Robust System Identification:
Finite-sample Guarantees and Connection to Regularization

Hyuk Park, Grani A. Hanasusanto, and Yingying Li

Department of Industrial and Enterprise Systems Engineering

University of Illinois Urbana-Champaign, United States

Abstract

We address the problem of identifying a stable linear time-invariant system from a single sample trajectory. The least squares estimate (LSE) is a commonly used algorithm for this purpose. However, LSE may exhibit poor identification errors when the number of samples is small. To mitigate the issue, we introduce the robust LSE, which integrates robust optimization techniques. We demonstrate that our robust LSE is equivalent to regularizing LSE using general Schatten p-norms. Moreover, we provide finite-sample analyses for the robust LSE, which can be directly transferred to the *regularized* LSE due to their equivalence. We showcase the empirical performance of our method in system identification tasks. Additionally, we combine our robust LSE with several online adaptive linear quadratic control algorithms and demonstrate that our method significantly outperforms existing approaches in regret.

Keywords: robust optimization; system identification; finite-sample analysis

1 Introduction

Many real-world engineering problems require learning unknown parameters from sample data. Specifically, in the control community, the problem of estimating the parameters of a dynamical system is referred to as *system identification*. System identification is crucial since accurate estimation of the unknown system is integral to developing safe and reliable control systems.

In this work, we focus on identifying linear time-invariant (LTI) systems from a single trajectory of fully observed states. One of the simplest system identification algorithms is the least squares method or the least squares estimate (LSE), which minimizes the squared prediction errors of the given samples. Due to the stochastic nature of the data, the performance of LSE cannot be deterministically guaranteed. Moreover, since the data comprises a single trajectory of states resulting from the evolution of the dynamical system, the samples are non-i.i.d. Recent works [1, 2] provide a non-asymptotic analysis of the least squares estimator (LSE), specifically addressing system identification errors with respect to a finite number of non-i.i.d. samples. Both works show that the error decays as fast as $\widetilde{\mathcal{O}}(1/\sqrt{T})$ where T denotes the number of samples.

Although these theoretical results are satisfactory, the empirical performance of LSE may suffer, especially when only a few samples are available. This limitation can be critical in applications where data collection is inherently limited or very expensive. To address this issue, we propose a new approach that combines robust optimization with LSE by formulating a min-max optimization problem, referred to as the *robust LSE problem*. We establish that the robust LSE problem can be formulated as a convex semidefinite program (SDP), making it tractably solvable. Furthermore, we provide finite-sample analysis for our approach and show that the robust LSE can achieve the same system identification error rate as LSE. Interestingly, we can show that our robust LSE problem is equivalent to the LSE problem with an additional regularization term given by the general Schatten p-norm. A few special cases of the Schatten p-norms have previously been introduced to regularize the LSE problems [3, 4]. However, to the best of our knowledge, there is no existing finite-sample analysis for the LSE problem with general Schatten p-norm regularization in the literature. Therefore, the equivalence provides new insights into the regularized LSE. The contributions of this paper are summarized as follows:

1. We introduce a novel system identification algorithm by combining robust optimization with the LSE problem. We further show that our robust LSE problem is equivalent to the LSE problem with the Schatten p-norm regularization term.

2. We provide finite-sample analyses, showing that our system identification error rate is $\widetilde{\mathcal{O}}(1/\sqrt{T})$. Due to the equivalence, our result provides a new finite-sample analysis for the regularized LSE.

3. We conduct numerical experiments to demonstrate the superior performance of our approach compared to LSE in system identification tasks. Additionally, we evaluate its effectiveness in adaptive control tasks by integrating our robust LSE with existing adaptive linear quadratic (LQ) control algorithms and showcasing significantly better performance compared to the existing algorithms.

Futher Literature Review

There has been a recent emergence of interest in deriving the finite-sample identification errors of LTI systems. Most works focus on analyzing performance of the *standard* LSE [1, 2, 5, 6, 7, 8, 9, 10]. One advantage of analyzing the standard LSE is that the system identification error term, which is the main interest of the analysis, can be analytically obtained using the solution to the LSE problem. This term can then be broken down in various ways, enabling different approaches to address the resulting components.

While theoretical guarantees for LSE appear promising, it often struggles in small data regimes, resulting in subpar estimates [4]. We employ robust optimization techniques [11] to enhance the resilience of LSE. The key idea of robust optimization is to find solutions that perform optimally against the worst-case realizations of uncertain data. The robust approach is not entirely new in the literature. In [12], a similar approach to ours is explored in a general learning context where i.i.d. samples are readily accessible, simplifying the statistical analysis. The authors in [10] also assume i.i.d. samples and utilize the standard LSE for system identification. They construct an uncertainty set of system parameters around the resulting estimate. They then solve a min-max problem, referred to as the robust LQR problem, to determine the best control input against the worst-case system parameter in the uncertainty set. In contrast, our approach directly formulates a min-max problem for system identification, seeking the best estimate against the worst-case realizations of the data. There is no existing work proposing the robust LSE problem formulation in this paper under the single trajectory assumption.

As stated in our contributions, the finite-sample analysis proposed in this paper is not limited to our method. It can be extended to the regularized LSE problem, where the regularization term is defined as the Schatten p-norm of a quadratic function of system parameters multiplied by a user-defined (scalar) tuning parameter—henceforth referred to as the regularization parameter. Special cases of the Schatten p-norm regularization are proposed in the literature [3, 4]. In [3], they introduce the squared Frobenius norm of system parameters with the regularization parameter set to a some small *fixed* value, hence the convergence of their estimate to the true system parameter is not guaranteed. In [4], they consider LTI systems with limited state observations, i.e., states cannot be directly observed. They introduce the nuclear norm of the Hankel matrix to their LSE problem and derive the finite-sample impulse response estimation errors only for the MISO (multi-input single-output) system under the assumption that i.i.d. samples are available.

