a) studying the estimation problem for linear systems relies in the stable case on concentration of measure arguments which also degrade as the mixing-time of the system grows.

We address these difficulties and offer a new statistical analysis of the ordinary least-squares (OLS) estimator of the dynamics $X_{t+1} = A_* X_t + \eta_t$ with no inputs, when the spectral radius of A_* is at most one ($\rho(A_*) \leq 1$, a regime known as *marginal stability*). Our results, detailed in Section 2, show that the statistical performance of OLS is determined by the minimum eigenvalue of the (finite-time) controllability Gramian $\Gamma_T = \sum_{s=0}^{T-1} A_*^s (A_*^\top)^s$. The controllability Gramian is a fundamental quantity in the theory of linear systems; the eigenvalues of the Gramian quantify how much white process noise $\eta_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2 I)$ can excite the system. We show that a larger $\lambda_{\min}(\Gamma_T)$ leads to faster estimation of A_* in operator norm, and we also prove that up to log factors the OLS estimator is minimax optimal. Furthermore, in Section 2.3 we offer similar statistical guarantees for a more general class of linear response time-series.

1.1. Related Work

Most directly related to our work is a recent series of papers by Faradonbeh et al. (2017a,b), who study the linear system identification problem by proving a non-asymptotic rate on the convergence of the OLS estimator to the true system matrices. In the regime where A_* is stable, Faradonbeh et al. recover a similar rate as our result. The major difference is that the dependence of their analysis on the spectral properties of A_* are qualitatively suboptimal, and difficult to interpret precisely. Their analysis is based on separately establishing concentration of the sample covariance

matrix $\sum_{t=1}^T X_t X_t^\top$ to the stationary covariance matrix and bounding the martingale difference term $\sum_{t=1}^T X_t \eta_t$. This decoupled analysis inevitably picks up a dependence on the condition number of the stationary covariance matrix, which means that as the system becomes more unstable, their bound deteriorates. Indeed, such a strategy is unable to provide any insight into the behavior of OLS when, for example, A_* is a scaled orthogonal matrix. On the other hand, our analysis does not decouple the two terms, and as a result our bounds only degrade in the *logarithm* of the condition number of the finite-time controllability Gramian Γ_T . Faradonbeh et al. (2017a) also provide a bound in the *unstable regime*, which we believe can be sharpened using our analysis techniques which couple the covariate- and noise-processes. We leave this to future work. Moreover, our analysis of one-dimensional, unstable systems corroborates the linear convergence behavior that Faradonbeh et al. (2017a) obtain for “explosive” systems, which are systems where *all* eigenvalues of A_* lie outside the complex unit disk.

Another closely related work is the scalar analysis by Rantzer (2018). In fact, our proof technique for scalar systems can be seen as an extension of his technique. The main difference is that by more carefully tracking the terms that appear in the moment generating function of the noise and covariate processes, we are able to discriminate behaviors that arise when A_* is stable versus unstable, and uncover a linear rate of convergence in the unstable regime.

Our result qualitatively matches the behavior of the rate given in Dean et al. (2017), in that the key spectral quantity governing the rate of convergence is the minimum eigenvalue of the finite-time controllability Gramian. The major difference is that the analysis in Dean et al. uses multiple independent trajectories, and discards all but the last state-transition in each trajectory. This decouples the covariates, and reduces the analysis to that of random design linear regression with independent covariates. We note, however, that the analysis in Dean et al. applies even when A_* is unstable.

More broadly, there has been recent interest in non-asymptotic analysis of linear system identification problems. Some of the earlier non-asymptotic literature in system identification include Campi and Weyer (2002) and Vidyasagar and Karandikar (2008). The results provided in this line of work are often quite conservative, featuring quantities which are exponential in the degree of the system. Furthermore, the rates given are often difficult to interpret. More recently, Shah et al. (2012) pose the problem of recovering a single-input, single-output (SISO) LTI system from linear measurements in the frequency domain as a sparse recovery problem, proving polynomial sample complexity for recovery in the \mathcal{H}_2 -norm. Hardt et al. (2016) show that under fairly restrictive assumptions on the A_* matrix, projected gradient descent recovers the state-space representation of an LTI system with only a polynomial number of samples. The analysis from both Shah et al. and Hardt et al. both degrade polynomially in $\frac{1}{1-\rho(A_*)}$, where $\rho(A_*)$ is the spectral radius of underlying A_* . On the other hand, Hazan et al. (2017) propose a new spectral filtering algorithm for online prediction of linear systems where the rates do not degenerate as $\rho(A_*) \rightarrow 1$, with the caveat that the analysis only applies to symmetric A_* matrices. Hazan et al. (2018) extends the analysis to diagonalizable matrices, but the obtained error rates are polynomial in problem parameters. Both works also consider the more general setting where X_t is observed indirectly via $Y_t = CX_t$ for an unknown observation matrix C . Moreover, the main metric of interest in both Hardt et al. (2016) and Hazan et al. (2017, 2018) is the prediction error. It is not clear how prediction error guarantees can be used in downstream robust control synthesis, whereas the operator norm bounds we provide can be used as direct inputs into robust synthesis for optimal control problems (Dean et al., 2017).

The most well-established technique in the statistics literature for dealing with non-independent, time-series data is the use of mixing-time arguments (Yu, 1994). In the machine learning literature,

mixing arguments have been used to develop generalization bounds (Mohri and Rostamizadeh, 2007a,b; Kuznetsov and Mohri, 2017; McDonald et al., 2017) which are analogous to the classical generalization bounds for i.i.d. data. As mentioned previously, a fundamental limitation of mixing-time arguments is that the bounds all degrade as the mixing-time increases. This has two implications for linear system identification: (a) none of these existing results can correctly capture the qualitative behavior as the A_* matrix reaches instability, and (b) these techniques cannot be applied to the regime where A_* is unstable, for which estimation is not only well-posed, but should be quite easy. It is for these reasons we do not pursue such arguments in this work.

2. Results

In this work, we consider both the specific problem of estimating linear dynamical systems, and a more general problem of linear estimation in time series. In both cases we measure the estimation error in the operator norm $\|\cdot\|_{\text{op}}$. In Section 2.1 we present upper bounds on the estimation error of the parameters A_* of a linear dynamical system, which hold for any A_* with $\rho(A_*) \leq 1$. In Section 2.2 we show that these upper bounds are nearly optimal in many regimes of interest. Finally, Section 2.3 states a general result, which applies to covariate processes with linear responses.

Notation: We let \mathcal{S}^{d-1} denote the unit sphere in \mathbb{R}^d . Given a matrix M we denote by M^\dagger its pseudoinverse. For a symmetric matrix $M \in \mathbb{R}^{d \times d}$, we let $\lambda_{\max}(M)$ and $\lambda_{\min}(M)$ denote its largest and smallest eigenvalues. If $M \in \mathbb{R}^{d \times d}$ and $M \succ 0$, we denote by \mathcal{S}_M the set of all points $x \in \mathbb{R}^d$ such that $\|M^{-1/2}x\|_2 = 1$.

2.1. Linear Dynamical Systems

We analyze the statistical performance of the OLS estimator of the parameter A_* from a single observed trajectory X_1, \dots, X_{T+1} satisfying $X_{t+1} = A_*X_t + \eta_t$, where $X_0 = 0$ and $\eta_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2 I_d)$:

$$\hat{A}(T) := \arg \min_{A \in \mathbb{R}^{d \times d}} \sum_{t=1}^T \frac{1}{2} \|X_{t+1} - AX_t\|_2^2. \quad (2.1)$$

Our bounds are stated in terms of the finite-time controllability Gramian of the system, denoted by $\Gamma_t := \sum_{s=0}^{t-1} (A_*^s)(A_*^s)^\top$, which captures the magnitude of the excitations induced by the process noise. Indeed, we can write X_t explicitly as

$$X_t = \sum_{s=1}^t A_*^{t-s} \eta_{s-1} \quad \text{which implies that } \mathbb{E}[X_t X_t^\top] = \sigma^2 \Gamma_t. \quad (2.2)$$

Hence, the expected covariance can be expressed in terms of the Gramians via $\mathbb{E}[\sum_{t=1}^T X_t X_t^\top] = \sigma^2 \cdot \sum_{t=1}^T \Gamma_t$. As is standard in analyses of least-squares, “larger” covariates/covariance matrices correspond to faster rates of learning. We are ready to state our first result, proved in Section 3:

Theorem 2.1 *Fix $\delta \in (0, 1/2)$ and consider the linear dynamical system $X_{t+1} = A_*X_t + \eta_t$, where A_* is a marginally stable matrix in $\mathbb{R}^{d \times d}$ (i.e. $\rho(A_*) \leq 1$), $X_0 = 0$, and $\eta_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2 I)$. Then there exist universal constants $c, C > 0$ such that*

$$\mathbb{P} \left[\left\| \hat{A}(T) - A_* \right\|_{\text{op}} > \frac{C}{\sqrt{T \lambda_{\min}(\Gamma_k)}} \sqrt{d \log \frac{d}{\delta} + \log \det(\Gamma_T \Gamma_k^{-1})} \right] \leq \delta, \quad (2.3)$$

for any $k \geq 1$ such that $\frac{T}{k} \geq c(d \log(d/\delta) + \log \det(\Gamma_T \Gamma_k^{-1}))$ holds.

Note that σ^2 does not appear in the bound from Theorem 2.1 because scaling the noise also rescales the covariates. In Appendix A, we show that for any marginally stable A_* , we can always choose a $k \geq 1$ provided T is sufficiently large. Therefore, even when $\rho(A_*) = 1$ and the system does not mix, we obtain finite-sample estimation guarantees which also guarantees consistency of estimation. In many cases, these rates are qualitatively no-worse than random-design linear regression with independent covariates (Theorem A.2 and Remark A).

In general, $\lambda_{\min}(\Gamma_k)$ is a nondecreasing function of the block length k . The intuition for this is that larger k takes into account more long-term excitations to lower bound the size of our covariance matrix. However, as we use longer blocks, our high probability bounds degrade. Thus, the optimal block length is the maximal value k which satisfies the condition in Theorem 2.1.

The dependence on the minimum eigenvalue of the Gramian $\lambda_{\min}(\Gamma_k)$ has two interpretations. From a *statistical* perspective, we have $\frac{1}{2k \cdot \sigma^2} \mathbb{E}[\sum_{t=1}^{2k} X_t X_t^\top] = \frac{1}{2k} \sum_{t=1}^{2k} \Gamma_t \succeq \frac{1}{2} \lambda_{\min}(\Gamma_k) \cdot I$. Thus, $\lambda_{\min}(\Gamma_k)$ gives a lower bound on the smallest eigenvalue value of the covariance matrix associated with the first $2k$ covariates. In fact, one can also show (see (3.12)) that for any $t_0 \geq 0$, we still have $\frac{1}{2k \cdot \sigma^2} \mathbb{E}[\sum_{t=t_0+1}^{t_0+2k} X_t X_t^\top | X_{t_0}] \succeq \frac{1}{2} \lambda_{\min}(\Gamma_k) \cdot I$. Theorem 2.1 thus states that the larger the expected covariance matrix, the faster A_* is estimated. Note that $\Gamma_k \succeq I$ for all $k \geq 1$.

The second interpretation is *dynamical*. The term $\lambda_{\min}(\Gamma_k)$ corresponds to the ‘‘excitability’’ of the system, which is the extent to which the process noise η_t influences future covariates. This can be seen from (2.2), where the slower $(A_*^{t_0})(A_*^{t_0})^\top$ decays as t_0 grows, the larger the contribution of η_{t-t_0-1} is. This is precisely the reason why linear systems with larger spectral radii mix slowly, and do not mix when $\rho(A_*) \geq 1$. In this light, Theorem 2.1 shows that with high-probability, the more a linear system is excited by the noise η_t , the easier it is to estimate the parameter matrix A_* . For stable systems with $\rho(A_*) < 1$, the following corollary removes the explicit dependence on the block length k for large values of T :

Corollary 2.2 *Suppose that $\rho(A_*) < 1$. Then the limit $\Gamma_\infty := \lim_{t \rightarrow \infty} \Gamma_t$ exists, and there is a time T_0 depending on A_* and δ such that the following holds w.p. $1 - \delta$ for all $T > T_0$:*

$$\left\| \widehat{A}(T) - A_* \right\|_{\text{op}} \leq \mathcal{O} \left(\sqrt{\frac{d \cdot \log\left(\frac{d}{\delta}\right)}{T \lambda_{\min}(\Gamma_\infty)}} \right). \quad (2.4)$$

The above corollary uses the fact that $\lim_{k \rightarrow \infty} \|A_*^k\|_{\text{op}}^{1/k} = \rho(A_*)$, which implies that the limit $\lim_{t \rightarrow \infty} \Gamma_t$ is finite when $\rho(A_*) < 1$. The rate of the convergence of Γ_t to Γ_∞ is related to the \mathcal{H}_∞ -norm of the linear system, a core concept in control theory. For an extended discussion on this relationship, we direct the reader to Tu et al. (2017). Corollaries A.2 and A.3 in the appendix give an analogue of Corollary 2.2 which holds even if $\rho(A_*) = 1$. We now explicitly describe the consequences of Theorem 2.1 for three illustrative classes of linear systems:

1. **Scalar linear system.** In this case the states X_t and the parameter A_* are scalars, and denoted $a_* = A_*$. For $|a_*| \leq 1$, we can apply Theorem 2.1 with block length $k = \mathcal{O}(T/\log(1/\delta))$. This then guarantees that $|\widehat{a} - a_*| \leq \mathcal{O} \left(\sqrt{\log(1/\delta) / \left(T \sum_{t=1}^{k_*} a_*^{2t} \right)} \right)$ with probability $1 - \delta$. In Appendix B, we show this statistical rate is minimax optimal (Theorem B.2). Moreover, we

offer a specialized analysis for the scalar case (Theorem B.1) which yields sharper constants and also applies to the unstable case $|a_*| > 1$, matching the lower bounds of Theorem B.2. Stated succinctly, our results in Appendix B imply that the OLS estimator satisfies with probability $1 - \delta$ error guarantees which can be categorized into three regimes:

$$|\hat{a} - a_*| = \begin{cases} \Theta \left(\sqrt{\frac{\log(1/\delta)(1-|a_*|)}{T}} \right) & \text{if } |a_*| \leq 1 - \frac{c \log(1/\delta)}{T}, \\ \Theta \left(\frac{\log(1/\delta)}{T} \right) & \text{if } 1 - \frac{c \log(1/\delta)}{T} < |a_*| \leq 1 + \frac{1}{T}, \\ \Theta \left(\frac{\log(1/\delta)}{|a_*|^T} \right) & \text{if } 1 + \frac{1}{T} \leq |a_*|. \end{cases}$$

White (1958) showed the same dependence on $|a_*|$ of the estimation error by characterizing the asymptotic distribution of $\hat{a} - a_*$ when appropriately scaled. However, our results offer finite sample guarantees.

2. **Scaled orthogonal systems.** Let us assume $A_* = \rho \cdot O$ for an orthogonal $d \times d$ matrix O and $|\rho| \leq 1$. Then, one can verify that $\Gamma_t = I \cdot \sum_{s=0}^{t-1} \rho^{2s}$ and that we can choose the block length $k = \mathcal{O} \left(\frac{T}{d \log(d/\delta)} \right)$. Therefore, Theorem 2.1 guarantees that with probability $1 - \delta$:

$$\|\hat{A} - A_*\|_{\text{op}} \leq \begin{cases} \mathcal{O} \left(\sqrt{(1-|\rho|) \cdot \frac{d \log(d/\delta)}{T}} \right) & \text{if } |\rho| \leq 1 - \frac{cd \log(d/\delta)}{T}, \\ \mathcal{O} \left(\frac{d \log(d/\delta)}{T} \right) & \text{if } 1 - \frac{cd \log(d/\delta)}{T} < |\rho|. \end{cases} \quad (2.5)$$

3. **Diagonalizable linear systems.** Let $A_* = SDS^{-1}$ define a diagonalizable linear system. We denote by $\underline{\rho}$ the smallest magnitude of an eigenvalue of A_* . In Appendix A, we show that we can choose the block length k such that $k \geq \frac{T}{cd \log \left(\frac{d \text{cond}(S)}{\delta} \right)}$. With this choice of k the OLS estimator satisfies (Corollary A.3)

$$\mathbb{P} \left[\|\hat{A} - A_*\|_{\text{op}} \leq \mathcal{O} \left(\sqrt{\frac{d \log(d \text{cond}(S)/\delta)}{T \left(1 + \text{cond}(S)^{-2} \sum_{s=0}^{k-1} \underline{\rho}^{2s} \right)}} \right) \right] \geq 1 - \delta$$

which could once again be split into a slow and fast rate, as in the examples presented above, depending on the size $\underline{\rho}$ of the least excitable mode of the system defined by A_* . Note that up to a factor of $\log(d \text{cond}(S)/\delta)$, the above bound is no worse than the worst-case rate for standard random-design least-squares in the operator norm (see Appendix G).