Notation

Bold lower-case letter $\boldsymbol{x}$ and upper-case letter $\boldsymbol{X}$ represent a vector and a matrix, respectively, while regular font x indicates a scalar. An $n \times n$ dimensional identity matrix is denoted as $\mathbb{I}_n$. For any real-valued $n \times m$ matrix $\boldsymbol{X} \in \mathbb{R}^{n \times m}$, $\|\boldsymbol{X}\|_p$ represents the Schatten p-norm of the matrix, which is defined as $\|\boldsymbol{X}\|_p = (\mathrm{tr}(\boldsymbol{X}^\top \boldsymbol{X})^{p/2})^{1/p}$. For several special cases of the Schatten p-norm, we may interchangeably use the following notations: nuclear norm $\|\cdot\|_* = \|\cdot\|_1$, Frobenius norm $\|\cdot\|_F = \|\cdot\|_2$, and operator norm $\|\cdot\| = \|\cdot\|_\infty$. In the entire paper, we will *not* use matrix norms induced by vector norms to prevent any confusion. We may use $\|\boldsymbol{x}\|_2$ for the Euclidean norm (i.e., ℓ^2 norm). In this case, the notation should be still clear since the norm is taken on a bold lower-case. For any square matrix $\boldsymbol{X} \in \mathbb{R}^{n \times n}$, the trace operation is denoted as $\mathrm{tr}(\boldsymbol{X})$. The spectral radius denoted $\rho(\boldsymbol{X})$ is the largest absolute value of the eigenvalues of $\boldsymbol{X}$. To denote a $n \times n$ positive (semi)definite matrix $\boldsymbol{X}$, we may interchangeably use $\boldsymbol{X} \in \mathbb{S}^n_{++}$ ($\boldsymbol{X} \in \mathbb{S}^n_{+}$) and $\boldsymbol{X} \succ 0$ ($\boldsymbol{X} \succeq 0$).

2 Problem Statement

In this work, our goal is to identify an unknown discrete-time linear time-invariant (LTI) system

$$\boldsymbol{x}_{t+1} = \boldsymbol{A}^\star \boldsymbol{x}_t + \boldsymbol{B}^\star \boldsymbol{u}_t + \boldsymbol{w}_t, \quad t = 0, \ldots, T-1. \tag{1}$$

Specifically, we want to recover the system parameter $\boldsymbol{\theta}^\star = [\boldsymbol{A}^\star \ \boldsymbol{B}^\star] \in \mathbb{R}^{n\times(n+m)}$ where $\boldsymbol{A}^\star$ is the state transition matrix and $\boldsymbol{B}^\star \in \mathbb{R}^{n\times m}$ is the control input matrix. $\boldsymbol{x}_t \in \mathbb{R}^n$, $\boldsymbol{u}_t \in \mathbb{R}^m$, and $\boldsymbol{w}_t \in \mathbb{R}^n$ are the state, the control input, and the noise at time t, respectively.

2.1 Assumptions

We state our assumptions about the unknown system (1).

A1. we consider a strictly stable system, i.e., $\rho(\boldsymbol{A}^\star) < 1$.

A2. the data, i.e., observed states of the system (1), is collected in a single trajectory of length $T+1$ denoted as $\{\boldsymbol{x}_t\}_{t=0}^T \in \mathbb{R}^{n(T+1)}$ with the initial state $\boldsymbol{x}_0 = \boldsymbol{0}$.

A3. let $\{\mathcal{F}_t\}_{t\geq 0}$ be a filtration and $\{\boldsymbol{x}_t\}_{t\geq 0}$ be a stochastic process such that $\boldsymbol{x}_t$ is $\mathcal{F}_{t-1}$ measurable.

A4. the noise $\boldsymbol{w}_t$ is a martingale difference sequence with respect to $\mathcal{F}_t$ with $\mathbb{E}[\boldsymbol{w}_t|\mathcal{F}_{t-1}] = \boldsymbol{0}$ and $\mathbb{E}[\boldsymbol{w}_t\boldsymbol{w}_t^\top|\mathcal{F}_{t-1}] = \boldsymbol{\Sigma}_w \succeq \sigma_w^2\mathbb{I}_n \succ \boldsymbol{0}$.

A5. Furthermore, we assume that $\boldsymbol{w}_t$ is a σ_w^2-conditionally sub-Gaussian random vector with respect to $\mathcal{F}_t$, i.e., for any unit vector $\boldsymbol{v} \in \mathbb{R}^n$, the inner-product $\boldsymbol{v}^\top\boldsymbol{w}_t$ is a σ_w^2-sub-Gaussian random variable conditionally on $\mathcal{F}_t$.

A6. the control input $\boldsymbol{u}_t$ is a σ_u^2-sub-Gaussian random vector with $\mathbb{E}[\boldsymbol{u}_t] = \boldsymbol{0}$ and $\mathbb{E}[\boldsymbol{u}_t\boldsymbol{u}_t^\top] = \sigma_u^2\mathbb{I}_m$. In other words, we inject sub-Gaussian exploration noise into the system to identify the system parameter $\boldsymbol{\theta}^\star$.

These are standard assumptions in the literature. In particular, (**A3.**)-(**A5.**) allow us to make use of tools from the self-normalized process [13], which appears in our theoretical analysis later. The main challenge in our theoretical analysis arises from the single trajectory assumption made in (**A2.**), as the trajectory comprises of non-i.i.d. samples. Due to this difficulty, some previous works rely on a more stringent assumption that T multiple independent trajectories are available, and they take only the last state in each trajectory to ensure that those T samples are i.i.d. Our main theoretical contribution lies in deriving finite-sample guarantees of the proposed method from non-i.i.d. samples, as discussed in Section 4.

2.2 Least Squares Estimate

The least squares estimate (LSE) is one of the commonly used algorithms for system identification. Given a single trajectory $\{\boldsymbol{x}_t\}_{t=0}^T$, the LSE denoted as $\overline{\boldsymbol{\theta}}_T$ minimizes the sum of the squares of the residuals:

$$\overline{\boldsymbol{\theta}}_T = \underset{\boldsymbol{\theta}:=[\boldsymbol{A}\ \boldsymbol{B}]}{\arg\min}\ \frac{1}{T}\sum_{t=0}^{T-1} \|\boldsymbol{x}_{t+1} - \boldsymbol{A}\boldsymbol{x}_t - \boldsymbol{B}\boldsymbol{u}_t\|_2^2 = \underset{\boldsymbol{\theta}}{\arg\min}\ \frac{1}{T}\sum_{t=0}^{T-1} \|\boldsymbol{x}_{t+1} - \boldsymbol{\theta}\boldsymbol{z}_t\|_2^2, \tag{2}$$

where we define $\boldsymbol{z}_t := [\boldsymbol{x}_t^\top \ \boldsymbol{u}_t^\top]^\top \in \mathbb{R}^{n+m}$ as an augmented vector of state and control input. Let us refer to the minimization in (2) as the LSE problem. Note that the objective function in (2) is quadratic in $\boldsymbol{\theta}$.