Remark 1 (Noise dependence) *As mentioned before, the estimation guarantee provided by Theorem 2.1 does not depend on the variance σ^2 of the noise η_t . For Gaussian noise with a general identity covariance $\eta_t \sim \mathcal{N}(0, \Sigma)$, one can rederive rates from our more general Theorem 2.4 to get a more precise dependence on Γ_t and Σ . Note that if the covariance Σ is known, an alternative estimator would be to choose \hat{A} to minimize a loss which takes Σ into account in the same way that one would for non-dynamic linear regression with heteroskedastic noise, e.g. $\hat{A}^\Sigma(T) := \arg \min_{A \in \mathbb{R}^{d \times d}} \sum_{t=1}^T \frac{1}{2} \|\Sigma^{-1/2} (X_{t+1} - AX_t)\|_2^2$.*

Remark 2 (Learning with input sequences) We can also consider the case where the linear system $X_{t+1} = A_*X_t + B_*u_t + \eta_t$ is driven by a known sequence of inputs u_0, u_1, \dots , with known B_* . Defining the control Gramian $\Gamma_t^{B_*} := \sum_{s=1}^t A_*^{t-s} B_* B_*^\top A_*^{t-s}$, the proof of Theorem 2.1 can be modified to show that, if the inputs are white noise $u_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_u^2 I)$, then there exist universal constants $c, C > 0$ such that, with probability $1 - \delta$,

$$\|\widehat{A}(T) - A_*\|_{\text{op}} \leq \frac{C\sigma^2}{\sqrt{T\lambda_{\min}(\sigma^2\Gamma_k + \sigma_u^2\Gamma_k^{B_*})}} \sqrt{d \log \left(\frac{1}{\delta} \frac{\text{tr}(\sigma^2\Gamma_T + \sigma_u^2\Gamma_T^{B_*})}{\lambda_{\min}(\sigma^2\Gamma_k + \sigma_u^2\Gamma_k^{B_*})} \right)}$$

for any k such that $\frac{T}{k} \geq cd \log \left(\frac{\text{tr}(\sigma^2\Gamma_T + \sigma_u^2\Gamma_T^{B_*})}{\delta\lambda_{\min}(\sigma^2\Gamma_k + \sigma_u^2\Gamma_k^{B_*})} \right)$. Process noise with covariance not equal to a multiple of the identity can be absorbed into B_* .

2.2. Lower Bounds for Linear System Identification

We have seen in Theorem 2.1 and in the subsequent examples that the estimation of linear dynamical systems is easier for systems which are easily excitable. It is natural to ask what is the best possible estimation rate one can hope to achieve. To make explicit the dependence of the lower bounds on the spectrum of Γ_t , we consider the minimax rate of estimation over the set $\rho \cdot \text{O}(d)$, where $\rho \in \mathbb{R}$ and $\text{O}(d)$ denotes the orthogonal group. In this case, we can define a *scalar* Gramian $\gamma_t(\rho) := \sum_{s=0}^{t-1} |\rho|^{2s}$, so that $\Gamma_t := \gamma_t(\rho) \cdot I$. We now show that the estimation rate of the OLS provided in Theorem 2.1 is optimal up to log factors for $|\rho| \leq 1 - \widetilde{\mathcal{O}}(d/T)$:

Theorem 2.3 Fix $d \geq 2$, $\rho \in \mathbb{R}$, $\delta \in (0, 1/4)$, and $\epsilon \leq \frac{\rho}{2048}$. Then, there exists a universal constant c_0 such for any estimator \widehat{A} ,

$$\sup_{\rho \in \text{O}(d)} \mathbb{P}_{\rho \text{O}} \left[\left\| \widehat{A}(T) - \rho \text{O} \right\|_{\text{op}} \geq \epsilon \right] \geq \delta \text{ for any } T \text{ such that } T\gamma_T(\rho) \leq \frac{c_0(d + \log(1/\delta))}{\epsilon^2},$$

where $\text{O}(d)$ is the orthogonal group of $d \times d$ real matrices.

This result is proved in Appendix F.1. We can interpret it by considering the following regimes:

$$\|\widehat{A} - A_*\|_{\text{op}} \geq \begin{cases} \Omega \left(\sqrt{\frac{(d + \log(1/\delta)) \cdot (1 - |\rho|)}{T}} \right) & \text{if } |\rho| \leq 1 - \frac{1}{T}, \\ \Omega \left(\frac{\sqrt{d + \log(1/\delta)}}{T} \right) & \text{if } 1 - \frac{1}{T} < |\rho| < 1 + \frac{1}{T} \\ \Omega \left(\frac{\sqrt{d + \log(1/\delta)}}{T|\rho|^T} \right) & \text{if } 1 + \frac{1}{T} \leq |\rho|. \end{cases}$$

Comparing to (2.5), we see that for $|\rho| \leq 1 - \widetilde{\mathcal{O}}(d/T)$, our upper and lower bounds coincide up to logarithmic factors. In the regime $\rho \in [1 - \widetilde{\mathcal{O}}(d/T), 1]$, our upper and lower bounds differ by a factor of $\widetilde{\mathcal{O}}(\sqrt{d + \log(1/\delta)})$.

2.3. General Time Series with Linear Responses

In this section, we consider a sequence of covariate-response pairs $(X_t, Y_t)_{t \geq 1}$, where $Y_t = A_* X_t + \eta_t$, with $Y_t, \eta_t \in \mathbb{R}^n$, $X_t \in \mathbb{R}^d$, and $A_* \in \mathbb{R}^{n \times d}$. The least squares estimator is then

$$\widehat{A}(T) := \arg \min_{A \in \mathbb{R}^{d \times d}} \sum_{t=1}^T \frac{1}{2} \|Y_t - AX_t\|_2^2. \quad (2.6)$$

We let $\mathcal{F}_t := \sigma(\eta_0, \eta_1, \dots, \eta_t, X_1, \dots, X_t)$ denote the filtration generated by the covariates and noise process. Note then that $Y_t \in \mathcal{F}_t$, but $Y_t \notin \mathcal{F}_{t-1}$. Further, we assume $\eta_t | \mathcal{F}_{t-1}$ is mean-zero, and σ^2 -sub-Gaussian (i.e., $\mathbb{E}[\exp(\lambda \eta_t) | \mathcal{F}_t] \leq e^{\sigma^2 \lambda^2 / 2}$). The linear dynamical systems sub-case is recovered from this general setting when $Y_t = X_{t+1}$.

To capture the excitation behavior observed in the case of linear systems we introduce a general martingale small-ball condition which quantifies the growth of the covariates X_t .

Definition 2.1 (Martingale Small-Ball) *Let $(Z_t)_{t \geq 1}$ be an $\{\mathcal{F}_t\}_{t \geq 1}$ -adapted random process taking values in \mathbb{R} . We say $(Z_t)_{t \geq 1}$ satisfies the (k, ν, p) -block martingale small-ball (BMSB) condition if, for any $j \geq 0$, one has $\frac{1}{k} \sum_{i=1}^k \mathbb{P}(|Z_{j+i}| \geq \nu | \mathcal{F}_j) \geq p$ almost surely. Given a process $(X_t)_{t \geq 1}$ taking values in \mathbb{R}^d , we say that it satisfies the $(k, \Gamma_{\text{sb}}, p)$ -BMSB condition for $\Gamma_{\text{sb}} \succ 0$ if, for any fixed $w \in \mathcal{S}^{d-1}$, the process $Z_t := \langle w, X_t \rangle$ satisfies $(k, \sqrt{w^\top \Gamma_{\text{sb}} w}, p)$ -BMSB.*

Such a small-ball condition is necessary for establishing a high-probability lower bound on $\lambda_{\min}(\sum_{t=1}^T X_t X_t^\top) = \min_{w \in \mathcal{S}^{d-1}} \sum_{t=1}^T \langle X_t, w \rangle^2$. The parameter Γ_{sb} plays the role of the Gramians Γ_t considered in the case of linear systems, and measures how excitable the covariates X_t are. As expected, the next result shows that a higher $\lambda_{\min}(\Gamma_{\text{sb}})$ leads to faster statistical estimation.

Theorem 2.4 *Fix $\epsilon, \delta \in (0, 1)$, $T \in \mathbb{N}$ and $0 \prec \Gamma_{\text{sb}} \preceq \bar{\Gamma}$. Then if $(X_t, Y_t)_{t \geq 1} \in (\mathbb{R}^d \times \mathbb{R}^n)^T$ is a random sequence such that (a) $Y_t = A_* X_t + \eta_t$, where $\eta_t | \mathcal{F}_t$ is σ^2 -sub-Gaussian and mean zero, (b) X_1, \dots, X_T satisfies the $(k, \Gamma_{\text{sb}}, p)$ -small ball condition, and (c) such that $\mathbb{P}[\sum_{t=1}^T X_t X_t^\top \not\prec T\bar{\Gamma}] \leq \delta$. Then if*

$$T \geq \frac{10k}{p^2} \left(\log \left(\frac{1}{\delta} \right) + 2d \log(10/p) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1}) \right), \quad (2.7)$$

we have

$$\mathbb{P} \left[\left\| \widehat{A}(T) - A_* \right\|_{\text{op}} > \frac{90\sigma}{p} \sqrt{\frac{n + d \log \frac{10}{p} + \log \det \bar{\Gamma} \Gamma_{\text{sb}}^{-1} + \log \left(\frac{1}{\delta} \right)}{T \lambda_{\min}(\Gamma_{\text{sb}})}} \right] \leq 3\delta. \quad (2.8)$$

The proof of Theorem 2.4 is outlined in Section 4, and technical details are deferred to Appendix D. We remark that the conclusion of Theorem 2.4 still holds if one replaces the $(k, \Gamma_{\text{sb}}, p)$ small-ball condition with any high probability lower bound of the form $\mathbb{P} \left(\sum_{t=1}^T X_t X_t^\top \succeq T\Gamma_{\text{sb}} \right) \leq \delta$.

2.4. Analysis Techniques

Let $\widehat{A} = \widehat{A}(T)$, let $\mathbf{X} \in \mathbb{R}^{T \times d}$ denote the matrix whose rows are X_t , and $\mathbf{E} \in \mathbb{R}^{T \times n}$ denote the matrix whose rows are η_t . Consider the compact SVD of \mathbf{X} and $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$, where $\mathbf{\Sigma}, \mathbf{V} \in \mathbb{R}^{d \times d}$ and $\mathbf{U} \in \mathbb{R}^{T \times d}$. Note then that we have $\widehat{A} - A_* = (\mathbf{X}^\dagger \mathbf{E})^\top$ which implies that

$$\|\widehat{A} - A_*\|_{\text{op}} = \|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \leq \sigma_d(\mathbf{X})^{-1} \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}}. \quad (2.9)$$

Here $\sigma_d((\mathbf{X})^{-1})$ denotes the d -th largest singular value of \mathbf{X} , which is precisely $\sqrt{1/\lambda_{\min}(\mathbf{X}^\top \mathbf{X})}$. The technical challenge arises from the fact that the singular space \mathbf{U}^\top and \mathbf{E} are correlated, and that the rows X_t of \mathbf{X} are also dependent. We upper bound $\|\mathbf{U}^\top \mathbf{E}\|_{\text{op}}$ with Lemma 4.14, a martingale-Chernoff bound that gives precise control on the deviations of sub-Gaussian martingale sequences in terms of random variance proxies. We explain this argument in more detail at the end of Section 4.

Our lower bound on $\sigma_{\min}(\mathbf{X}) = \sqrt{\lambda_{\min}(\sum_{t=1}^T X_t X_t^\top)}$ eschews mixing-time arguments in favor of a careful modification of Mendelson's small-ball method (Mendelson, 2014). We divide our covariates into size- k blocks $\{X_{(\ell-1)k+1}, \dots, X_{\ell k}\}$, such that for any fixed $w \in \mathcal{S}^{d-1}$, the quantity $\sum_{\ell=1}^k \langle X_{(\ell-1)k+1}, w \rangle^2$ can be lower bounded by the $(k, \Gamma_{\text{sb}}, p)$ -BMSB condition. Proposition 2.5 below (proved in Appendix E.1) then implies that we have $\sum_{t=1}^T \langle X_t, w \rangle^2 \gtrsim T w^\top \Gamma_{\text{sb}} w$ with probability at least $1 - \exp(-cT/k)$ for some constant c :

Proposition 2.5 *Let $\{Z_t\}_{t \geq 1}$ be a scalar process that satisfies the (k, ν, p) -BMSB condition. Then*

$$\mathbb{P} \left[\sum_{i=1}^T Z_i^2 \leq \frac{\nu^2 p^2}{8} k \lfloor T/k \rfloor \right] \leq e^{-\frac{\lfloor T/k \rfloor p^2}{8}}. \quad (2.10)$$

Once T is large enough, these high-probability bounds can be used to derive a uniform bound over $w \in \mathcal{S}^{d-1}$ via a discretization argument (Lemma 4.1). In general there is a trade-off between the size of the blocks k and the probability guarantee obtained: a larger block size leads to a larger parameter ν and a faster rate, but it degrades the probability guarantee.

3. Theorem 2.1 as a corollary of Theorem 2.4

In this section, we show how to obtain Theorem 2.1 as a fairly straightforward consequence of our meta-theorem, Theorem 2.4. By assumption, our noise process satisfies the σ^2 -sub-Gaussian tail condition. Moreover, we see that

$$\begin{aligned} \mathbb{P} \left[\mathbf{X}^\top \mathbf{X} \not\geq \frac{\sigma^2 d}{\delta} T \Gamma_T \right] &= \mathbb{P} \left[\lambda_{\max}((T \Gamma_T)^{-1/2} \mathbf{X}^\top \mathbf{X} (T \Gamma_T)^{-1/2}) \geq \frac{\sigma^2 d}{\delta} \right] \\ &\leq \frac{\delta}{\sigma^2 d} \cdot \mathbb{E}[\lambda_{\max}((T \Gamma_T)^{-1/2} \mathbf{X}^\top \mathbf{X} (T \Gamma_T)^{-1/2})] \\ &\leq \frac{\delta}{d \sigma^2} \cdot \mathbb{E}[\text{tr}((T \Gamma_T)^{-1/2} \mathbf{X}^\top \mathbf{X} (T \Gamma_T)^{-1/2})] \leq \delta \end{aligned}$$

where the last inequality follows because $\mathbb{E}[\mathbf{X}^\top \mathbf{X}] = \sigma^2 \sum_{t=1}^T \Gamma_t \preceq \sigma^2 T \Gamma_T$. Hence, we can set $\bar{\Gamma} = \frac{\sigma^2 d}{\delta} \Gamma_T$. Noting $\log \det((\frac{d}{\delta} \sigma^2 \Gamma_T)(\sigma^2 \Gamma_{\lfloor k/2 \rfloor})^{-1}) = d \log d/\delta + \log \det(\Gamma_T \Gamma_{\lfloor k/2 \rfloor}^{-1})$, it suffices to verify that $(X_t)_{t \geq 1}$ satisfies the $(k, \sigma^2 \Gamma_{\lfloor k/2 \rfloor}, \frac{3}{20})$ -BMSB condition.