Therefore, we can rewrite the LSE problem as

$$\min_{\boldsymbol{\theta}} \frac{1}{T}\sum_{t=0}^{T-1}\begin{bmatrix}\boldsymbol{x}_{t+1}\\ \boldsymbol{z}_t\end{bmatrix}^\top \begin{bmatrix}\mathbb{I}_n & -\boldsymbol{\theta}\\ -\boldsymbol{\theta}^\top & \boldsymbol{\theta}^\top\boldsymbol{\theta}\end{bmatrix}\begin{bmatrix}\boldsymbol{x}_{t+1}\\ \boldsymbol{z}_t\end{bmatrix} = \min_{\boldsymbol{\theta}} \operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\widehat{\boldsymbol{\Omega}}_T\right), \tag{3}$$

where $\boldsymbol{G}(\boldsymbol{\theta}) = \begin{bmatrix}\mathbb{I}_n & -\boldsymbol{\theta}\\ -\boldsymbol{\theta}^\top & \boldsymbol{\theta}^\top\boldsymbol{\theta}\end{bmatrix}$ and $\widehat{\boldsymbol{\Omega}}_T = \frac{1}{T}\sum_{t=0}^{T-1}\begin{bmatrix}\boldsymbol{x}_{t+1}\\ \boldsymbol{z}_t\end{bmatrix}\begin{bmatrix}\boldsymbol{x}_{t+1}\\ \boldsymbol{z}_t\end{bmatrix}^\top$.

We can express the *true* LSE problem by substituting $\widehat{\boldsymbol{\Omega}}_T$ in (3) with its expectation, namely, $\boldsymbol{\Omega}_T^\star = \mathbb{E}[\widehat{\boldsymbol{\Omega}}_T]$. Then, one can obtain the system parameter by solving the true LSE problem:

$$\boldsymbol{\theta}^\star = \arg\min_{\boldsymbol{\theta}} \operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}_T^\star\right). \tag{4}$$

From (4), it is clear that obtaining the true system parameter requires knowledge of $\boldsymbol{\Omega}_T^\star$ while an empirical estimate $\widehat{\boldsymbol{\Omega}}_T$ exhibits some estimation errors which are dependent of the data. In other words, a poor estimate $\widehat{\boldsymbol{\Omega}}_T$ may lead to inferior performance, which is commonly the case when the sample size T is small or, in our context, a short trajectory of observations $\{\boldsymbol{x}_t\}_{t=0}^T$. In fact, statistical analyses regarding the LSE performance revolve around understanding the sum of the outer product $\sum_{t=0}^{T-1}\boldsymbol{x}_{t+1}\boldsymbol{x}_{t+1}^\top$, known as the Gram matrix. The expected Gram matrix $\mathbb{E}[\sum_{t=0}^{T-1}\boldsymbol{x}_{t+1}\boldsymbol{x}_{t+1}^\top]$, which coincides with the first diagonal block in $\boldsymbol{\Omega}_T^\star$, can be nicely represented as a matrix-valued function of the unknown system parameter $\boldsymbol{\theta}^\star$:

$$\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{x}_{t+1}\boldsymbol{x}_{t+1}^\top\right] = \sum_{t=0}^{T-1}\boldsymbol{\Gamma}_t(\boldsymbol{\theta}^\star) = \sum_{t=0}^{T-1}\sum_{s=0}^{t}({\boldsymbol{A}^\star}^\top)^s(\sigma_u^2\boldsymbol{B}^\star{\boldsymbol{B}^\star}^\top + \boldsymbol{\Sigma}_w)(\boldsymbol{A}^\star)^s. \tag{5}$$

Of course, the expected Gram matrix (5) is not accessible to us since it requires $\boldsymbol{\theta}^\star$.

3 Robust Least Squares Estimate

As pointed out earlier, $\widehat{\boldsymbol{\Omega}}_T$ based on T non-i.i.d. samples may fail to accurately estimate $\boldsymbol{\Omega}_T^\star$ when T is small. Even with sufficiently large T, $\overline{\boldsymbol{\theta}}_T$ in (2) may perform poorly in practical applications where our assumptions about the system (1) are not satisfied, e.g., the data can be contaminated by some unmodelled factors such as non-zero or time-correlated noise processes. To address the issue, we formulate a robust version of the LSE problem to obtain the robust estimate denoted as $\widehat{\boldsymbol{\theta}}_T$:

$$\widehat{\boldsymbol{\theta}}_T = \arg\min_{\boldsymbol{\theta}} \max_{\boldsymbol{\Omega}\in\mathcal{U}_T^{p,\epsilon}} \operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}\right) \text{ where } \mathcal{U}_T^{p,\epsilon} = \left\{\boldsymbol{\Omega}\in\mathbb{S}_+^{2n+m} : \|\boldsymbol{\Omega}-\widehat{\boldsymbol{\Omega}}_T\|_p \le \epsilon\right\}. \tag{6}$$

The proposed approach (6) first constructs the uncertainty set $\mathcal{U}_T^{p,\epsilon}$ which contains all positive semidefinite matrices $\boldsymbol{\Omega}$ that are within a distance of $\epsilon \ge 0$ from the estimate $\widehat{\boldsymbol{\Omega}}_T$ in the Schatten p-norm. Then, it seeks a minimizer $\widehat{\boldsymbol{\theta}}_T$ that performs best under the worst-case matrix $\boldsymbol{\Omega}$ in $\mathcal{U}_T^{p,\epsilon}$. However, the min-max problem in (6) is difficult to solve directly since the objective function involves a maximization problem. In the following, we introduce an equivalent semidefinite program (SDP) for the robust LSE problem.

Theorem 1. *For any given uncertainty set parameters $p \geq 1$ (as in the Schatten p-norm) and $\epsilon \geq 0$, the robust LSE problem in (6) can be equivalently reformulated as the SDP*

$$\begin{aligned} \min \quad & \mathrm{tr}(\boldsymbol{\Gamma}\widehat{\boldsymbol{\Omega}}_T) + \epsilon\|\boldsymbol{\Gamma}\|_q \\ \text{s.t.} \quad & \boldsymbol{\theta} \in \mathbb{R}^{n\times(n+m)}, \quad \boldsymbol{\Gamma} \in \mathbb{S}_+^{2n+m}, \quad \boldsymbol{H} \in \mathbb{S}_+^{n+m}, \\ & \boldsymbol{\Gamma} \succeq \begin{bmatrix} \mathbb{I}_n & -\boldsymbol{\theta} \\ -\boldsymbol{\theta}^\top & \boldsymbol{H} \end{bmatrix}, \\ & \begin{bmatrix} \mathbb{I}_n & \boldsymbol{\theta} \\ \boldsymbol{\theta}^\top & \boldsymbol{H} \end{bmatrix} \succeq \mathbf{0}, \end{aligned} \tag{7}$$

where $\|\cdot\|_q$ is the dual Schatten norm of $\|\cdot\|_p$, that is, q such that $\frac{1}{p} + \frac{1}{q} = 1$.