Proposition 3.1 *Consider the linear dynamical system $X_{t+1} = A X_t + \eta_t$, where $X_0 \in \mathbb{R}^d$ and $\eta_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2 I)$, and let $\Gamma_t := \sum_{s=0}^{t-1} (A^s)(A^s)^\top$. Then, for $1 \leq k \leq T$, the process $(X_t)_{t \geq 1}$ satisfies the*

$$\left(k, \sigma^2 \Gamma_{\lfloor k/2 \rfloor}, \frac{3}{20} \right) \text{-block martingale small-ball condition.} \quad (3.11)$$

Proof Let $1 \leq k' \leq k$. Note that, for $t \geq 1$, $X_{s+t} | \mathcal{F}_s \sim \mathcal{N}(\langle w, A^t X_0 \rangle, \sigma^2 w^\top \Gamma_t w)$. Hence, for $t \geq k'$, one has

$$\mathbb{P} \left(|\langle w, X_{s+t} \rangle| \geq \sigma \sqrt{w^\top \Gamma_{k'} w} | \mathcal{F}_s \right) \stackrel{(i)}{\geq} \mathbb{P} \left(|\langle w, X_t \rangle| \geq \sigma \sqrt{w^\top \Gamma_t w} | \mathcal{F}_s \right) \stackrel{(ii)}{\geq} \frac{3}{10},$$

where (i) uses the fact that $\Gamma_t \succeq \Gamma_{k'}$ for $t \geq k'$, and (ii) follows from the Paley-Zygmund lower bound,

$$\forall t \in \mathbb{R}, \quad \mathbb{P}_{Z \sim \mathcal{N}(0, \sigma^2)} [|t + Z| \geq \sigma] \geq \mathbb{P}[|Z| \geq \sigma] \geq 3/10. \quad (3.12)$$

Therefore, we have

$$\frac{1}{k} \sum_{t=1}^k \mathbb{P} \left(|\langle w, X_t \rangle| \geq \sigma \sqrt{w^\top \Gamma_{k'} w} \right) \geq \frac{1}{k} \sum_{t=k'}^k \mathbb{P} \left(|\langle w, X_t \rangle| \geq \sigma \sqrt{w^\top \Gamma_{k'} w} \right) \geq \frac{3}{10} \frac{k - k' + 1}{k}.$$

Finally, choosing $k' = \lfloor k/2 \rfloor$ yields the desired conclusion. \blacksquare

4. Proof of Theorem 2.4

Again, we let $\mathbf{X} \in \mathbb{R}^{T \times d}$ denote the matrix whose rows are X_t , and $\mathbf{E} \in \mathbb{R}^{T \times n}$ denote the matrix whose rows are η_t , and consider the compact SVD of \mathbf{X} and $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$, where $\mathbf{\Sigma}, \mathbf{V} \in \mathbb{R}^{d \times d}$ and $\mathbf{U} \in \mathbb{R}^{T \times d}$. Recalling (2.9), we have $\|\hat{A}(T) - A_\star\|_{\text{op}} \leq \sigma_d(\mathbf{X})^{-1} \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}}$. Let K denote a threshold parameter to be chosen later. Then $\|\hat{A}(T) - A_\star\|_{\text{op}} \leq \sigma_d(\mathbf{X})^{-1} \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}}$ implies the following set-theoretic inclusions,

$$\begin{aligned} & \left\{ \|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right\} \cap \left\{ \mathbf{X} \mathbf{X}^\top \succeq \frac{k \lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}}{16} \right\} \\ & \subseteq \left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \sigma_{\min}(\mathbf{X}) \right\} \cap \left\{ \mathbf{X} \mathbf{X}^\top \succeq \frac{k \lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}}{16} \right\} \\ & \subseteq \left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} \geq K \right\} \cap \left\{ \mathbf{X} \mathbf{X}^\top \succeq \frac{k \lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}}{16} \right\}. \end{aligned}$$

Now define the following events

$$\mathcal{E}_1 := \left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} \geq K \right\}, \quad \mathcal{E}_2 := \left\{ \mathbf{X}^\top \mathbf{X} \succeq \frac{k \lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}}{16} \right\}, \quad \mathcal{E}_3 := \left\{ \mathbf{X}^\top \mathbf{X} \not\preceq \bar{\Gamma} \right\}.$$

Then we have

$$\begin{aligned} & \mathbb{P} \left[\|\hat{A}(T) - A_\star\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right] \leq \mathbb{P} \left[\left\{ \|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right\} \right] \\ & \leq \mathbb{P} \left[\left\{ \|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right\} \cap \mathcal{E}_3^c \right] + \mathbb{P}[\mathcal{E}_3] \\ & \leq \mathbb{P} \left[\left\{ \|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k} \lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right\} \cap \mathcal{E}_2 \cap \mathcal{E}_3^c \right] + \mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c] + \mathbb{P}[\mathcal{E}_3] \\ & \leq \mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c] + \mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c] + \mathbb{P}[\mathcal{E}_3]. \end{aligned}$$

By assumption $\mathbb{P}[\mathcal{E}_3] \leq \delta$. Our task is to show that both $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c]$ and $\mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c]$ are upper bounded by δ for the choice of $K = 20\sigma\sqrt{n + d\log(10/p) + \log\det(\bar{\Gamma}\Gamma_{\text{sb}}^{-1}) + \log(1/\delta)}$. Both bounds are proven in detail in Appendix D, but here we state the main technical arguments required for their proof. All supplementary technical results (Lemma 4.1, Lemma 4.2 and Proposition 2.5) are proven in Appendix E.

Our bound on $\mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c]$ comes directly from the BMSB assumption. Applying Proposition 2.5, we have that for any fixed $w \in \mathcal{S}^{d-1}$, $\mathbb{P}\left[\sum_{t=1}^T \langle w, X_t \rangle^2 \leq \frac{w^\top \Gamma_{\text{sb}} w p^2}{8} k \lfloor T/k \rfloor\right] \leq e^{-\frac{\lfloor T/k \rfloor p^2}{8}}$. To obtain a Lowner lower-bound $\mathbf{X}^\top \mathbf{X}$, we need to strengthen the above pointwise bound into a lower bound on $\inf_{w \in \mathcal{S}^{d-1}} \sum_{t=1}^T \langle w, X_t \rangle^2$. This is achieved through the following covering lemma, proved in Appendix E.3.1:

Lemma 4.1 *Let $Q \in \mathbb{R}^{T \times d}$ and consider matrices $0 \prec \Gamma_{\min} \preceq \Gamma_{\max} \in \mathbb{R}^{d \times d}$. Let \mathcal{T} be a $1/4$ -net of $\mathcal{S}_{\Gamma_{\min}}$ in the metric $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$. Then, if $\inf_{w \in \mathcal{T}} w^\top Q^\top Q w \geq 1$ and $Q^\top Q \preceq \Gamma_{\max}$, we have*

$$Q^\top Q \succeq \Gamma_{\min}/2. \quad (4.13)$$

Choosing $Q = \mathbf{X}\mathbf{X}^\top$, this lemma gives us a bound on the granularity at which we need to cover \mathcal{S}^{d-1} in terms of a uniform Lowner upper bound $\Gamma_{\max} = T\bar{\Gamma}$, and pointwise Lowner lower bound with $\Gamma_{\min} = \frac{\Gamma_{\text{sb}} p^2}{8} k \lfloor T/k \rfloor$. The details are worked through in Appendix D.1. The larger Γ_{\max} is compared to Γ_{\min} in a Lowner sense, the larger the cardinality of the net \mathcal{T} must be. Crucially, this term enters logarithmically into our final bound via the relative volume $\log\det(\Gamma_{\max}\Gamma_{\min}^{-1})$ of the ellipsoids induced by Γ_{\max} and Γ_{\min} .

Lastly, we bound the probability of $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c$, which is the event that $\{\|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} \geq K\}$ when the spectrum of $\mathbf{X}^\top \mathbf{X}$ is bounded in some desired range. \mathbf{U} is difficult to control directly. Instead, we work with quantities in terms of \mathbf{X} , namely $\|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} = \sup_{v \in \mathcal{S}^{n-1}, w \in \mathbb{R}^d} \frac{w^\top \mathbf{X} \mathbf{E} v}{\|\mathbf{X} w\|}$. The key idea here now is to use a martingale-Chernoff bound to show that, for any fix $w \in \mathcal{S}^{d-1}$ and $v \in \mathcal{S}^{n-1}$, either $w^\top \mathbf{X}^\top \mathbf{E} v$ concentrates like a $\sigma^2 \|\mathbf{X} w\|_2^2$ -sub-Gaussian random variable, or $w^\top \mathbf{X}^\top \mathbf{E} v$ is much smaller than the lower bound on $\sigma_{\min}(\mathbf{X})$ under $\mathcal{E}_2 \cap \mathcal{E}_3^c$.

We emphasize that our bound controls $w^\top \mathbf{X}^\top \mathbf{E} v$ in terms of $\sigma^2 \|\mathbf{X} w\|_2^2$, which is a proxy for variance. This is subtle yet powerful because it yields an immediate cancellation between the numerator and denominator of $\frac{w^\top \mathbf{X} \mathbf{E} v}{\|\mathbf{X} w\|}$. In particular, this implies that $\mathbb{P}(\{\frac{w^\top \mathbf{X} \mathbf{E} v}{\|\mathbf{X} w\|} \gtrsim \log(1/\delta)\} \cap \mathcal{E}_2 \cap \mathcal{E}_3^c) \lesssim \delta$. This lets us reduce our problem to finding an appropriate discretization (see Lemma D.2). We stress that an approach which bounds $w^\top \mathbf{X} \mathbf{E} v$ and $\|\mathbf{X} w\|$ separately would be considerably less sharp, and would degrade for slower-mixing systems. Our data-dependent concentration bound is a consequence of the following technical lemma, which we apply with $Z_t := w^\top X_t$ and $W_t = \eta_t^\top v$.

Lemma 4.2 *Let $\{\mathcal{F}_t\}_{t \geq 0}$ be a filtration, and $\{Z_t\}_{t \geq 1}$ and $\{W_t\}_{t \geq 1}$ be real-valued processes adapted to \mathcal{F}_t and \mathcal{F}_{t+1} respectively. Moreover, assume $W_t | \mathcal{F}_t$ is mean zero and σ^2 -sub-Gaussian. Then, for any positive real numbers $\alpha, \beta, \beta_-, \beta_+$ we have*

$$(a) \quad \mathbb{P}\left[\left\{\sum_{t=1}^T Z_t W_t \geq \alpha\right\} \cap \left\{\sum_{t=1}^T Z_t^2 \leq \beta\right\}\right] \leq \exp\left(-\frac{\alpha^2}{2\sigma^2\beta}\right). \quad (4.14)$$

$$(b) \quad \mathbb{P}\left[\left\{\frac{\sum_{t=1}^T Z_t W_t}{\sqrt{\sum_{t=1}^T Z_t^2}} > \alpha\right\} \cap \left\{\sum_{t=1}^T Z_t^2 \in [\beta_-, \beta_+]\right\}\right] \leq \log\left[\frac{\beta_+}{\beta_-}\right] \exp\left(\frac{-\alpha^2}{6\sigma^2}\right) \quad (4.15)$$

5. Discussion and future work

In this paper, we analyzed the performance of the OLS estimator for the estimation of linear dynamics $X_{t+1} = A_*X_t + \eta_t$ from a single trajectory X_0, X_1, \dots, X_T , as a special case of linear estimation in time series. We show that, up to logarithmic factors, the OLS estimator attains an information-theoretic lower bound for $\rho(A_*) < 1$, provided that $T \gtrsim \frac{d}{1-\rho(A_*)}$. Moreover, we present an analysis that eschews both mixing and concentration arguments for estimation in time series. We believe that there are several promising directions for future work:

- Our lower and upper bounds do not perfectly match, even when $\rho(A_*) < 1$. We believe resolving these discrepancies may shed greater insight into learning in dynamical systems.
- While our guarantees are stated in the operator norm, control applications may require more granular notions of error which vary for different modes of A_* . Developing error bounds which capture the error rate at each mode may result in more applicable bounds for control applications downstream.
- While our analysis can accommodate an unknown B_* as a consequence of Theorem 2.4, the resulting rates do not distinguish between the error in the estimation of A_* and that of B_* . In future, we hope to develop sharp error rates for A_* and B_* individually, similar to Dean et al. (2017) in the independent covariates setting.
- Our convergence rates degrade for systems with $\rho(A_*) > 1$, whereas we know from Faradonbeh et al. (2017a) that a large class of these systems are still identifiable with OLS. Is there a unified analysis for systems with stable and unstable modes?
- In many systems, we do not observe X_t directly, but only view CX_t for a matrix $C \in \mathbb{R}^{n_o \times n}$, where $n_o \leq n$. Hazan et al. (2017, 2018) provide filtering techniques to minimize regret for diagonalizable matrices; it would be interesting to understand the sample complexity for estimating arbitrary matrices with these limited observations.
- Ultimately, we would like to understand what sequences of control inputs u_t yield the most accurate estimation of the system (A_*, B_*) . This would inform adaptive algorithms which adjust the sequence u_t in a sequential fashion, and online algorithms which ensure low regret relative to a given cost functional over time.

Acknowledgements

MS is supported by an NSF GRFP award. As part of the RISE lab, MS and HM are supported in part by NSF CISE Expeditions Award CCF-1730628, DHS Award HSHQDC-16-3-00083, and gifts from Alibaba, Amazon Web Services, Ant Financial, CapitalOne, Ericsson, GE, Google, Huawei, Intel, IBM, Microsoft, Scotiabank, Splunk and VMware. BR is generously supported in part by NSF award CCF-1359814, ONR awards N00014-14-1-0024 and N00014-17-1-2191, the DARPA Fundamental Limits of Learning (Fun LoL) Program, and an Amazon AWS AI Research Award. BR is generously supported in part by NSF award CCF-1359814, ONR awards N00014-17-1-2191, N00014-17-1-2401, and N00014-17-1-2502, the DARPA Fundamental Limits of Learning (Fun LoL) and Lagrange Programs, and an Amazon AWS AI Research Award. This work was supported in part by the Mathematical Data Science program of the Office of Naval Research under grant number N00014-15-1-2670.

References

- S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. 2013.
- Marco C Campi and Erik Weyer. Finite Sample Properties of System Identification Methods. *IEEE Transactions on Automatic Control*, 47(8), 2002.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the Sample Complexity of the Linear Quadratic Regulator. *arXiv:1710.01688*, 2017.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite Time Identification in Unstable Linear Systems. *arXiv:1710.01852*, 2017a.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite Time Analysis of Optimal Adaptive Policies for Linear-Quadratic Systems. *arXiv:1711.07230*, 2017b.
- Moritz Hardt, Tengyu Ma, and Benjamin Recht. Gradient Descent Learns Linear Dynamical Systems. *arXiv:1609.05191*, 2016.
- Elad Hazan, Karan Singh, and Cyril Zhang. Learning Linear Dynamical Systems via Spectral Filtering. In *Neural Information Processing Systems*, 2017.
- Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral Filtering for General Linear Dynamical Systems. 2018.
- Daniel Hsu, Sham M. Kakade, and Tong Zhang. Random Design Analysis of Ridge Regression. *Foundations of Computational Mathematics*, 14, 2014.
- Anthony W Knapp. *Representation Theory of Semisimple Groups: An Overview Based on Examples (PMS-36)*. Princeton University Press, 2016.
- Vitaly Kuznetsov and Mehryar Mohri. Generalization Bounds for Non-Stationary Mixing Processes. *Machine Learning*, 106(1), 2017.
- Daniel J. McDonald, Cosma R. Shalizi, and Mark Schervish. Nonparametric Risk Bounds for Time-Series Forecasting. *Journal of Machine Learning Research*, 18, 2017.
- Shahar Mendelson. Learning without Concentration. In *Conference on Learning Theory*, 2014.
- Mehryar Mohri and Afshin Rostamizadeh. Stability Bounds for Non-i.i.d. Processes. In *Neural Information Processing Systems*, 2007a.
- Mehryar Mohri and Afshin Rostamizadeh. Rademacher Complexity Bounds for Non-I.I.D. Processes. In *Neural Information Processing Systems*, 2007b.
- Anders Rantzer. Concentration Bounds for Single Parameter Adaptive Control. In *American Control Conference*, 2018.
- Parikshit Shah, Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. Linear System Identification via Atomic Norm Regularization. In *Conference on Decision and Control*, 2012.

Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht. Non-Asymptotic Analysis of Robust Control from Coarse-Grained Identification. *arXiv:1707.04791*, 2017.

Roman Vershynin. Introduction to the Non-Asymptotic Analysis of Random Matrices. *arXiv:1011.3027*, 2011.

Mathukumalli Vidyasagar and Rajeeva L Karandikar. A learning theory approach to system identification and stochastic adaptive control. *Journal of Process Control*, 18(3), 2008.

John S White. The limiting distribution of the serial correlation coefficient in the explosive case. *The Annals of Mathematical Statistics*, pages 1188–1197, 1958.

Bin Yu. Rates of Convergence for Empirical Processes of Stationary Mixing Sequences. *The Annals of Probability*, 22(1), 1994.

Appendix A. Existence of k for $\rho(A_*) \leq 1$

Proposition A.1 *Let $A_* = SJS^{-1}$ and $T \geq d$, where J has block sizes b_1, \dots, b_L . Then,*

$$\log \det(\Gamma_T \Gamma_k^{-1}) \leq 2d \log \text{cond}(S) + d \log \frac{T}{k} + 4 \log T \sum_{\ell: b_\ell \geq 2} b_\ell^2 \quad (1.16)$$

where $\text{cond}(S)$ denotes the complex condition number of S , namely $\sqrt{\lambda_{\max}(S^*S)/\lambda_{\min}(S^*S)}$.

The above proposition directly implies consistency for whenever $\rho(A_*) \leq 1$:

Corollary A.2 *[Consistency of Least Squares] There exists a universal constants $C, c > 0$ such that for any time horizon T , $\delta \in (0, 1/2)$, and any A_* with $\rho(A_*) \leq 1$ and Jordan decomposition $A_* = SJS^{-1}$, where J has block sizes b_1, \dots, b_L , $k \in \mathbb{N}$ satisfies the conditions of Theorem 2.1 provided that*

$$\frac{T}{k} \geq c \left(d \log \left(\frac{d \text{cond}(S) T}{k \delta} \right) + \log T \sum_{\ell: b_\ell \geq 2} b_\ell^2 \right), \quad (1.17)$$

Taking $k = 1$, implies a minimax-rate of estimation of

$$\|A - A^*\|_{\text{op}} \leq C \sqrt{\frac{\log(dT \text{cond}(S)/\delta) + \log T \sum_{\ell: b_\ell \geq 2} b_\ell^2}{T}} \quad (1.18)$$

which holds as long as $T \geq c(d \log(\text{cond}(S) \frac{d}{k\delta}) + \sum_{\ell: b_\ell \geq 1} b_\ell^2 \log(\sum_{\ell: b_\ell \geq 1} b_\ell^2))$

Remark *Observe that if the Jordan decomposition is such that $\sum_{\ell: b_\ell \geq 1} b_\ell^2 \lesssim d$, then our minimax rate coincides with the minimax rate of linear regression with isotropic covariates up to logarithmic factors. Moreover, we above that*

For diagonalizable matrices, the rates can be made more explicit:

Corollary A.3 *There exists a universal constants $C, c > 0$ such that the following holds. Fix any time horizon T , $\delta \in (0, 1/2)$, let $A_* = SDS^{-1}$ be diagonalizable with $\rho(A_*) \leq 1$ and minimum eigenvalue-magnitude $\underline{\rho}$. Then, $k \in \mathbb{N}$ satisfies the conditions of Theorem 2.1 provided that*

$$\frac{T}{k} \geq cd \log(\text{cond}(S)/\delta) \quad (1.19)$$

This implies that for $T \geq cd \log(\text{cond}(S)/\delta)$

$$\|A - A^*\|_{\text{op}} \leq C \sqrt{\frac{d \log(\text{cond}(S)/\delta)}{T \Gamma_{\lfloor T/cd \log(\text{cond}(S)/\delta) \rfloor}}} \quad (1.20)$$

$$\leq \sqrt{\frac{d \log(\text{cond}(S)/\delta)}{T(1 + \text{cond}(S))^{-2} \sum_{s=1}^{\lfloor T/cd \log(\text{cond}(S)/\delta) \rfloor - 1} \underline{\rho}^{2s}}} \quad (1.21)$$

A.1. Proof of Corollary A.2

By Theorem 2.1, T and k must satisfy the inequality

$$T/k \geq c(d \log \frac{d}{\delta} + \log \det \Gamma_T \Gamma_k^{-1}) \quad (1.22)$$

Equation (1.17) follows directly from Proposition A.1. Specializing to $k = 1$, Proposition A.1 and Theorem 2.1 immediately imply (1.18). To upper bound the burn-in time for T , we note that by Proposition A.1, the condition (1.22) holds as soon as

$$T \geq c'(d \log \frac{d \text{cond}(S)}{\delta}) \text{ and } T \geq c' \log T(d + \sum_{\ell: b_\ell \geq 2} b_\ell^2) \quad (1.23)$$

for a universal constant c' . We can bound $d + \sum_{\ell: b_\ell \geq 2} b_\ell^2 \leq \sum_\ell b_\ell^2$, where the latter sum is over all Jordan blocks. We now invoke the following lemma, which we prove shortly:

Lemma A.4 *Let $\alpha \geq 1$. Then for any $T \in \mathbb{N}$, $T \geq \alpha \log T$ as soon as $T \geq 2\alpha \log 4\alpha$*

The lemma implies that it is enough to ensure $T \geq c'(d \log \frac{d \text{cond}(S)}{\delta})$ and that $T \geq 2c' \log T(4 \sum_\ell b_\ell^2)$, both of which can be ensured by choosing the constant c in Corollary A.2 to be sufficiently large.

Proof [Proof of Lemma A.4] Taking derivatives, $T \mapsto T - \alpha \log T$ is increasing in T for all $T \geq \alpha$. Hence, it suffices to show that for $T = 2\alpha \log 4\alpha$, $T \geq \alpha \log T$. Observe that for this choice of α ,

$$\begin{aligned} \alpha \log T &= \alpha \log(2\alpha \log(4\alpha)) \\ &\leq \alpha \log((2 \log 4) \cdot \alpha + 2\alpha^2) \\ &\leq \alpha \log((2 \log 4 + 2)\alpha^2) \quad \text{since } \alpha \geq 1 \\ &= 2\alpha \log \sqrt{2 \log 4 + 2} \leq 2\alpha \log 4. \end{aligned}$$

■

A.2. Proof of Corollary A.3

By Theorem 2.1, T and k must satisfy the inequality

$$T/k \geq c(d \log \frac{d}{\delta} + \log \det \Gamma_T \Gamma_k^{-1})$$

Using the upper bound on $\log \det \Gamma_T \Gamma_k^{-1}$ from Proposition A.1, it is enough to ensure that

$$\frac{T}{k} \geq c'(d \log \frac{d \text{cond}(S)}{\delta}) \text{ and } \frac{T}{k} \geq c' d \log \frac{T}{k} \quad (1.24)$$

for some universal constant c' (note that the term $\sum_{\ell: b_\ell \geq 2} b_\ell^2$ vanishes for diagonalizable A_*). By inflating c' , we may assume $c' \geq 1$. Applying Lemma A.4 with change of variables $T \leftarrow T/k$, (1.24) holds as long as $T/k \geq 2c' d \log(4c' d)$, and $T/k \geq c'(d \log \frac{d \text{cond}(S)}{\delta})$, which holds as long as

$$T/k \geq c'''(d \log \frac{d \text{cond}(S)}{\delta}))$$

for some constant c''' . This proves (1.19). We then see that that (1.20) in Corollary A.3 is an immediate consequence of A.3, and (1.20) follows from the Lowner Lower bound, with $A_* = SDS^{-1}$:

$$\begin{aligned}
 \Gamma_T &= I + \sum_{t=1}^{T-1} (A_*^t)(A_*^t)^* \\
 &= I + \sum_{t=1}^{T-1} (SD^t S^{-1})(SD^t S^{-1})^* \\
 &= I + S \sum_{t=1}^{T-1} SD^t S^{-1} S^{-*} D^{t*} S^* \\
 &\succeq I + \lambda_{\min}(S^{-1} S^{-*}) \sum_{t=1}^{T-1} SD^t D^{t*} S^* \\
 &\succeq I + \lambda_{\min}(S^{-1} S^{-*}) S S^* \left(\sum_{t=1}^{T-1} \underline{\rho}^{2t} \right) \\
 &\succeq I + \lambda_{\min}(S^{-1} S^{-*}) \lambda_{\max}(S S^*) \left(\sum_{t=1}^{T-1} \underline{\rho}^{2t} \right) I \\
 &= I(1 + \text{cond}(S)^{-2} \left(\sum_{t=1}^{T-1} \underline{\rho}^{2t} \right))
 \end{aligned}$$

A.3. Proof of Proposition A.1

Let $A_* = SJS^{-1}$, where J is a Jordan-Block matrices with blocks J_1, \dots, J_L of sizes b_1, \dots, b_L . Note that even those A_* is real, S and J may be complex valued, so we shall use adjoints instead of transposes. We can compute

$$\Gamma_t = S \sum_{s=0}^{t-1} (J^s) S^{-1} S^{-*} J^{s*} S^\top$$

Hence,

$$\begin{aligned}
 &\log \det(\Gamma_T \Gamma_k^{-1}) \\
 &= \log \det \left\{ S \sum_{s=0}^{T-1} J^s S^{-1} S^{-*} J^{s*} S^\top \left(S \sum_{s=0}^{k-1} J^s S^{-1} S^{-*} J^{s*} S^\top \right)^{-1} \right\} \\
 &= \log \det \left\{ S \sum_{s=0}^{T-1} J^s S^{-1} S^{-*} J^{s*} \left(\sum_{s=0}^{k-1} (J^s) S^{-1} S^{-*} J^{s*} \right)^{-1} S^{-1} \right\} \\
 &= \log \det \left\{ \sum_{s=0}^{T-1} J^s S^{-1} S^{-*} J^{s*} \left(\sum_{s=0}^{k-1} (J^s) S^{-1} S^{-*} J^{s*} \right)^{-1} \right\}
 \end{aligned}$$

Lower bounding $S^{-1}S^{-*} \lesssim \sigma_{\min}(S)^{-1}$ and $S^{-1}S^{-*} \gtrsim 1/\sigma_{\max}(S)^2$, we can upper bound the above by

$$\begin{aligned} \log \det(\Gamma_T \Gamma_k^{-1}) &\leq \log \det\left\{\text{cond}(S)^2 \sum_{s=0}^{T-1} J^s J^{s*} \left(\sum_{s=0}^{k-1} (J^s) J^{s*}\right)^{-1}\right\} \\ &= 2d \log \text{cond}(S) + \log \det\left\{\sum_{s=0}^{T-1} (J^s) J^{s*} \left(\sum_{s=0}^{k-1} J^s J^{s*}\right)^{-1}\right\} \end{aligned}$$

To continue the bound, write the Jordan matrices $J = \text{block}(J_1, \dots, J_L)$ as block diagonal matrices. Then

$$\begin{aligned} \log \det(\Gamma_T \Gamma_k^{-1}) &\leq \log \det\left\{\text{cond}(S)^2 \sum_{s=0}^{T-1} (J^s) J^{s*} \left(\sum_{s=0}^{k-1} J^s J^{s*}\right)^{-1}\right\} \\ &= 2d \log \text{cond}(S) \sum_{b=1}^B \log \det\left\{\sum_{s=0}^{T-1} J_\ell^s J_\ell^{s*} \left(\sum_{s=0}^{k-1} J_\ell^s J_\ell^{s*}\right)^{-1}\right\} \end{aligned}$$

If $J_\ell = a_\ell$ is a Jordan matrix with block size equal to 1, then

$$\begin{aligned} \log \det\left\{\sum_{s=0}^{T-1} (J_\ell^s) J_\ell^{s*} \left(\sum_{s=0}^{k-1} J_\ell^s J_\ell^{s*}\right)^{-1}\right\} &= \log \frac{\sum_{s=0}^{T-1} |a_\ell|^{2s}}{\sum_{s=0}^{k-1} |a_\ell|^{2s}} \\ &\leq \log \frac{\lceil T/k \rceil \sum_{s=0}^{k-1} |a_\ell|^{2s}}{\sum_{s=0}^{k-1} a_\ell^{2s}} = \log(\lceil T/k \rceil) \end{aligned}$$

where the inequality uses the fact that a_ℓ^{2s} is decreasing. If $\dim(J_\ell) > 1$, we shall use the following lemma:

Lemma A.5 *Let $A \succeq 0$ be a $d \times d$ Complex Hermitian matrix. Then $A \preceq d \text{Diag}(A)$, where $\text{Diag}(A)$ is the diagonal matrices whose diagonal entries are those of A .*

Proof We can write

$$\begin{aligned} d \text{Diag}(A) - A &= (d-1) \text{Diag}(A) + \sum_{1 \leq i \neq j \leq d} A_{ij} e_i e_j^\top + \overline{A_{ij}} e_j e_i^\top \\ &= \sum_{1 \leq i \neq j \leq d} A_{ii} e_i e_i^* + A_{ij} (e_i e_j^\top + e_j e_i^\top) + A_{jj} e_j e_j^\top \succeq 0 \end{aligned}$$

■

We can then bound

$$\begin{aligned}
 \log \det \left\{ \sum_{s=0}^{T-1} J_\ell^s J_\ell^{s*} \left(\sum_{s=0}^{T-1} (J_\ell^s J_\ell^{s*})^{-1} \right) \right\} &\stackrel{(i)}{\leq} \log \det \left\{ \sum_{s=0}^{T-1} J_\ell^s J_\ell^{s*} \right\} \\
 &\stackrel{(ii)}{\leq} \log \det \left\{ \dim(J_\ell) \text{Diag} \left(\sum_{s=0}^{T-1} (J_\ell^s J_\ell^{s*}) \right) \right\} \\
 &= \dim(J_\ell) \log \dim(J_\ell) + \sum_{i=1}^{\dim(J_\ell)} \log \left(\sum_{s=0}^{T-1} (J_\ell^s J_\ell^{s*})_{ii} \right) \\
 &= \dim(J_\ell) \log \dim(J_\ell) + \sum_{i=1}^{\dim(J_\ell)} \log \left(\sum_{s=0}^{T-1} \sum_j (J_\ell^s)_{ij}^2 \right)
 \end{aligned}$$

where (i) uses that $\sum_{s=0}^{T-1} J_\ell^s J_\ell^{s*} \succeq I$, and (ii) uses Lemma A.5. We can then compute that if J_ℓ has diagonals a_ℓ ,

$$(J_\ell^s)_{i,j} = \begin{cases} \binom{s}{j-i} a_\ell^{s-(j-i) \vee 0} & i \leq j \\ 0 & \text{otherwise} \end{cases}$$

So that

$$\begin{aligned}
 \sum_j (J_\ell^s)_{ij}^2 &= \sum_{j \geq 1} \binom{s}{j-i}^2 |a_\ell|^{2(T-(j-i) \vee 0)} \\
 &= \dim(J_\ell)^2 s^{2(\dim(J_\ell)-i)}
 \end{aligned}$$

Hence,

$$\begin{aligned}
 \sum_{i=1}^{\dim(J_\ell)} \log \left(\sum_{s=0}^{T-1} \sum_j (J_\ell^s)_{ij}^2 \right) &\leq \sum_{i=1}^{\dim(J_\ell)} \log \left(\dim(J_\ell)^2 \sum_{s=0}^{T-1} s^{2(\dim(J_\ell)-i)} \right) \\
 &\leq \sum_{i=1}^{\dim(J_\ell)} \log \left(\dim(J_\ell)^2 T^{2(\dim(J_\ell)-i)+1} \right) \\
 &= 2 \dim(J_\ell) \log(\dim(J_\ell)) + \sum_{i=1}^{\dim(J_\ell)} 2(\dim(J_\ell) - i) + 1 \log(T) \\
 &= 2 \dim(J_\ell) \log(\dim(J_\ell)) + \log(T) \cdot \sum_{i=1}^{\dim(J_\ell)} i \\
 &\leq 2 \dim(J_\ell) \log(\dim(J_\ell)) + \dim(J_\ell)^2 \log T \\
 &\leq 4 \dim(J_\ell)^2 \log T
 \end{aligned}$$

where the last line uses that $T \geq d \geq \dim(J_\ell)$, and that $\dim(J_\ell) \geq 2$.