Proof. Dualizing the inner maximization problem with the constraint $\|\boldsymbol{\Omega}-\widehat{\boldsymbol{\Omega}}_T\|_p \leq \epsilon$ given by our uncertainty set, we have

$$\max_{\boldsymbol{\Omega}\succeq\mathbf{0}} \min_{\lambda\geq 0} \mathrm{tr}(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}) + \lambda\epsilon - \lambda\|\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}}_T\|_p$$

$$= \max_{\boldsymbol{\Omega}\succeq\mathbf{0}} \min_{\lambda\geq 0} \mathrm{tr}(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}) + \lambda\epsilon - \max_{\|\boldsymbol{\Gamma}\|_q\leq\lambda} \mathrm{tr}\left(\boldsymbol{\Gamma}\left(\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}}_T\right)\right) \tag{8}$$

$$= \max_{\boldsymbol{\Omega}\succeq\mathbf{0}} \min_{\lambda\geq 0} \mathrm{tr}(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}) + \lambda\epsilon + \min_{\|\boldsymbol{\Gamma}\|_q\leq\lambda} \mathrm{tr}\left(\boldsymbol{\Gamma}\left(\widehat{\boldsymbol{\Omega}}_T - \boldsymbol{\Omega}\right)\right) \tag{9}$$

$$= \min_{\substack{\lambda\geq 0,\\ \|\boldsymbol{\Gamma}\|_q\leq\lambda}} \mathrm{tr}\left(\boldsymbol{\Gamma}\widehat{\boldsymbol{\Omega}}_T\right) + \lambda\epsilon + \max_{\boldsymbol{\Omega}\succeq\mathbf{0}} \mathrm{tr}((\boldsymbol{G}(\boldsymbol{\theta}) - \boldsymbol{\Gamma})\boldsymbol{\Omega}) \tag{10}$$

$$= \min_{\substack{\lambda\geq 0,\\ \|\boldsymbol{\Gamma}\|_q\leq\lambda}} \mathrm{tr}\left(\boldsymbol{\Gamma}\widehat{\boldsymbol{\Omega}}_T\right) + \lambda\epsilon \ \text{ s.t. } \boldsymbol{\Gamma} \succeq \begin{bmatrix} \mathbb{I}_n & -\boldsymbol{\theta} \\ -\boldsymbol{\theta}^\top & \boldsymbol{\theta}^\top\boldsymbol{\theta} \end{bmatrix} \tag{11}$$

$$= \min_{\substack{\lambda\geq 0,\\ \|\boldsymbol{\Gamma}\|_q\leq\lambda,\\ \boldsymbol{H}\succeq\mathbf{0}}} \mathrm{tr}\left(\boldsymbol{\Gamma}\widehat{\boldsymbol{\Omega}}_T\right) + \lambda\epsilon \ \text{ s.t. } \boldsymbol{\Gamma} \succeq \begin{bmatrix} \mathbb{I}_n & -\boldsymbol{\theta} \\ -\boldsymbol{\theta}^\top & \boldsymbol{H} \end{bmatrix}, \begin{bmatrix} \mathbb{I}_n & \boldsymbol{\theta} \\ \boldsymbol{\theta}^\top & \boldsymbol{H} \end{bmatrix} \succeq \mathbf{0}. \tag{12}$$

In the first equality (8), we use the definition of the dual norm for $\lambda\|\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}}_T\|_p$. As in the second equality (9), we can convert the maximization to a minimization since $\max f(\cdot) = -\min -f(\cdot)$. The third equality (10) exploits strong duality by following the standard results of the convex analysis (see Theorem 1, Chapter 8 in [14]). The feasible set of $(\lambda, \boldsymbol{\Gamma})$ defined in (9) is a convex set, and the objective function of the inner minimization problem is convex in $(\lambda, \boldsymbol{\Gamma})$. Furthermore, we can show the existence of an interior point in the feasible set, that is, there always exists some $\boldsymbol{\Gamma}$ such that the following strict inequality holds: $\|\boldsymbol{\Gamma}\|_q < \lambda$ for any $\lambda > 0$. Hence, strong duality holds. Then, the maximization over $\boldsymbol{\Omega}$ in (10) leads to a restriction of the feasible set which is given by the constraint in (11). In other words, $(\boldsymbol{G}(\boldsymbol{\theta}) - \boldsymbol{\Gamma})$ in (10) needs to be negative semidefinite. In the last equality, we linearize the quadratic term $\boldsymbol{\theta}^\top\boldsymbol{\theta}$ by following Lemma 4 in [15]. Then, we can combine the minimization in (12) with the minimization over $\boldsymbol{\theta}$ in (7). Finally, reversing the epigraphic reformulation $\|\boldsymbol{\Gamma}\|_q \leq \lambda$ in the equality (12) yields the problem formulation (7), which is a semidefinite program. □

Note that the Schatten p-norm defined in (6) corresponds to the Schatten q-norm in the objective function in (7). For any $q \geq 1$, the reformulation (7) is a convex SDP. In particular, for several choices of q such as $q = 1, 2, \infty$, it is readily solvable by off-the-shelf commercial solvers.

Interestingly, the robust LSE problem exhibits an equivalence to the regularized LSE problem as shown in the following corollary.

Corollary 1. *For any given uncertainty set parameters $p \geq 1$ and $\epsilon \geq 0$, the robust LSE problem* (6) *is equivalent to the LSE problem with the Schatten q-norm regularization term as follows:*

$$\min_{\boldsymbol{\theta}} \operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\widehat{\boldsymbol{\Omega}}_T\right) + \epsilon\|\boldsymbol{G}(\boldsymbol{\theta})\|_q. \tag{13}$$

Proof. By reversing the epigraphic reformulation $\|\boldsymbol{\Gamma}\|_q \leq \lambda$ in (11), we have

$$\min_{\boldsymbol{\Gamma},\boldsymbol{\theta}} \ \operatorname{tr}\left(\boldsymbol{\Gamma}\widehat{\boldsymbol{\Omega}}_T\right) + \epsilon\|\boldsymbol{\Gamma}\|_q \ \text{ s.t. } \boldsymbol{\Gamma} \succeq \underbrace{\begin{bmatrix} \mathbb{I}_n & -\boldsymbol{\theta} \\ -\boldsymbol{\theta}^\top & \boldsymbol{\theta}^\top\boldsymbol{\theta} \end{bmatrix}}_{=\boldsymbol{G}(\boldsymbol{\theta})}. \tag{14}$$