Appendix B. Specialized Analysis in Scalar Linear Systems

In this appendix, we present specialized upper and lower bounds in the case of scalar systems. Specifically, we consider $x_{t+1} = a_*x_t + \eta_t$, where $\eta_t \sim \mathcal{N}(0, \sigma^2)$, and $x_0 = 0$. Our upper bound has sharp, explicit constants, and captures the correct qualitative behavior for unstable scalar systems:

Theorem B.1 *Let $\epsilon \in (0, 1)$ and $\delta \in (0, 1/2)$. Then $\mathbb{P}[|\hat{a}(T) - a_*| \leq \epsilon] \geq 1 - \delta$ as long as*

$$T \geq \begin{cases} \frac{8}{\epsilon} \log\left(\frac{2}{\delta}\right) + \frac{4}{\epsilon^2} (1 - (|a_*| - \epsilon)^2) \log\left(\frac{2}{\delta}\right) & a_* \leq 1 + \epsilon \\ \max\left\{ \frac{8}{(|a_*| - \epsilon)^2 - 1} \log\left(\frac{2}{\delta}\right), \frac{4 \log(\frac{1}{\epsilon})}{\log(|a_*| - \epsilon)} + 8 \log\left(\frac{2}{\delta}\right) \right\} & a_* > 1 + \epsilon. \end{cases}$$

We match the upper bound with a lower bound which shows that our rates are optimal. Unlike the d -dimension case, our lower bound considers ‘local alternatives’ rather than scaled orthogonal matrices¹:

Theorem B.2 (1-D Lower Bound) *Fix an $a_* \in \mathbb{R}$, and define $\Gamma_T := \sum_{t=1}^T a_*^{2t}$. Fix an alternative $a' \in \{a_* - 2\epsilon, a_* + 2\epsilon\}$, and $\delta \in (0, 1/4)$. Then for any estimator \hat{a} ,*

$$\sup_{a \in \{a_*, a'\}} \mathbb{P}_a [|\hat{a}(T) - a_*| \geq \epsilon] \geq \delta \text{ for any } T \text{ such that } T\Gamma_T \leq \frac{\log(1/2\delta)}{8\epsilon^2}.$$

Theorem B.1 is proven in Section C below, and Theorem B.2 is proven in Section F.1.

Appendix C. Proof of Theorem B.1

To prove Theorem B.1, we write the error $E = \hat{a} - a = \frac{\sum_{t=0}^{T-1} x_t \eta_t}{\sum_{t=0}^{T-1} x_t^2}$. Since we are interested in upper bounding the probability that $|E| > \epsilon$ it suffices to show that the following two probabilities are small:

$$\mathbb{P}\left(\epsilon \sum_{t=0}^{T-1} x_t^2 - \sum_{t=0}^{T-1} x_t \eta_t < 0\right) \quad \text{and} \quad \mathbb{P}\left(\epsilon \sum_{t=0}^{T-1} x_t^2 + \sum_{t=0}^{T-1} x_t \eta_t < 0\right).$$

These probabilities are upper bounded by a standard Chernoff bound

$$\mathbb{P}\left(\epsilon \sum_{t=0}^{T-1} x_t^2 \pm \sum_{t=0}^{T-1} x_t \eta_t < 0\right) \leq \inf_{\lambda \leq 0} \mathbb{E} \exp\left(\lambda \epsilon \sum_{t=0}^{T-1} x_t^2 \pm \lambda \sum_{t=0}^{T-1} x_t \eta_t\right). \quad (3.25)$$

We will apply this equation with $\lambda = -\epsilon$, controlling its magnitude with following lemma, proved in Section C.1 below:

Lemma 3 *Let a, ν, μ , and x be real numbers with $\nu < 1$ and let $\eta \sim \mathcal{N}(0, 1)$. Then*

$$\mathbb{E}_\eta \exp\left(\frac{\nu}{2}(ax + \eta)^2 + \mu x \eta\right) = \frac{\exp\left(x^2 \frac{\nu a^2 + 2\nu a \mu + \mu^2}{2(1-\nu)}\right)}{\sqrt{1-\nu}}.$$

1. In one dimension, the orthogonal matrices are just the set $\{-1, 1\}$, and this precludes packing ‘nearby’ orthogonal matrices as in the d -dimensional case

With this lemma in hand, we can construct a recursive sequence which upper bounds $|a - \hat{a}|$ with high probability:

Proposition C.1 *Let a be a real number and for $\alpha \in \mathbb{R}_+$ and $\epsilon \in (0, 1)$ define recursively the sequence ρ_t by $\rho_{T-1} = 1$ and*

$$\rho_t = \begin{cases} 1 + r\rho_{t+1} & \rho_{t+1} \leq \alpha/\epsilon^2, \\ \alpha/\epsilon^2 & \rho_{t+1} > \alpha/\epsilon^2. \end{cases} \quad \text{where } r = \frac{(|a| - \epsilon)^2}{1 + \alpha}.$$

With this notation, $\mathbb{P}(|\hat{a} - a| \leq \epsilon) \leq 2 \exp\left(-\frac{\epsilon^2}{2(1+\alpha)} \sum_{t=1}^{T-1} \rho_t\right)$.

Proof The proof of this result is similar to the proof of the Azuma-Hoeffding inequality. It requires upper-bounding the MGF introduced in (3.25) by inductively applying the tower property of conditional expectation. We detail the proof in Section C.3. \blacksquare

The proof of Theorem B.1 is concluded in Section C.2, where we upper bound the sum $\sum_{t=1}^{T-1} \rho_t$, and solve for T .

C.1. Proof of Lemma 3

$$\begin{aligned} \mathbb{E}_\eta \exp\left(\frac{\nu}{2}(ax + \eta)^2 + \mu x \eta\right) &= e^{\frac{\nu}{2}a^2x^2} \mathbb{E}_\eta e^{\frac{\nu}{2}\eta^2 + \eta x(\nu a + \mu)} \\ &= \frac{e^{\frac{\nu}{2}a^2x^2}}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{\frac{\nu-1}{2}\eta^2 + \eta x(\nu a + \mu)} d\eta = e^{\frac{\nu}{2}a^2x^2} \frac{e^{x^2 \frac{(\nu a + \mu)^2}{2(1-\nu)}}}{\sqrt{1-\nu}} = \frac{\exp\left(x^2 \frac{\nu a^2 + 2\nu a \mu + \mu^2}{2(1-\nu)}\right)}{\sqrt{1-\nu}}. \end{aligned}$$

C.2. Proof of Theorem B.1

Once again we let $a \geq 0$ for simplicity and recall from Proposition C.1 that we denote $r = (a - \epsilon)^2 / (1 + \alpha)$. We study the case $a \leq 1$ first. Let us consider the sequence ρ_t introduced in Proposition C.1 with $\alpha = 2\epsilon$ and note that

$$1 + r + \dots + r^t \leq 1 + (1 + 2\epsilon)^{-1} + \dots + (1 + 2\epsilon)^{-t} \leq \frac{1}{1 - (1 + 2\epsilon)^{-1}} \leq \frac{2}{\epsilon},$$

which shows that for all t we have $\rho_{T-1-t} = 1 + r + \dots + r^t$ and hence $\sum_{t=1}^{T-1} \rho_t = \sum_{t=1}^{T-1} \frac{1-r^t}{1-r} = \frac{T}{1-r} - \frac{\sum_{t=0}^{T-1} r^t}{1-r}$. Since $T/2 \geq 1 + r + r^2 + \dots + r^{T-1}$ when $T \geq 6/\epsilon$, we obtain that $\sum_{t=1}^{T-1} \rho_t \geq \frac{T}{2(1-r)}$, which, together with Proposition C.1, it implies that

$$\mathbb{P}(|\hat{a} - a| \leq \epsilon) \leq 2 \exp\left(-\frac{\epsilon^2 T}{4(1+2\epsilon)(1-r)}\right) = 2 \exp\left(-\frac{\epsilon^2 T}{4(1+2\epsilon - (a-\epsilon)^2)}\right).$$

The first part of the corollary follows immediately.

We turn to the case $|a| > 1 + \epsilon$. Once again we assume $a > 0$ for simplicity. Recall that we have the freedom to choose any $\alpha \in \mathbb{R}_+$ for defining the sequence ρ_t . Since $a > 1 + \epsilon$, if we choose $\alpha < (a - \epsilon)^2 - 1$ we guarantee that $r > 1$. To satisfy this inequality we choose

$\alpha = ((a - \epsilon)^2 - 1)/2$. Then, with this choice of α , the sequence ρ_t grows exponentially to α/ϵ^2 . More precisely, by construction, since

$$\left[\frac{(a - \epsilon)^2}{1 + \alpha} \right]^{T-2} = \left[\frac{2(a - \epsilon)^2}{1 + (a - \epsilon)^2} \right]^{T-2} \geq (a - \epsilon)^{T-2},$$

ρ_1 is guaranteed to be equal to α/ϵ^2 as long as $(a - \epsilon)^{T-2} \geq \alpha/\epsilon^2$. This last inequality holds when $T \geq \frac{\log\left(\frac{(a-\epsilon)^2-1}{2\epsilon^2}\right)}{\log(a-\epsilon)} + 2$. In particular, if we choose T to be at least double the right-hand side of the previous expression, then at least half of the terms ρ_t are equal to α/ϵ^2 , implying

$$\mathbb{P}(|\hat{a} - a| \leq \epsilon) \leq 2 \exp\left(-\frac{\alpha T}{4(1 + \alpha)}\right).$$

The conclusion now follows easily.

C.3. Proof of Proposition C.1

We restrict ourselves to the case $a \geq 0$ (the case $a < 0$ can be analyzed analogously), and hence $r = (a - \epsilon)^2/(1 + \alpha)$. We upper bound the MGF (3.25) when $\lambda = -\epsilon$. Note that

$$\begin{aligned} \mathbb{E} \exp\left(-\epsilon^2 \sum_{t=0}^{T-1} x_t^2 \pm \epsilon \sum_{t=0}^{T-1} x_t \eta_t\right) &= \mathbb{E} \left[e^{-\epsilon^2 \sum_{t=0}^{T-1} x_t^2 \pm \epsilon \sum_{t=0}^{T-2} x_t \eta_t} \mathbb{E}_{\eta_{T-1}} \left[e^{\pm \epsilon x_{T-1} \eta_{T-1}} \mid \mathcal{F}_{T-1} \right] \right] \\ &= \mathbb{E} \left[e^{-\epsilon^2 \sum_{t=0}^{T-2} x_t^2 \pm \epsilon \sum_{t=0}^{T-3} x_t \eta_t} \mathbb{E} \left[e^{-\frac{\epsilon^2}{2} x_{T-1}^2 \pm \epsilon x_{T-2} \eta_{T-2}} \mid \mathcal{F}_{T-2} \right] \right]. \end{aligned}$$

Then, from Lemma 3 we can upper bound the MGF by induction on k by

$$\mathbb{E} \left[e^{-\epsilon^2 \sum_{t=0}^{T-k-1} x_t^2 - \epsilon \sum_{t=0}^{T-k-2} x_t \eta_t} \mathbb{E} \left[e^{-\frac{\epsilon^2 \beta_{T-k}}{2} x_{T-k}^2 - \epsilon x_{T-k-1} \eta_{T-k-1}} \mid \mathcal{F}_{T-k-1} \right] \right] \prod_{j=T-k+1}^{T-1} (1 + \epsilon^2 \beta_j)^{-1/2},$$

where β_t is any positive sequence such that $\beta_{T-1} = 1$ and for $1 \leq t < T - 1$ it satisfies $\beta_t \leq 1 + \frac{\beta_{t+1}(a-\epsilon)^2}{1 + \epsilon^2 \beta_{t+1}}$. It is straightforward to check that the sequence ρ_t defined in the proposition statement above satisfies this recursive inequality for any $\alpha \in (0, 1)$. Therefore, we obtain the upper bound

$$\begin{aligned} \mathbb{E} \exp\left(-\epsilon^2 \sum_{t=0}^{T-1} x_t^2 - \epsilon \sum_{t=0}^{T-1} x_t \eta_t\right) &\leq \prod_{t=1}^{T-1} (1 + \epsilon^2 \rho_t)^{-1/2} = \exp\left(\sum_{t=1}^{T-1} -\frac{1}{2} \log(1 + \epsilon^2 \rho_t)\right) \\ &\leq \exp\left(\sum_{t=1}^{T-1} -\frac{\epsilon^2 \rho_t}{2(1 + \epsilon^2 \rho_t)}\right) \leq \exp\left(-\frac{\epsilon^2}{2(1 + \alpha)} \sum_{t=1}^{T-1} \rho_t\right). \end{aligned}$$

Appendix D. Proof of Theorem 2.4

In this section, we conclude the technical aspects of the proof of Theorem 2.4. Recall the definition of the events

$$\mathcal{E}_1 := \left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} \geq K \right\}, \quad \mathcal{E}_2 := \left\{ \mathbf{X}^\top \mathbf{X} \succeq \frac{k \lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}}{16} \right\}, \quad \mathcal{E}_3 := \left\{ \mathbf{X}^\top \mathbf{X} \not\preceq \bar{\Gamma} \right\}.$$

As we recall from Section 4, if we can show that $\mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c]$ and $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c]$ where bounded above by δ , we have

$$\mathbb{P} \left[\|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{4K}{p\sqrt{k}\lfloor T/k \rfloor \lambda_{\min}(\Gamma_{\text{sb}})} \right] \leq 3\delta \quad (4.26)$$

Observe that our condition on k implies that necessarily $k \leq T/10$, so that $k\lfloor T/k \rfloor \geq T - k \geq 9/10T$. Hence, we will have established

$$\mathbb{P} \left[\|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}} \geq \frac{10}{9} \cdot \frac{4K}{p\sqrt{T}\lambda_{\min}(\Gamma_{\text{sb}})} \right] \leq 3\delta, \quad (4.27)$$

and substiting in

$$K = 20\sigma \sqrt{n + d \log \frac{10}{p} + \log \det(\bar{\Gamma}\Gamma_{\text{sb}}^{-1}) + \log(1/\delta)}. \quad (4.28)$$

proves the theorem.

D.1. Bounding $\mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c]$

Substituting the definitions of \mathcal{E}_2 and \mathcal{E}_3 ,

$$\mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c] = \mathbb{P} \left[\left\{ \mathbf{X}^\top \mathbf{X} \not\leq \frac{k\lfloor T/k \rfloor p^2 \nu^2}{16} \right\} \cap \left\{ \mathbf{X}^\top \mathbf{X} \preceq \Gamma_{\max} \right\} \right].$$

Proposition 2.5 and Equation (2.7) imply

$$\forall w \in \mathbb{R}^d, \mathbb{P} \left[\|\mathbf{X}w\|^2 \leq \frac{k\lfloor T/k \rfloor p^2 w^\top \Gamma_{\text{sb}} w}{8} \right] \leq \exp \left(-\frac{\lfloor T/k \rfloor p^2}{8} \right) \quad (4.29)$$

We apply Lemma 4.1 with $Q = \mathbf{X}$, with $\Gamma_{\max} \leftarrow T\bar{\Gamma}$, $\Gamma_{\min} \leftarrow k\lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}/8$, and \mathcal{T} a net $1/4$ -net of $\mathcal{S}_{\Gamma_{\min}}$ in the norm $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$. We shall use the following estimate of $|\mathcal{T}|$:

Lemma D.1 *Let $0 \prec \Gamma_{\min} \preceq \Gamma_{\max}$, and let \mathcal{T} be a minimal $\epsilon \leq 1/2$ -net of $\mathcal{S}_{\Gamma_{\min}}$ in the norm $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$. Then, $\log |\mathcal{T}| \leq d \log(1 + \frac{2}{\epsilon}) + \log \det(\Gamma_{\max} \Gamma_{\min}^{-1})$.*

Proof The covering number of $\mathcal{S}_{\Gamma_{\min}}$ in the norm $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$ is the same as the covering number of the shell of the ellipsoid $E := \{w : w^\top \Gamma_{\min}^{-1/2} \Gamma_{\max} \Gamma_{\min}^{-1/2} w\}$ in the norm $\|(\cdot)\|_2$. Consider a maximal ϵ -separated set of $\text{bd}(E)$. Letting \mathcal{B} denote the unit ball in \mathbb{R}^d , a standard volumetric argument shows

$$\begin{aligned} |\mathcal{T}| &\leq \frac{\text{vol}(\frac{2}{\epsilon}\mathcal{B} + E)}{\text{vol}(\frac{2}{\epsilon})} \\ &\stackrel{(i)}{\leq} \frac{\text{vol}((\frac{2}{\epsilon} + 1)E)}{\text{vol}(\mathcal{B})} = \left(\frac{2}{\epsilon} + 1\right)^d \frac{\text{vol}(E)}{\text{vol}(\mathcal{B})} \\ &= \left(\frac{2}{\epsilon} + 1\right)^d \det(\Gamma_{\min}^{-1/2} \Gamma_{\max} \Gamma_{\min}^{-1/2}) \\ &= \left(\frac{2}{\epsilon} + 1\right)^d \det(\Gamma_{\max} \Gamma_{\min}^{-1}) \end{aligned}$$

where (i) uses the inclusion $\mathcal{E} \subset \mathcal{E}$. ■

For $\epsilon = 1/4$ and our choice of Γ_{\min} , we have

$$\log |\mathcal{T}| = d \log\left(1 + \frac{2}{\epsilon}\right) + \log \det(\Gamma_{\max} \Gamma_{\min}^{-1}) \quad (4.30)$$

$$\begin{aligned} &= d \log(9) + d \log(8T/(\lfloor T/k \rfloor k)p^2) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1}) \\ &\leq d \log(9) + d \log(72/10p^2) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1}) \quad \text{since } k \leq T/2 \\ &\leq 2d \log(10/p) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1}). \end{aligned} \quad (4.31)$$

Hence, we conclude that

$$\begin{aligned} \mathbb{P}[\mathcal{E}_2^c \cap \mathcal{E}_3^c] &= \mathbb{P}\left[\left\{\mathbf{X}^\top \mathbf{X} \not\leq \frac{k \lfloor T/k \rfloor p^2 \nu^2}{16}\right\} \cap \{\mathbf{X}^\top \mathbf{X} \preceq \Gamma_{\max}\}\right] \\ &\leq \mathbb{P}\left[\left\{\exists w \in \mathcal{T} : \|\mathbf{X}w\|^2 < \frac{k \lfloor T/k \rfloor T p^2 w^\top \Gamma_{\text{sb}} w}{8}\right\} \cap \{\mathbf{X}^\top \mathbf{X} \preceq \Gamma_{\max}\}\right] \\ &\leq \exp\left(-\frac{\lfloor T/k \rfloor p^2}{8} + 2d \log(10/p) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1})\right) \\ &\leq \exp\left(-\frac{T p^2}{10k} + 2d \log(10/p) + \log \det(\bar{\Gamma} \Gamma_{\text{sb}}^{-1})\right) \end{aligned}$$

D.2. Bounding $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c]$

We shall need the following discretization lemma, proved in Appendix E.3.2:

Lemma D.2 *Let $Q \in \mathbb{R}^{n \times m}$ have full column rank, $q \in \mathbb{R}^n$, let $0 < \Gamma_{\min} \preceq Q^\top Q \preceq \Gamma_{\max}$, and let \mathcal{T} be a $1/4$ -net of $\mathcal{S}_{\Gamma_{\min}}$ in the norm $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$. Then,*

$$\sup_{w \in \mathcal{S}^{m-1}} \frac{\langle Qw, q \rangle}{\|Qw\|} \leq 2 \max_{w \in \mathcal{T}} \frac{\langle Qw, q \rangle}{\|Qw\|} \quad (4.32)$$

To control the size of $\|\mathbf{U}^\top \mathbf{E}\|_{\text{op}}$ we use a variational formulation of the operator norm and two coverings. Let $\mathbb{R}_*^d := \mathbb{R}^d - \{0\}$. Note that if \mathcal{T}_1 is a $1/2$ -net of \mathcal{S}^{n-1} (over the v), then,

$$\begin{aligned} \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} &\leq \sup_{v \in \mathcal{S}^{n-1}, w \in \mathbb{R}_*^d} \frac{w^\top \mathbf{U}^\top \mathbf{E}v}{\|w\|} \\ &\leq 2 \max_{v \in \mathcal{T}_1} \left(\sup_{w \in \mathbb{R}_*^d} \frac{w^\top \mathbf{U}^\top \mathbf{E}v}{\|w\|} \right) \quad \text{since } \mathcal{T}_1 \text{ is a } 1/2\text{-net} \\ &= 2 \max_{v \in \mathcal{T}_1} \left(\sup_{w \in \mathbb{R}_*^d} \frac{w^\top \mathbf{V} \Sigma \mathbf{U}^\top \mathbf{E}v}{\|\Sigma \mathbf{V}^\top w\|} \right) \quad \text{since } \Sigma \mathbf{V}^\top \text{ is full rank on } \mathcal{E}_2 \\ &= 2 \max_{v \in \mathcal{T}_1} \left(\sup_{w \in \mathbb{R}_*^d} \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} \right) = 2 \max_{v \in \mathcal{T}_1} \left(\sup_{w \in \mathcal{S}^{d-1}} \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} \right), \end{aligned}$$

where the second-to-last equation uses $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$ and $\|\Sigma\mathbf{V}^\top w\| = \|\mathbf{U}\Sigma\mathbf{V}^\top w\| = \|\mathbf{X}w\|$. Define $\mathcal{E}_g := \mathcal{E}_2 \cap \mathcal{E}_3^c$. We see then that on \mathcal{E}_g , we have

$$\frac{k\lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}^2}{16} \preceq \mathbf{X}^\top \mathbf{X} \preceq T\Gamma_{\text{max}}.$$

We now apply Lemma D.2 with $Q = \mathbf{X}$, with

$$\Gamma_{\text{max}} \leftarrow T\bar{\Gamma} \quad \text{and} \quad \Gamma_{\text{min}} \leftarrow k\lfloor T/k \rfloor p^2 \Gamma_{\text{sb}}/16 \quad (4.33)$$

and \mathcal{T}_2 a net $1/4$ -net of $\mathcal{S}_{\Gamma_{\text{min}}}$ in the norm $\|\Gamma_{\text{max}}^{1/2}(\cdot)\|_2$. This yields

$$\begin{aligned} & \mathbb{P} \left[\left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} > K \right\} \cap \mathcal{E}_g \right] \\ & \leq \mathbb{P} \left[\left\{ \max_{v \in \mathcal{T}_1} \sup_{w \in \mathcal{S}^{d-1}} \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/2 \right\} \cap \mathcal{E}_g \right] \\ & \leq \mathbb{P} \left[\left\{ \max_{v \in \mathcal{T}_1} \max_{w \in \mathcal{T}_2} \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/4 \right\} \cap \mathcal{E}_g \right] \\ & \leq |\mathcal{T}_1| |\mathcal{T}_2| \sup_{v \in \mathcal{T}_1, w \in \mathcal{T}_2} \mathbb{P} \left[\left\{ \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/4 \right\} \cap \mathcal{E}_g \right]. \end{aligned}$$

To obtain a pointwise bound on $\mathbb{P} \left[\left\{ \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/4 \right\} \cap \mathcal{E}_g \right]$ we use Lemma 4.2, with $Z_t = \langle X_t, w \rangle$, $W_t = \langle \eta_t, v \rangle$, and the bounds $\beta_- = w^\top \Gamma_{\text{min}} w$, and $\beta_+ = w^\top \Gamma_{\text{max}} w$, Γ_{min} , Γ_{max} are as in Equation (4.33). We can then bound

$$\begin{aligned} & |\mathcal{T}_1| |\mathcal{T}_2| \sup_{v \in \mathcal{S}^{d-1}, w \in \mathcal{S}^{d-1}} \mathbb{P} \left[\left\{ \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/4 \right\} \cap \mathcal{E}_g \right] \\ & \leq |\mathcal{T}_1| |\mathcal{T}_2| \sup_{v \in \mathcal{S}^{d-1}, w \in \mathcal{S}^{d-1}} \mathbb{P} \left[\left\{ \frac{w^\top \mathbf{X}^\top \mathbf{E}v}{\|\mathbf{X}w\|} > K/4 \right\} \cap \left\{ \|\mathbf{X}w\|^2 \in \left[\frac{k\lfloor T/k \rfloor \nu^2 p^2}{16}, \lambda_+ \right] \right\} \right] \\ & \leq |\mathcal{T}_1| |\mathcal{T}_2| \log_+ \left\lceil \frac{\beta_+}{\beta_-} \right\rceil \exp(-K^2/96\sigma^2) \\ & \stackrel{(i)}{\leq} \exp(n \log 5 + d \log 9 + \log \det(\frac{32}{p^2} \Gamma_{\text{max}} \Gamma_{\text{min}}^{-1})) \log_+ \left\lceil \frac{\beta_+}{\beta_-} \right\rceil \exp(-K^2/96\sigma^2). \end{aligned}$$

where (i) uses the standard metric entropy bound for the sphere(see, e.g. Vershynin (2011)), and an analogous computation to (4.31). Now since $w \neq 0$, bound $\log[x] \leq \log 1 + x \leq x$, and computing as in (4.31) yields the bound

$$\begin{aligned} \log \left\lceil \frac{\beta_+}{\beta_-} \right\rceil & \leq \frac{\beta_+}{\beta_-} \leq \sup_{w \in \mathbb{R}^d - \{0\}} \frac{w^\top \Gamma_{\text{max}} w}{w^\top \Gamma_{\text{min}} w} \\ & = \|\Gamma_{\text{min}}^{-1/2} \Gamma_{\text{max}} \Gamma_{\text{min}}^{-1/2}\| = \lambda_{\text{max}}(\Gamma_{\text{max}} \Gamma_{\text{min}}^{-1}) \\ & = \lambda_{\text{max}}\left(\frac{32}{p^2} \bar{\Gamma} \Gamma_{\text{sb}}^{-1}\right) \\ & \leq \exp(\log \lambda_{\text{max}}\left(\frac{32}{p^2} \bar{\Gamma} \Gamma_{\text{sb}}^{-1}\right)) \\ & \stackrel{(i)}{=} \exp(\log \det\left(\frac{32}{p^2} \bar{\Gamma} \Gamma_{\text{sb}}^{-1}\right)). \end{aligned}$$

where (i) uses the fact that $\frac{32}{p^2}\bar{\Gamma}\Gamma_{\text{sb}}^{-1}$ has the same eigenvalues as $\frac{32}{p^2}\Gamma_{\text{sb}}^{-1/2}\bar{\Gamma}\Gamma_{\text{sb}}^{-1/2} \succeq \frac{32}{p^2} \succeq I$. All together, we have

$$\begin{aligned} & \mathbb{P} \left[\left\{ \|\mathbf{U}^\top \mathbf{E}\|_{\text{op}} > K \right\} \cap \mathcal{E}_g \right] \\ & \leq \exp(n \log 5 + d \log 9 + 2 \log \det(\frac{32}{p^2}\bar{\Gamma}\Gamma_{\text{sb}}^{-1})) \exp(-K^2/96\sigma^2) \\ & \leq \exp(n \log 5 + 2d \log \frac{96}{p^2} + 2 \log \det(\bar{\Gamma}\Gamma_{\text{sb}}^{-1})) \exp(-K^2/96\sigma^2) \\ & \leq \exp(4(n + d \log \frac{10}{p} + \log \det(\bar{\Gamma}\Gamma_{\text{sb}}^{-1}))) \exp(-(K/10\sigma)^2) \end{aligned}$$

Hence, we guarantee that $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3^c] \leq \delta$ if we choose

$$K = 20\sigma \sqrt{n \log + d \log \frac{10}{p} + \log \det(\bar{\Gamma}\Gamma_{\text{sb}}^{-1}) + \log(1/\delta)}.$$

Appendix E. Proof of Technical Results

E.1. Proof of Proposition 2.5

To exploit the (k, ν, p) block martingale small-ball condition we partition the sequence of random variables Z_1, Z_2, \dots, Z_T into $\lfloor T/k \rfloor$ blocks of size k (we discard the remainder terms). For simplicity we denote $S = \lfloor T/k \rfloor$. We consider the random variables

$$B_j = \mathbb{I} \left(\sum_{i=1}^k Z_{jk+i}^2 \geq \frac{\nu^2 pk}{2} \right) \quad \text{for } 0 \leq j \leq S-1.$$

Given this notation, we can use the Chernoff bound to obtain

$$\mathbb{P} \left[\sum_{i=1}^T Z_i^2 \leq \frac{\nu^2 p^2}{8} kS \right] \leq \mathbb{P} \left[\sum_{j=0}^{S-1} B_j \leq \frac{p}{4} S \right] \leq \inf_{\lambda \leq 0} e^{-\frac{p}{4} S} \mathbb{E} e^{\lambda \sum_{j=0}^{S-1} B_j}. \quad (5.34)$$

The first inequality above uses the trivial inequality $\sum_{i=1}^k Z_{jk+i}^2 \geq \frac{\nu^2 pk}{2} B_j$.

For upper bounding the MGF on the right hand side we will use the tower property with respect to the filtration \mathcal{F}_{jk} for j from $S-1$ to 0. Before turning to that computation it is valuable to lower bound the conditional expectations $\mathbb{E}[B_j | \mathcal{F}_{jk}]$:

$$\begin{aligned} \mathbb{E}[B_j | \mathcal{F}_{jk}] &= \mathbb{P} \left[\sum_{i=1}^k Z_{jk+i}^2 \geq \frac{\nu^2 pk}{2} \mid \mathcal{F}_{jk} \right] \stackrel{(a)}{\geq} \mathbb{P} \left[\frac{1}{k} \sum_{i=1}^k \mathbb{I}(|Z_{jk+i}| \geq \nu) \geq \frac{p}{2} \mid \mathcal{F}_{jk} \right] \\ &\stackrel{(b)}{\geq} \frac{(p/2)}{1 - (p/2)} \geq \frac{p}{2}, \end{aligned}$$

where (a) uses the trivial inequality $\frac{1}{\nu^2} Z_{jk+i}^2 \geq \mathbb{I}(|Z_{jk+i}| \geq \nu)$, and (b) uses inequality follows from the (k, ν, p) -BMSB condition and the following claim is straightforward.

Claim: Let Z be a random variable supported in $[0, 1]$ almost surely such that $\mathbb{E}[Z] \geq p$ for some $p \in (0, 1)$. Then, for all $t \in [0, p]$, $\mathbb{P}[Z \geq t] \geq \frac{p-t}{1-t}$.

From this lower bound on $\mathbb{E}[B_j | \mathcal{F}_{jk}]$, using $\lambda \leq 0$, we get

$$\begin{aligned} \mathbb{E}\left[e^{\lambda B_j} | \mathcal{F}_{jk}\right] &= e^\lambda \mathbb{P}[B_j = 1 | \mathcal{F}_{jk}] + \mathbb{P}[B_j = 0] = (e^\lambda - 1)\mathbb{E}[B_j | \mathcal{F}_{jk}] + 1 \\ &\leq (e^\lambda - 1)\frac{p}{2} + 1. \end{aligned}$$

Now, by inductively conditioning on \mathcal{F}_{jk} , we can upper bound

$$\begin{aligned} \mathbb{E}e^{\lambda \sum_{j=0}^{S-1} B_j} &= \mathbb{E}\left[e^{\lambda \sum_{j=0}^{S-2} B_j} \mathbb{E}\left[e^{\lambda B_{S-1}} | \mathcal{F}_{(S-1)k}\right]\right] \leq \left((e^\lambda - 1)\frac{p}{2} + 1\right) \mathbb{E}\left[e^{\lambda \sum_{j=0}^{S-2} B_j}\right] \\ &\leq \left((e^\lambda - 1)\frac{p}{2} + 1\right)^S. \end{aligned}$$

We now plug in this upper bound in Equation 5.34 and optimize for λ . From the first order optimality condition it is easy to see that the optimal choice of λ is

$$\lambda_\star = \log\left(\frac{1 - p/2}{2 - p/2}\right).$$

Plugging in λ_\star back in Equation 5.34, after some elementary calculus and algebraic manipulations, we find the desired conclusion.