Suppose that $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C} \succeq \mathbf{0}$ and $\boldsymbol{A} \succeq \boldsymbol{B}$. Then, the following is true: $\operatorname{tr}(\boldsymbol{AC}) \geq \operatorname{tr}(\boldsymbol{BC})$. Recall that positive semidefinite inequality $\succeq$ implies ordering on matrices known as Loewner's ordering. One property of the Loewner's ordering is that $\boldsymbol{A} \succeq \boldsymbol{B} \Rightarrow \sigma_i(\boldsymbol{A}) \geq \sigma_i(\boldsymbol{B})$ for all i where $\sigma_i(\cdot)$ denotes the i-th singular value of the corresponding matrix (note that the converse is not necessarily true). Also, by definition, the Schatten q-norm is equivalent to the ℓ^q-norm of the vector of singular values, i.e., $\|\boldsymbol{A}\|_q = \|[\sigma_1(\boldsymbol{A}), \ \dots \ , \ \sigma_n(\boldsymbol{A})]^\top\|_q = (\sum_{i=1}^n |\sigma_i(\boldsymbol{A})|^q)^{1/q}$. Using these facts, we can conclude that $\boldsymbol{\Gamma} = \boldsymbol{G}(\boldsymbol{\theta})$ holds when $\boldsymbol{\Gamma}$ and $\boldsymbol{\theta}$ are the minimizer of the problem (14). Hence, the problem (14) is equivalent to (13). □

A few remarks are in order. If the nuclear norm (i.e., $q = 1$) is used in (13), then we have $\epsilon\|\boldsymbol{G}(\boldsymbol{\theta})\|_* = \epsilon\|\boldsymbol{\theta}\|_F^2 + \epsilon n$. Thus, the regularization term simplifies to a squared Frobenius norm regularization on $\boldsymbol{\theta}$, and the resulting problem constitutes a tractable quadratic program. Corollary 1 further draws an interesting connection between the robust LSE and the regularized LSE in [3]. In that work, the regularization parameter ϵ is set to a small *fixed* value. However, it lacks a clear explanation of how the regularization impacts the performance of the LSE since it is introduced merely to ensure the invertibility of the Gram matrix $\sum_{t=0}^{T-1} \boldsymbol{x}_{t+1}\boldsymbol{x}_{t+1}^\top$. In this case, a convergence rate on the system identification error cannot be established. Our result, therefore, not only provides a justification for the use of squared Frobenius norm regularization but also guidance on how to control the parameter as the sample size T increases. The recent work [4] uses a Hankel nuclear norm regularization to identify low-order linear systems using i.i.d. trajectories. They recognize that the regularization term yields better performance than the unregularized LSE when the number of samples is small. However, to the best of our knowledge, there is no finite-sample analysis for the LSE with general Schatten norm regularization under the single trajectory assumption. In Section 4, we provide finite-sample analyses for the robust LSE problem, which ultimately results in system identification errors.

4 Performance Guarantees

In this section, we discuss the finite-sample guarantees of our robust LSE under the single trajectory assumption. The ultimate goal of our analysis is to achieve the optimal rate of system identification errors in terms of the number of samples T, which is known to be $\widetilde{\mathcal{O}}(1/\sqrt{T})$ under the i.i.d. assumption. For the *unregularized* LSE (2), henceforth referred to as the *standard* LSE, there have been many works providing such an optimal rate. Unlike the regularized LSE, the main advantage of analyzing the standard LSE is that the system identification error term can be analytically obtained using the solution to the LSE problem in (2), namely, $\overline{\boldsymbol{\theta}}_T - \boldsymbol{\theta}^\star = (\sum_{t=0}^{T-1} \boldsymbol{w}_t \boldsymbol{z}_t^\top)(\sum_{t=0}^{T-1} \boldsymbol{z}_t \boldsymbol{z}_t^\top)^{-1}$. Then, the error term can be decomposed in various ways to facilitate different analyses [1, 2, 5]. However, these decomposition methods do not apply to our robust LSE since the identification error term for the robust LSE, i.e., $\widehat{\boldsymbol{\theta}}_T - \boldsymbol{\theta}^\star$ is no longer expressed in a convenient analytical form, hence, requires a different analysis.

We first provide the finite-sample coverage guarantee of our uncertainty set in (6), which is eventually used as the main ingredient for our system identification error analysis.

Proposition 1. *(Finite-sample coverage guarantee). For any significance level $\delta \in (0, 1]$, we have*

$$\mathbb{P}\left[\boldsymbol{\Omega}_T^\star \in \mathcal{U}_T^p(\delta)\right] \geq 1 - \delta, \tag{15}$$

Here, $\mathcal{U}_T^p(\delta)$ is the uncertainty set for the robust LSE problem in (6) *defined as follows:*

$$\mathcal{U}_T^p(\delta) = \left\{\boldsymbol{\Omega} \in \mathbb{S}_+^{2n+m} : \|\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}}_T\|_p \leq \epsilon(\delta)\right\} \text{ and } \epsilon(\delta) = \widetilde{\mathcal{O}}(1/\sqrt{T}). \tag{16}$$

Proof. Proving (15) amounts to showing that the distance between $\boldsymbol{\Omega}_T^\star$ and $\widehat{\boldsymbol{\Omega}}_T$ is small with high probability (w.h.p.): $\|\boldsymbol{\Omega}_T^\star - \widehat{\boldsymbol{\Omega}}_T\|_p \leq \epsilon(\delta)$ with probability (w.p.) at least $1 - \delta$. Here, we derive the upper bound $\epsilon(\delta)$ for $p = \infty$, i.e., the case where the norm in (16) defined by the Schatten ∞-norm (equivalently, operator norm $\|\cdot\| = \|\cdot\|_\infty$). Due to the equivalence of norms, it is easy to show similar bounds for any $p \geq 1$ with different dimensional factors.