E.2. Proof of Martingale Concentration (Lemma 4.2)

For ease of notation we denote $S_t = \sum_{s=1}^t Z_s W_s$ and $R_t = \sum_{s=1}^t Z_s^2$.

(a) Using a Chernoff argument, we have

$$\begin{aligned} \mathbb{P}[\{S_T \geq \alpha\} \cap \{R_T \leq \beta\}] &= \inf_{\lambda > 0} \mathbb{P}\left[\{e^{\lambda S_T} \geq e^{\lambda \alpha}\} \cap \{R_T \leq \beta\}\right] \\ &= \inf_{\lambda > 0} \mathbb{P}\left[e^{\lambda S_T} \mathbb{I}(R_T \leq \beta) \geq e^{\lambda \alpha}\right] \\ &\leq \inf_{\lambda > 0} e^{-\lambda \alpha} \mathbb{E}[e^{\lambda S_T} \mathbb{I}(R_T \leq \beta)] \\ &= \inf_{\lambda > 0} e^{-\lambda \alpha} \cdot e^{\lambda^2 \sigma^2 \beta / 2} \mathbb{E}[e^{\lambda S_T - \lambda^2 \sigma^2 \beta / 2} \mathbb{I}(R_T \leq \beta)] \\ &\leq \inf_{\lambda > 0} e^{-\lambda \alpha} \cdot e^{\lambda^2 \sigma^2 \beta / 2} \mathbb{E}[e^{\lambda S_T - \lambda^2 \sigma^2 R_T / 2}]. \end{aligned}$$

Now, we claim that $\mathbb{E}[e^{\lambda S_T - \lambda^2 \sigma^2 R_T / 2}] \leq 1$. Indeed, by the tower rule and the assumption that $W_t | \mathcal{F}_t$ is a zero mean σ -sub-Gaussian r.v., we have

$$\begin{aligned} \mathbb{E}[\exp(\lambda S_T - \lambda^2 \sigma^2 R_T / 2)] &= \mathbb{E}[\mathbb{E}[\exp(\lambda S_T - \lambda^2 \sigma^2 R_T / 2) | \mathcal{F}_T]] \\ &\leq \mathbb{E}[\exp(\lambda S_{T-1} - \lambda^2 \sigma^2 R_{T-1} / 2) \mathbb{E}[e^{\lambda Z_T W_T - \lambda^2 \sigma^2 Z_T^2 / 2} | \mathcal{F}_T]] \\ &\leq \mathbb{E}[\exp(\lambda S_{T-1} - \lambda^2 \sigma^2 R_{T-1} / 2)] \\ &\quad \vdots \\ &\leq \mathbb{E}[\exp(\lambda S_1 - \lambda^2 \sigma^2 R_1 / 2)] \leq 1. \end{aligned} \tag{5.35}$$

Hence,

$$\mathbb{P}[\{S_t \geq \alpha\} \cap \{R_T \leq \beta\}] \leq \inf_{\lambda > 0} e^{-\lambda \alpha} e^{\lambda^2 \sigma^2 \beta / 2} = e^{-\alpha^2 / 2\sigma^2 \beta}.$$

(b) Let $B := \log \lceil \frac{\beta_+}{\beta_-} \rceil$. Then

$$\begin{aligned} \mathbb{P}[\{S_T > \alpha \sqrt{R_T}\} \cap \{\beta_- \leq R_T \leq \beta_+\}] &\leq \mathbb{P}[\{S_T > \alpha \sqrt{R_T}\} \cap \{\beta_- \leq R_T \leq e^B \beta_-\}] \\ &= \sum_{i=0}^{B-1} \mathbb{P}[\{S_T > \alpha \sqrt{R_T}\} \cap \{e^i \beta_- \leq R_T \leq e^{i+1} \beta_-\}] \\ &\leq \sum_{i=0}^{B-1} \mathbb{P}[\{S_T > \alpha \sqrt{e^i \beta_-}\} \cap \{e^i \beta_- \leq R_T \leq e^{i+1} \beta_-\}] \\ &\leq \sum_{i=0}^{B-1} \mathbb{P}[\{S_T > \alpha \sqrt{e^i \beta_-}\} \cap \{R_T \leq e^{i+1} \beta_-\}] \\ &\stackrel{(i)}{\leq} \sum_{i=0}^{B-1} \exp\left(\frac{-\alpha^2 e^i \beta_-}{2e^{i+1} \sigma^2 \beta_-}\right) \\ &= B \exp\left(\frac{-\alpha^2}{2e\sigma^2}\right) \leq \log \left\lceil \frac{\beta_+}{\beta_-} \right\rceil \exp\left(\frac{-\alpha^2}{6\sigma^2}\right). \end{aligned}$$

Above, (i) follows from part (a) of the claim.

E.3. Proofs of Covering Results

E.3.1. PROOF OF LEMMA 4.1

Consider the transformed matrix $\Gamma_{\min}^{-1/2} Q \Gamma_{\min}^{-1/2}$. It suffices to show that under the assumptions of Lemma 4.1,

$$\inf_{w \in \mathcal{S}_{\Gamma_{\min}}} \|Qw\| \geq 3/4, \quad (5.36)$$

since then $Q^T Q \succeq (3/4)^2 \Gamma_{\min} \succeq \Gamma_{\min}/2$. Now $w, v \in \mathbb{R}^d$, we can bound

$$\|Q \Gamma_{\max}^{-1/2} (w - v)\| \leq \|\Gamma_{\max}^{1/2} (w - v)\|$$

since $Q^T Q \preceq \Gamma_{\max}$. In particular, if \mathcal{T} is a $1/2$ -net of $\mathcal{S}_{\Gamma_{\min}}$ in $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$, then

$$\inf_{w \in \mathcal{S}_{\Gamma_{\min}}} \|Qw\| \geq \inf_{w \in \mathcal{T}} \|Qw\| - \frac{1}{4} \geq \frac{1}{4}.$$

where the last step follows from the assumption that $\inf_{w \in \mathcal{T}} \|Qw\| \geq 1$.

E.3.2. PROOF OF LEMMA D.2

Define the map $\phi(w) := \frac{Qw}{\|Qw\|}$. We shall prove that for all $v, w \in \mathcal{S}^{m-1}$, one has that

$$\|\phi(w) - \phi(v)\| \leq \frac{2\|Q(v-w)\|}{\|Qv\|} \quad (5.37)$$

Note observe that, if $0 \prec \Gamma_{\min} \preceq Q^\top Q \preceq \Gamma_{\max}$. Hence, each $w \in \mathcal{S}_{\Gamma_{\min}}$ can be written as $\Gamma_{\min}^{-1/2} w'$ for $w' \in \mathcal{S}^{d-1}$, we have that

$$\inf_{w \in \mathcal{S}_{\Gamma_{\min}}} \|Qw\|^2 = \inf_{w' \in \mathcal{S}^{d-1}} (w')^\top \Gamma_{\min}^{-1/2} Q^\top Q \Gamma_{\min}^{-1/2} w' \geq \|w'\|_2^2 = 1$$

and that

$$\|Qw\| = w^\top Q^\top Q w \leq w^\top \Gamma_{\max} w = \|\Gamma_{\max}^{1/2} w\|^2 \quad (5.38)$$

Thus, for all $v, w \in \mathcal{S}_{\Gamma_{\min}}^{m-1}$,

$$\|\phi(w) - \phi(v)\| \leq \frac{2\|Q(v-w)\|}{\|Qv\|} \leq 2\|\Gamma_{\max}^{1/2}(v-w)\|$$

Since $\Gamma_{\min}^{-1/2} \succ 0$ is full rank, we have

$$\sup_{w \in \mathcal{S}^{d-1}} \frac{\langle Qw, q \rangle}{\|Qw\|} = \sup_{w \in \mathcal{S}_{\Gamma_{\min}}^{d-1}} \frac{\langle Qw, q \rangle}{\|Q_{\min} w\|}$$

and since \mathcal{T} is a $1/4$ -net of \mathcal{S}^{d-1} in the norm $\|\Gamma_{\max}^{1/2}(\cdot)\|_2$, (5.37) implies the above is at most

$$2 \sup_{w \in \mathcal{T}} \frac{\langle Qw, q \rangle}{\|Q_{\min} w\|}.$$

It remains to check (5.37).

$$\begin{aligned} \|\phi(v) - \phi(w)\| &= \left| \frac{Qv}{\|Qv\|} - \frac{Qw}{\|Qw\|} \right| \\ &\leq +\|Qw\| \left| \frac{1}{\|Qv\|} - \frac{1}{\|Qw\|} \right| \\ &\leq \frac{\|Q(v-w)\|}{\|Qv\|} + \frac{|\|Qw\| - \|Qv\||}{\|Qv\|} \\ &\leq 2 \frac{\|Q(w-v)\|}{\|Qv\|}. \end{aligned}$$

Appendix F. Lower Bounds

F.1. Proof of Information Theoretic Lower Bounds, Theorems B.2 and 2.3

In this section we prove Theorem B.2 and 2.3. We shall use $\mathbb{P}_A^{(T)}$ to denote the law of the iterates $X_{t+1} = AX_t + \eta_t$, where $\eta_t \sim \mathcal{N}(0, I)$, for $t = 1, 2, \dots, T$. We shall prove Theorems B.2 and 2.3 using Birge's Inequality, a bound which is qualitatively similar to Fano's inequality, but yields sharp high-probability lower bounds in low-dimensional settings.

Lemma F.1 (Variant of Birge's Inequality) *Let $\mathcal{E}_0, \mathcal{E}_1, \dots, \mathcal{E}_N$ be disjoint events, $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_N$ be probability laws, and let $\min_i \mathbb{P}(\mathcal{E}_i^c) \leq 1/2$. Then, for any $\delta \in (0, 1/2)$,*

$$\sum_{i=1}^N \text{KL}(\mathbb{P}_i, \mathbb{P}_0) \geq (1 - 2\delta) \log(N/2\delta). \quad (6.39)$$

In particular, fix an $\epsilon > 0$ and $\delta \in (0, 1/2)$, and suppose that for a finite set $\mathcal{N} \subset \mathbb{R}^{n \times n}$, all $A_1 \neq A_2 \in \mathcal{N}$ satisfy $\|A_1 - A_2\|_{\text{op}} \geq 2\epsilon$. Then if $\inf_{\hat{A}} \sup_{A \in \mathcal{N}} \mathbb{P}_A[\|A - \hat{A}(T)\|_{\text{op}} \geq \epsilon] \leq \delta$, then T, δ and $|\mathcal{N}|$ satisfy the following inequality for any $A_0 \in \mathcal{N}$:

$$\sup_{A \in \mathcal{N} - \{A_0\}} \text{KL}(\mathbb{P}_A^{(T)}, \mathbb{P}_{A_0}^{(T)}) \geq (1 - 2\delta) \log\left(\frac{|\mathcal{N}| - 1}{2\delta}\right). \quad (6.40)$$

We prove Lemma F.1 from a standard statement of Birge's inequality from (Boucheron et al., 2013, Theorem 4.20), in Section F.5. Lemma F.1 relates the probability of error to the KL-divergences between laws $\mathbb{P}_A^{(T)}$ in 2ϵ -separated set \mathcal{N} . Thus, our first step will be to compute the term $\text{KL}(\mathbb{P}_A^{(T)}, \mathbb{P}_{A_0}^{(T)})$. This amounts to a straightforward computation, carried out in Section F.4.2.

Lemma F.2 *Let $O, O' \in \text{O}(d)$. Then, $\text{KL}(\mathbb{P}_\rho^{(T)}, \mathbb{P}_{A_0}^{(T)}) = \|\rho O - A_0\|_{\mathbb{F}}^2 \cdot \sum_{t=1}^T \gamma_t(\rho)$, where we recall $\gamma_t(\rho) = \sum_{s=0}^{t-1} |\rho|^{2s}$.*

We are now in a position to prove the lower bound in one-dimension:

F.2. 1-D Lower Bound: Proof of Theorem B.2

Proof Fix an $\rho \in \mathbb{R}$, and let $\rho' \in \{\rho - 2\epsilon, \rho + 2\epsilon\}$. Viewing ρ, ρ' as matrices in $\mathbb{R}^{1 \times 1}$, we have Lemma F.2, implies $\text{KL}(\mathbb{P}_\rho^{(T)}, \mathbb{P}_{\rho'}^{(T)}) = 4\epsilon^2 \cdot \sum_{t=1}^T \gamma_t(\rho)$. Then, applying Lemma F.1 with $A_0 = \rho$ and $\mathcal{N} = \{\rho, \rho'\}$, we have for if $\sup_{a \in \rho, \rho'} \mathbb{P}_a[|\hat{a}(T) - a| < \epsilon] \leq \delta$, then, $\epsilon^2 \cdot \sum_{t=1}^T \gamma_t(\rho) \geq (1 - 2\delta) \log\left(\frac{1}{2\delta}\right)$. Hence, we need $T \gamma_T(\rho) \geq \sum_{t=1}^T \gamma_t(\rho) \geq \frac{1}{4\epsilon} (1 - 2\delta) \log\left(\frac{1}{2\delta}\right)$. ■

F.3. d -Dimensional Lower Bound

If we chose \mathcal{N} to be a 2ϵ -packing of the set $\rho \text{O}(d)$, then Lemma F.1 and Lemma F.2 imply that, for any estimate \hat{A} such that

$$\sup_{O \in \text{O}(d)} \mathbb{P}_{\rho O}[\|\hat{A}(T) - \rho O\|_{\text{op}} \geq \epsilon] \geq \delta \text{ for any } T \text{ such that } (1 - 2\delta) \log\left(\frac{|\mathcal{N}|}{2\delta}\right) \geq \left(\sum_{t=1}^T \gamma_t(\rho)\right) \cdot \max_{\rho O, \rho O' \in \mathcal{N}} \|\rho O - \rho O'\|_{\mathbb{F}}^2.$$

In light of the above, our goal will be to construct a 2ϵ -packing \mathcal{N} such that $\inf_{\rho O, \rho O' \in \mathcal{N}} \|\rho O - \rho O'\|_{\mathbb{F}}^2$ is as small as possible. This is achieved by the following proposition, which lifts a $1/2$ -packing of the unit ball in $d - 1$ -dimensions to a packing \mathcal{N}_0 of $O(d)$, proved in Section F.3:

Proposition F.3 *Fix an $\epsilon_0 \leq 1/256$, and let \mathcal{T} be an $1/2$ -packing of $B_{d-1}(1)$. Then, there exists a set $\overline{\mathcal{N}} \subset O(d)$ with $|\mathcal{N}_0| = |\mathcal{T}|$ and, for all $A_1 \neq A_2 \in \overline{\mathcal{N}}$,*

$$\|A_1 - A_2\|_{\text{op}} \geq \epsilon_0/4 \quad \text{and} \quad \|A_1 - A_2\|_{\mathbb{F}} \leq 4\epsilon_0. \quad (6.41)$$

We now reparameterize the above proposition with $\epsilon_0 = \frac{8\epsilon}{\rho}$. Let \mathcal{T} be a maximal $1/2$ -packing of $\mathcal{B}_{d-1}(1)$; a standard fact shows that $|\mathcal{T}| \geq 2^{d-1}$. Then, as long as $\epsilon \leq \frac{\rho}{2048}$, $\mathcal{N} = \rho\overline{\mathcal{N}}$ is a 2ϵ -packing of the set $\rho O(d)$, and for all $\rho O, \rho O' \in \mathcal{N}$, $\|A_1 - A_2\|_{\mathbb{F}} \leq 32\epsilon$,

$$\begin{aligned} & \sup_{O \in O(d)} \mathbb{P}_{\rho O}[\|\widehat{A}(T) - \rho O\|_{\text{op}} \geq \epsilon] \geq \delta \text{ for any} \\ & T \text{ such that } (1 - 2\delta) \log \frac{2^d}{4\delta} \geq \left(\sum_{t=1}^T \gamma_t\right) \cdot (32\epsilon)^2. \end{aligned}$$

In particular, for $\delta \leq 1/4$, we see that there exists a universal constant c_0 such that $(1 - 2\delta) \log \frac{2^d}{4\delta} \geq c_0(d + \log(1/\delta))$, and hence for $c = c_0/32^2$, we see that

$$\sup_{\rho O \in O(d)} \mathbb{P}[\|\widehat{A}(T) - \rho O\|_{\text{op}} \geq \epsilon] \geq \delta \text{ for any } T : \frac{c_0(d + \log(1/\delta))}{\epsilon^2} \geq \sum_{t=1}^T \gamma_t(\rho).$$

Bounding $\sum_{t=1}^T \gamma_t(\rho) \leq T\gamma_T(\rho)$ concludes the proof.