Note that $\widehat{\boldsymbol{\Omega}}_T$ can be explicitly expressed as follows:

$$\widehat{\boldsymbol{\Omega}}_T = \frac{1}{T}\sum_{t=0}^{T-1} \begin{bmatrix} \boldsymbol{x}_{t+1} \\ \boldsymbol{x}_t \\ \boldsymbol{u}_t \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{t+1} \\ \boldsymbol{x}_t \\ \boldsymbol{u}_t \end{bmatrix}^\top \tag{17}$$

$$= \frac{1}{T}\sum_{t=0}^{T-1} \begin{bmatrix} \boldsymbol{x}_{t+1}\boldsymbol{x}_{t+1}^\top & \left(\boldsymbol{A}^\star \boldsymbol{x}_t + \boldsymbol{B}^\star \boldsymbol{u}_t + \boldsymbol{w}_t\right)\boldsymbol{x}_t^\top & \left(\boldsymbol{A}^\star \boldsymbol{x}_t + \boldsymbol{B}^\star \boldsymbol{u}_t + \boldsymbol{w}_t\right)\boldsymbol{u}_t^\top \\ \boldsymbol{x}_t\left(\boldsymbol{A}^\star \boldsymbol{x}_t + \boldsymbol{B}^\star \boldsymbol{u}_t + \boldsymbol{w}_t\right)^\top & \boldsymbol{x}_t\boldsymbol{x}_t^\top & \boldsymbol{x}_t\boldsymbol{u}_t^\top \\ \boldsymbol{u}_t\left(\boldsymbol{A}^\star \boldsymbol{x}_t + \boldsymbol{B}^\star \boldsymbol{u}_t + \boldsymbol{w}_t\right)^\top & \boldsymbol{u}_t\boldsymbol{x}_t^\top & \boldsymbol{u}_t\boldsymbol{u}_t^\top \end{bmatrix}. \tag{18}$$

Similarly, $\boldsymbol{\Omega}_T^\star$ is expectation of (18), i.e., $\boldsymbol{\Omega}_T^\star = \mathbb{E}[\widehat{\boldsymbol{\Omega}}_T]$. Hence, using (18), we can establish the following inequalities:

$$\|\boldsymbol{\Omega}_T^\star - \widehat{\boldsymbol{\Omega}}_T\| \leq 2\left(1+\|\boldsymbol{A}^\star\|\right) \underbrace{\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right\|}_{(a)}$$

$$+ 2\underbrace{\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\right\|}_{(b)} + 2\left(1+\|\boldsymbol{A}^\star\|+\|\boldsymbol{B}^\star\|\right)\underbrace{\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{u}_t\boldsymbol{x}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{u}_t\boldsymbol{x}_t^\top\right\|}_{(c)}$$

$$+ \left(1+2\|\boldsymbol{B}^\star\|\right)\underbrace{\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{u}_t\boldsymbol{u}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{u}_t\boldsymbol{u}_t^\top\right\|}_{(d)} + 2\underbrace{\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right\|}_{(e)}.$$

Our goal is to bound each of the terms (a)-(e), and then combine the results to complete the proof.

(a):

Notice that we analyze the difference between the Gram matrix and its expectation with factor $(1/T)$. Similar results are discussed in [1]. First, we introduce the preparatory result in [1].

Suppose $\rho(\boldsymbol{A}) < 1$ for a matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. Consider a $t \times t$ block Toeplitz matrix

$$\boldsymbol{H}_t = \begin{bmatrix} \mathbb{I}_n & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{A} & \mathbb{I}_n & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \boldsymbol{A}^{t-1} & \boldsymbol{A}^{t-2} & \cdots & \mathbb{I}_n \end{bmatrix} \in \mathbb{R}^{nt\times nt}. \tag{19}$$

Then, for any $t \geq 1$, there exists a finite constant $\mathcal{J}(\boldsymbol{A}) > 0$ that only depends on $\boldsymbol{A}$ such that

$$\|\boldsymbol{H}_t\| \leq \mathcal{J}(\boldsymbol{A}) := \sum_{s=0}^{+\infty}\|\boldsymbol{A}^s\|, \tag{20}$$

where $\mathcal{J}(\boldsymbol{A})$ is specifically the limit of a matrix power series $\sum_{s=0}^{t}\|\boldsymbol{A}^s\|$.

The authors in [1] analyze the sample complexity of the unregularized LSE where an unknown system is uncontrolled. (i.e., identifying only $\boldsymbol{A}^\star$). We can derive a similar result to Lemma 2 in [1].

Under an i.i.d. sub-Gaussian exploration noise, our dynamic system can be written as $\boldsymbol{x}_{t+1} = \boldsymbol{A}^\star\boldsymbol{x}_t + \boldsymbol{\eta}_t$ where $\boldsymbol{\eta}_t$ is a zero mean noise with a covariance matrix $\boldsymbol{\Sigma}_\eta := \mathbb{E}[\boldsymbol{\eta}_t\boldsymbol{\eta}_t^\top] = \sigma_u^2\boldsymbol{B}^\star\boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w$. Then, we can define vectorized states of the system up to time T:

$$\begin{bmatrix} \boldsymbol{x}_1 \\ \vdots \\ \boldsymbol{x}_T \end{bmatrix} = \mathbf{H}_T\boldsymbol{C}_\eta^{1/2}\boldsymbol{\xi} \in \mathbb{R}^{nT} \text{ where } \boldsymbol{C}_\eta = \mathbb{E}\left[\begin{bmatrix} \boldsymbol{\eta}_0 \\ \vdots \\ \boldsymbol{\eta}_{T-1} \end{bmatrix}\begin{bmatrix} \boldsymbol{\eta}_0 \\ \vdots \\ \boldsymbol{\eta}_{T-1} \end{bmatrix}^\top\right] = \begin{bmatrix} \boldsymbol{\Sigma}_\eta & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Sigma}_\eta \end{bmatrix} \in \mathbb{S}_+^{nT}$$

$$\text{and } \boldsymbol{\xi} = \begin{bmatrix} \boldsymbol{\xi}_0 \\ \vdots \\ \boldsymbol{\xi}_{T-1} \end{bmatrix} \in \mathbb{R}^{nT} \text{ is isotropic, i.e., } \mathbb{E}[\boldsymbol{\xi}\boldsymbol{\xi}^\top] = \mathbb{I}_{nT} \tag{21}$$

To simplify the notation, let us define the reciprocal of the square root matrix of the expected Gram matrix as follows:

$$\boldsymbol{M} := \left(\sum_{t=0}^{T-1} \boldsymbol{\Gamma}_t(\boldsymbol{\theta}^\star)\right)^{-1/2} = \left(\sum_{t=0}^{T-1}\sum_{s=0}^{t} (\boldsymbol{A}^{\star\top})^s(\sigma_u^2 \boldsymbol{B}^\star \boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w)(\boldsymbol{A}^\star)^s\right)^{-1/2}.$$