F.4. Proof of Proposition F.3

We now construction of the packing $\overline{\mathcal{N}}$. If we define the set $\text{Skew}(d) := \{X \in \mathbb{R}^{d \times d} : X^\top = -X\}$, and recall the matrix exponential $\exp(X) = \sum_{j=0}^{\infty} X^j/j!$, a well-known theorem in Lie Theory ensures that $\exp(\text{Skew}(d)) \subset O(d)$ (see, e.g. [Knapp \(2016\)](#)). Moreover, \exp is an approximate isometry (in both $\|\cdot\|_{\text{op}}$ and $\|\cdot\|_{\mathbb{F}}$) from a small neighborhood of $0 \in \text{Skew}(d)$ to a small neighborhood of the identity $I \in O(d)$. Hence, our strategy will be to construct a packing in $\text{Skew}(d)$, and then push it to $O(d)$ under the \exp mapping.

Formally, given $\epsilon \leq 1/256$, and a $1/2$ pcking of $B_{d-1}(1)$ \mathcal{T} , define for $w \in \mathcal{T}$ the matrix

$$M(w) := \epsilon \left(e_1(0, w)^\top + (0, w)e_1^\top \right) \in \text{Skew}(d),$$

where e_1 denotes the first canonical basis vector in \mathbb{R}^d . Observe that $\|M(w)\|_{\mathbb{F}} = \sqrt{2\|w\|^2} = \sqrt{2}\|w\|$ and, since the singular value of $M(w)$ are paired, we have $\|M(w)\|_{\text{op}} = \|w\|$. Hence, for every $w_1 \neq w_2 \in \mathcal{B}_{d-1}(1)$, we have

$$\begin{aligned} \|M(w_1 - w_2)\|_{\text{op}} &= \epsilon\|w_1 - w_2\|_2 \geq \epsilon/2 \text{ and} \\ \|M(w_1 - w_2)\|_{\mathbb{F}} &= \sqrt{2}\epsilon\|w_1 - w_2\|_2 \leq 2\sqrt{2}\epsilon. \end{aligned}$$

Now, we define our packing $\overline{\mathcal{N}}$ formally as

$$\overline{\mathcal{N}} := \{\exp(M(w)) : w \in \mathcal{T}\}.$$

We now introduce the following lemma, proved in Section F.4.1, which precisely describes the extent to which $\exp(\cdot)$ is an isometry:

Lemma F.4 *Let $\|\cdot\|$ be a sub-multiplicative norm (e.g., $\|\cdot\|_{\text{op}}$ or $\|\cdot\|_F$), and $X, Y \in \mathbb{R}^{d \times d}$. Then,*

$$\|\exp(X + Y) - \exp(X) - Y\| \leq e^{2K} - 1 - 2K, \text{ where } K = \max\{\|X\|, \|Y\|\}. \quad (6.42)$$

We apply the above with $Y = M(w_1) - M(w_2) = M(w_1 - w_2)$, and $X = M(w_2)$. Then, $X + Y = M(w_1)$, $\max\{\|X\|_{\text{op}}, \|Y\|_{\text{op}}\} \leq 2\epsilon$, and $\max\{\|X\|_F, \|Y\|_F\} \leq 2\sqrt{2}\epsilon$. Hence, Lemma F.4 implies that

$$\begin{aligned} \|\exp(M(w_1)) - \exp(M(w_2)) - M(w_1 - w_2)\|_{\text{op}} &\leq e^{8\epsilon} - 1 - 8\epsilon \text{ and} \\ \|\exp(M(w_1)) - \exp(M(w_2)) - M(w_1 - w_2)\|_F &\leq e^{8\sqrt{2}\epsilon} - 1 - 8\sqrt{2}\epsilon. \end{aligned} \quad (6.43)$$

We can upper bound both displays in (6.43) using the following short technical lemma:

Lemma F.5 *Let $t \in [0, \log 2]$. Then $e^t - 1 - t \leq t^2$.*

Proof Let $f(t) = e^t - 1 - t$, and $g(t) = t^2$. Then, $f(0) = f'(0) = g(0) = g'(0) = 0$. Moreover, $f''(t) = e^t$, and $g''(t) = 2$. Hence, as long as $0 \leq t \leq \log 2$, $f(t) = \int_0^t \int_0^u f''(s) ds du \leq \int_0^t \int_0^u g''(s) ds du = g(t)$. \blacksquare

Hence for $\epsilon \leq \log 2 / 4\sqrt{2} \leq 1/256$, (6.43) and Lemma F.5 combine to imply

$$\begin{aligned} \|\exp(M(w_1)) - \exp(M(w_2)) - M(w_1 - w_2)\|_{\text{op}} &\leq 64\epsilon^2 \text{ and} \\ \|\exp(M(w_1)) - \exp(M(w_2)) - M(w_1 - w_2)\|_F &\leq 128\epsilon^2. \end{aligned}$$

Hence, by the triangle inequality, for $\epsilon \leq 1/256$

$$\begin{aligned} \|\exp(M(w_1)) - \exp(M(w_2))\|_{\text{op}} &\geq \|M(w_1 - w_2)\|_{\text{op}} - 64\epsilon^2 \\ &= \|w_1 - w_2\| - 64\epsilon^2 \geq \epsilon/2 - 64\epsilon^2 \geq \epsilon/4, \end{aligned}$$

and, again, for $\epsilon \leq 1/256$

$$\begin{aligned} \|\exp(M(w_1)) - \exp(M(w_2))\|_F &\leq \|M(w_1 - w_2)\|_F + 128\epsilon^2 \\ &= \sqrt{2}\|w_1 - w_2\| + 128\epsilon^2 \leq 2\sqrt{2}\epsilon + \epsilon/2 \leq 4\epsilon. \end{aligned}$$

F.4.1. PROOF OF LEMMA F.4

Let $M_{i,j}(X, Y)$ denote the homogenous monomial of degree j such consisting of the $\binom{j}{i}$ -products of X i -times, and Y $j - i$ -times. Note then that $M_{j,j}(X, Y) = X^j$, so that $(X + Y)^j - X^j = -X^j + \sum_{i=0}^{j-1} M_{i,j}(X, Y) = \sum_{i=0}^{j-1} M_{i,j}(X, Y)$. Moreover, by the sub-multiplicativity of $\|\cdot\|$, we

have $\|M(X, Y)\| \leq \binom{j}{i} \|X\|_2^i \|Y\|_2^{j-i} = \binom{j}{i} K^j$.

$$\begin{aligned}
 \|\exp(X + Y) - \exp(X) - Y\| &= \|(j!)^{-1} \sum_{j=2}^{\infty} (X + Y)^j - X^j\|_{\text{op}} \\
 &= \left\| \sum_{j=2}^{\infty} (j!)^{-1} \sum_{i=0}^j M_{i,j}(X, Y) \right\| \\
 &\leq \sum_{j=2}^{\infty} (j!)^{-1} \sum_{i=0}^{j-1} \|M_{i,j}(X, Y)\| \\
 &\leq \sum_{j=2}^{\infty} (j!)^{-1} \sum_{i=0}^{j-1} \binom{j}{i} K^j \\
 &\leq \sum_{j=2}^{\infty} (j!)^{-1} (2K)^j = e^{2K} - 1 - 2K.
 \end{aligned}$$

F.4.2. PROOF OF LEMMA F.2

For a matrix M , let M_i denote the i -th row of M , and let $M^{\otimes 2} := M^\top M$

$$\begin{aligned}
 \text{KL}(\mathbb{P}_{\rho O}^{(T)}, \mathbb{P}_A^{(T)}) &= \mathbb{E}_{\rho O} \left[\sum_{t=1}^T \sum_{i=1}^n \langle (\rho O - A)_i, X_t \rangle^2 \right] \\
 &= \mathbb{E}_{\epsilon_1, \dots, \epsilon_T} \left[\sum_{t=1}^T \sum_{i=1}^n \left\langle (\rho O - A)_i, \sum_{s=1}^t \rho^{t-s} O^{t-s} \epsilon_s \right\rangle^2 \right] \\
 &= \sum_{t=1}^T \sum_{i=1}^n \left\langle (\rho O - A)_i, \mathbb{E}_{\epsilon_1, \dots, \epsilon_T} \left[\left(\sum_{s=1}^t \rho^{t-s} O^{t-s} \epsilon_s \right)^{\otimes 2} \right] (\rho O - A)_i \right\rangle.
 \end{aligned}$$

We may now compute that, for any $t \in [T]$,

$$\begin{aligned}
 &\mathbb{E}_{\epsilon_1, \dots, \epsilon_T} \left[\left(\sum_{s=1}^t \rho^{t-s} O^{t-s} \epsilon_s \right)^{\otimes 2} \right] \\
 &= \sum_{s=1}^t \mathbb{E}[\rho^{2(t-s)} (O^{2(t-s)})^\top O^{2(t-s)} \epsilon_s^2] + \sum_{1=s \neq s' \leq t} \mathbb{E}[\rho^{2t-s-s'} (O^{2(t-s)})^\top O^{2(t-s')} \epsilon_s \epsilon_{s'}] \\
 &= \sum_{s=1}^t \mathbb{E}[\rho^{2(t-s)} (O^{2(t-s)})^\top O^{2(t-s)}] = \sum_{s=1}^t \rho^{2(t-s)} I.
 \end{aligned}$$

Hence, we have,

$$\begin{aligned}
 \text{KL}(\mathbb{P}_{\rho O}^{(T)}, \mathbb{P}_A^{(T)}) &= \sum_{t=1}^T \sum_{i=1}^n \left\langle (\rho O - A)_i, \left(\sum_{s=1}^t \rho^{2(t-s)} \right) I \cdot (\rho O - A)_i \right\rangle \\
 &= \sum_{t=1}^T \left(\sum_{s=1}^t \rho^{2(t-s)} \right) \sum_{i=1}^n \|(\rho O - A)_i\|_2^2 \\
 &= \|\rho O - A\|_F^2 \sum_{t=1}^T \left(\sum_{s=1}^t \rho^{2(t-s)} \right) = \|\rho O - A\|_F^2 \left(\sum_{t=1}^T \sum_{s=0}^{t-1} \rho^{2s} \right).
 \end{aligned}$$

F.5. Proof of Lemma F.1

Birge's inequality states that $\sum_{i=1}^N \text{KL}(\mathbb{P}_i, \mathbb{P}_0) \geq (1 - \delta) \log(N \frac{1-\delta}{\delta}) + \delta \log(\frac{\delta}{1-\delta/N})$ (Boucheron et al., 2013). Observe that $\delta \log(\frac{\delta}{1-\delta/N}) \geq \delta \log \frac{\delta}{N(1-\delta)} = -\delta \log \frac{N(1-\delta)}{\delta}$. Hence $\sum_{i=1}^N \text{KL}(\mathbb{P}_i, \mathbb{P}_0) \geq (1 - 2\delta) \log(\frac{(1-\delta)N}{\delta}) \geq (1 - 2\delta) \log \frac{N}{2\delta}$ for $\delta < 1/2$. For the second statement, choose $\mathcal{E}_A := \{\|A - \hat{A}(T)\|_{\text{op}} < \epsilon\}$ for $A \in \mathcal{N}$. Since \mathcal{N} is 2ϵ -separated in $\|\cdot\|_{\text{op}}$, all \mathcal{E}_i are disjoint. Hence, for any $A_0 \in \mathcal{N}$

$$(1 - 2\delta) \log(|\mathcal{N}|/2\delta) \leq \frac{1}{|\mathcal{N}| - 1} \sum_{A \in \mathcal{N} - \{A_0\}} \text{KL}(\mathbb{P}_A^{(T)}, \mathbb{P}_{A_0}^{(T)}) \leq \sup_{A \in \mathcal{N} - \{A_0\}} \text{KL}(\mathbb{P}_A^{(T)}, \mathbb{P}_{A_0}^{(T)}). \quad (6.44)$$

Since A_0 was arbitrary, we may pass to an inf over all $A_0 \in \mathcal{N}$.

Appendix G. Analysis of Standard LS

Here, we show that given a regression of the form $Y_t = A_* X_t + \eta_t$ for $Y_t, X_t \in \mathbb{R}^d$, $A_* \in \mathbb{R}^{d \times d}$, where η_t are an i.i.d. Gaussian noise sequence, then $\|\hat{A} - A_*\|_{\text{op}} \gtrsim \sqrt{d/\lambda_{\min}(\sum_{t=1}^T X_t X_t^\top)}$ is in fact necessary. This is a consequence of the fact that the operator-norm error involves a supremum over all directions in \mathcal{S}^{d-1} . Formally,

Theorem G.1 *Let $X_1, \dots, X_T \in \mathbb{R}^d$ be an arbitrary dynamical process, let $\eta_1, \dots, \eta_T \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I_d)$, and independence of X_1, \dots, X_T . Then, given observations $Y_t = A_* X_t + \eta_t$ for $A_* \in \mathbb{R}^{d \times d}$ and $t \in \{1, \dots, T\}$, the least-squares estimator \hat{A} satisfies the lower bound*

$$\mathbb{E} \left[\|\hat{A} - A_*\|_{\text{op}}^2 \mid X_1, \dots, X_T \right] \geq \frac{d}{\lambda_{\min} \left(\sum_{t=1}^T X_t X_t^\top \right)}$$

Proof Conditioning on X_1, \dots, X_T , we may assume without loss of generality that \mathbf{X} is deterministic. We let $\mathbf{X} \in \mathbb{R}^{T \times d}$ denote the matrix whose rows are X_t , and $\mathbf{E} \in \mathbb{R}^{T \times d}$ denote the matrix whose rows are ϵ_t . Then, $\hat{A}(T) - A_* = \mathbf{X}^\dagger \mathbf{E}$. Moreover, if $v_* \in \arg \max_{v: \|v\|_2=1} \|v^\top \mathbf{X}^\dagger\|_2$, then $\|v_*^\top \mathbf{X}^\dagger\|_2 = \lambda_{\min}(\sum_{t=1}^T X_t X_t^\top)^{-1/2}$. Moreover, v_* depends only on \mathbf{X} which by construction is

independent of \mathbf{E} . Hence, $v^\top \mathbf{X}^\dagger \mathbf{E} \sim \mathcal{N}(0, \lambda_{\min}(\sum_{t=1}^T X_t X_t)^{-1} \cdot I_d)$. And hence,

$$\begin{aligned} \mathbb{E} \left[\|\widehat{A}(T) - A_*\|_{\text{op}}^2 \right] &= \mathbb{E} \left[\|\mathbf{X}^\dagger \mathbf{E}\|_{\text{op}}^2 \right] = \mathbb{E} \left[\sup_{v \in \mathcal{S}^{d-1}} \|v^\top \mathbf{X}^\dagger \mathbf{E}\|_2^2 \right] \\ &\geq \mathbb{E} \left[\|v_*^\top \mathbf{X}^\dagger \mathbf{E}\|_2^2 \right] \\ &= \mathbb{E} \left[\|w\|^2 \mid w \sim \mathcal{N}(0, \lambda_{\min}(\sum_{t=1}^T X_t X_t)^{-1} \cdot I_d) \right] = d \lambda_{\min} \left(\sum_{t=1}^T X_t X_t \right)^{-1}. \end{aligned}$$

■