Then, we can establish the following equalities:

$$\|\boldsymbol{M}^\top \sum_{t=0}^{T-1} \boldsymbol{x}_t\boldsymbol{x}_t^\top \boldsymbol{M} - \mathbb{I}_n\| = \sup_{\|\boldsymbol{u}\|_2\le 1} \left| \boldsymbol{u}^\top \left(\boldsymbol{M}^\top \sum_{t=0}^{T-1} \boldsymbol{x}_t\boldsymbol{x}_t^\top \boldsymbol{M} - \mathbb{I}_n\right)\boldsymbol{u}\right| \tag{22}$$

$$= \sup_{\|\boldsymbol{u}\|_2\le 1} \left| \|\sum_{t=0}^{T-1} \boldsymbol{x}_t^\top \boldsymbol{M}\boldsymbol{u}\|_2^2 - \mathbb{E}\left[\left\|\Sigma_{t=0}^{T-1}\boldsymbol{x}_t^\top \boldsymbol{M}\boldsymbol{u}\right\|_2^2\right]\right| \tag{23}$$

$$= \sup_{\|\boldsymbol{u}\|_2\le 1} \left| \|\boldsymbol{\Sigma}_{\boldsymbol{Mu}}^\top \mathbf{H}_{\boldsymbol{T}} \mathbf{C}_\eta^{1/2}\boldsymbol{\xi}\|_2^2 - \mathbb{E}\left[\|\boldsymbol{\Sigma}_{\boldsymbol{Mu}}^\top \mathbf{H}_{\boldsymbol{T}} \mathbf{C}_\eta^{1/2}\boldsymbol{\xi}\|_2^2\right]\right| \tag{24}$$

$$= \sup_{\|\boldsymbol{u}\|_2\le 1} \left| \|\boldsymbol{\Sigma}_{\boldsymbol{Mu}}^\top \mathbf{H}_{\boldsymbol{T}} \mathbf{C}_\eta^{1/2}\boldsymbol{\xi}\|_2^2 - \|\boldsymbol{\Sigma}_{\boldsymbol{Mu}}^\top \mathbf{H}_{\boldsymbol{T}} \mathbf{C}_\eta^{1/2}\|_F^2\right|, \tag{25}$$

where $\boldsymbol{\Sigma}_{\boldsymbol{Mu}} = \begin{bmatrix} \boldsymbol{Mu} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{Mu} \end{bmatrix} \in \mathbb{R}^{nT\times T}$ in (24) is a block diagonal matrix.

The first equality (22) is the variational form of the operator norm. In the last equality (25), we use the fact that $\mathbb{E}[\|\boldsymbol{D}\boldsymbol{\xi}\|_2^2] = \mathrm{tr}(\boldsymbol{D}^\top\boldsymbol{D}\mathbb{E}[\boldsymbol{\xi}\boldsymbol{\xi}^\top]) = \|\boldsymbol{D}\|_F^2$ for an isotropic vector $\boldsymbol{\xi}$. The objective function in (25) can be written as $|\ |\boldsymbol{\xi}^\top\boldsymbol{W}\boldsymbol{\xi}| - |\mathbb{E}[\boldsymbol{\xi}^\top\boldsymbol{W}\boldsymbol{\xi}]|\ |$ where $(\boldsymbol{\xi}^\top\boldsymbol{W}\boldsymbol{\xi})_{\boldsymbol{W}\in\mathcal{W}}$ indexed by a set of matrices $\mathcal{W}$ is referred to as a chaos process.

We omit the remaining steps since they are identical to the proof of Lemma 2 in [1] once we recognize that (25) is the supremum of a chaos process. The main idea for the remaining steps is that the Hanson-Wright inequality [16] provides the concentration bound on (25) when $\boldsymbol{u}$ is fixed. Then, we can use the ϵ-net argument, i.e., discretizing the feasible region $\mathcal{U} = \{\boldsymbol{u} : \|\boldsymbol{u}\|_2 \le 1\}$ and combining the bounds for all $\boldsymbol{u} \in \mathcal{U}(\epsilon)$ where $\mathcal{U}(\epsilon)$ is an ϵ-net of $\mathcal{U}$. Following this idea, for $\delta \in (0, 1]$, we have

$$\Pr\left[\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right\| \le \epsilon_{(a)}(\delta)\right] \ge 1-\delta, \text{ where}$$

$$\epsilon_{(a)}(\delta) = \sigma_w^2 \max\left\{\frac{\sqrt{\|\boldsymbol{M}^{-1}\|\|\boldsymbol{H}_T\|^2\,\|\mathbf{C}_\eta\|\left(\log\left(\frac{2}{\delta}\right)+c_2 n\right)}}{\sqrt{c_1}T}, \frac{\|\boldsymbol{H}_T\|^2\,\|\mathbf{C}_\eta\|\left(\log\left(\frac{2}{\delta}\right)+c_2 n\right)}{c_1 T}\right\}. \tag{26}$$

Note that $\|\boldsymbol{H}_T\|$ in (26) can be further bounded by some finite constant $\mathcal{J}(\boldsymbol{A}^\star)$ due to the preparatory result (20). However, we have not made the explicit dependence of $\epsilon_{(a)}(\delta)$ in terms of T yet as $\|\boldsymbol{M}^{-1}\|$ in (26) grows with T. We defer the discussion to where the bounds on (b) and (c) are established since the same issue arises.

(b) and (c):

The same technique is applied to (b) and (c). Hence, we only show the derivation for (b). Note that since the noise term $\boldsymbol{w}_t$ is independent of $\boldsymbol{x}_t$, the expectation in (b) is a zero matrix. Hence, we only need to analyze $(1/T)\|\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\|$. Assuming $\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top$ is invertible (at the moment), we can break (b) into the product of two terms as follows:

$$\begin{aligned}\frac{1}{T}\Big\|\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\Big\| &= \frac{1}{T}\left\|\left(\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\right)\left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{-1/2}\left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{1/2}\right\| \\ &\le \frac{1}{T}\underbrace{\left\|\left(\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\right)\left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{-1/2}\right\|}_{\text{self-normalized martingale}}\underbrace{\left\|\left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{1/2}\right\|}_{\text{persistent excitation term}}. \qquad (27)\end{aligned}$$

As denoted, the stochastic process in (27) is referred to as the self-normalized martingale whose finite-sample bounds are already analyzed in [3]. Hence, we can invoke the following results to obtain the bound on the self-normalized term.

Suppose that $\boldsymbol{V}_T = \sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top + \boldsymbol{V}$ where $\boldsymbol{V} = c\lfloor T/2\rfloor\boldsymbol{\Gamma}_1(\boldsymbol{\theta}^\star)$ is a positive definite matrix with a universal constant $c > 0$, ensuring the invertibility of $\boldsymbol{V}_T$. Then, for $\delta \in (0,1]$, we have

$$\mathbb{P}\left[\left\|\left(\sum_{t=0}^{T-1}\boldsymbol{w}_t\boldsymbol{x}_t^\top\right)\left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{-1/2}\right\| \le 4\sqrt{\|\boldsymbol{\Sigma}_w\|\log\left(\sqrt{\frac{\det(\boldsymbol{V}_T)}{\det(\boldsymbol{V})}}\cdot\frac{5^n}{\delta}\right)}\right] \ge 1-\delta \qquad (28)$$

$$\text{as long as } T \ge \mathcal{O}\left(n\log\left(\frac{n}{\delta}\right) + \log\left(\frac{\det\boldsymbol{\Gamma}_T(\boldsymbol{\theta}^\star)}{\det\boldsymbol{\Gamma}_1(\boldsymbol{\theta}^\star)}\right)\right). \qquad (29)$$

Note that $\boldsymbol{V}_T$ in (28) is the only term that has dependence on T and it increases at most logarithmically as T grows. We make a few comments before proceeding: i) the bound (28) has to be probabilistic since the invertibility (i.e., positive definiteness) of $\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top$ cannot be guaranteed deterministically; ii) the lower bound on T in (29), i.e., the minimum number of samples that ensures the invertibility of $\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top$ w.h.p., is called the burn-in time. Here, we use the big-O notation for the burn-in time only because we want to streamline the exposition. We make the quantity explicit in the proof of Theorem 2 under sufficient conditions.

Subsequently, we derive an upper bound on the persistent excitation term in (27). Note that the term is similar to one in (26). Hence, we can establish the following inequalities:

$$\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right]^{1/2} - \left(\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right)^{1/2}\right\| \le n^{\frac{1}{4}}\sqrt{\left\|\mathbb{E}\left[\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right] - \sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top\right\|} \le n^{\frac{1}{4}}\sqrt{\epsilon_{(a)}(\delta)} \qquad (30)$$

w.p. at least $1-\delta$.

In the first inequality, we use the following fact: $\|\boldsymbol{A}^{1/2}-\boldsymbol{B}^{1/2}\| \leq \sqrt{\|\boldsymbol{A}-\boldsymbol{B}\|_F} \leq n^{\frac{1}{4}}\sqrt{\|\boldsymbol{A}-\boldsymbol{B}\|}$ for any $\boldsymbol{A},\boldsymbol{B}\in\mathbb{S}_+^n$. The second inequality follows from (26). By the reverse triangle inequality, we can further derive the following upper bound on the persistent excitation term:

$$\left\| \left(\sum_{t=0}^{T-1} \boldsymbol{x}_t \boldsymbol{x}_t^\top \right)^{1/2} \right\| \leq \underbrace{\left\| \mathbb{E}\left[\sum_{t=0}^{T-1} \boldsymbol{x}_t \boldsymbol{x}_t^\top \right]^{1/2} \right\|}_{=\boldsymbol{M}^{-1}} + n^{\frac{1}{4}}\sqrt{\epsilon_{(a)}(\delta)}. \tag{31}$$

Recall that we have not addressed the term $\|\boldsymbol{M}^{-1}\|$ in $\epsilon_{(a)}(\delta)$. In fact, the term $\|\mathbb{E}[\sum_{t=0}^{T-1}\boldsymbol{x}_t\boldsymbol{x}_t^\top]^{1/2}\|$ in (31) is equivalent to $\|\boldsymbol{M}^{-1}\|$ as denoted above. Using the definition of the expected Gram matrix (5), we obtain the following inequalities:

$$\begin{aligned}
&\left\| \mathbb{E}\left[\sum_{t=0}^{T-1} \boldsymbol{x}_t \boldsymbol{x}_t^\top \right]^{1/2} \right\| = \left\| \mathbb{E}\left[\sum_{t=0}^{T-1} \boldsymbol{x}_t \boldsymbol{x}_t^\top \right] \right\|^{1/2} = \left\| \sum_{t=0}^{T} \boldsymbol{\Gamma}_t\left(\theta^\star\right) \right\|^{1/2} \\
=&\left\| \sum_{t=0}^{T}\sum_{s=0}^{t} \left(\boldsymbol{A}^\star\right)^s \left(\sigma_u^2 \boldsymbol{B}^\star \boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w\right) \left(\boldsymbol{A}^{\star\top}\right)^s \right\|^{1/2} \\
\leq&\left\| T\sum_{s=0}^{+\infty} \left(\boldsymbol{A}^\star\right)^s \left(\sigma_u^2 \boldsymbol{B}^\star \boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w\right) \left(\boldsymbol{A}^{\star\top}\right)^s \right\|^{1/2} \\
\leq&\sqrt{T}\left\| \sigma_u^2 \boldsymbol{B}^\star \boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w \right\|^{1/2} \left\| \sum_{s=0}^{+\infty} \left(\boldsymbol{A}^\star\right)^s \right\| \\
=&\sqrt{T}\left\| \sigma_u^2 \boldsymbol{B}^\star \boldsymbol{B}^{\star\top} + \boldsymbol{\Sigma}_w \right\|^{1/2} J\left(\boldsymbol{A}^\star\right) \qquad (32) \\
=&\mathcal{O}(\sqrt{T}).
\end{aligned}$$

The first equality holds since the expected Gram matrix is positive semidefinite and (32) follows from the preparatory result (20). Here, we emphasize $\|\boldsymbol{M}^{-1}\|$ grows at the rate of $\mathcal{O}(\sqrt{T})$. Hence, combining (28) and (31) with the factor $1/T$ yields that (b) is upper-bounded by $\widetilde{\mathcal{O}}(1/\sqrt{T})$. Moreover, since $\|\boldsymbol{M}^{-1}\| = \mathcal{O}(\sqrt{T})$, we can claim that $\epsilon_{(a)}(\delta)$ in (26) is $\mathcal{O}(1/T^{1/4})$.

(d) and (e):

They can be addressed by the standard concentration inequality for a covariance matrix (see Theorem 6.5 in [17]). For (d), under i.i.d. sub-Gaussian exploration noise, we can claim that there exist universal constants $c_1, c_2, c_3 > 0$ such that

$$\mathbb{P}\left[\left\| \mathbb{E}\left[\sum_{t=1}^{T} \boldsymbol{u}_t \boldsymbol{u}_t^\top\right] - \sum_{t=1}^{T} \boldsymbol{u}_t \boldsymbol{u}_t^\top \right\| \leq \epsilon_{(c)}(\delta) \right] \geq 1-\delta, \tag{33}$$

where $\epsilon_{(c)}(\delta) = \sigma_u^2 \cdot c_1 \left(\sqrt{\frac{m}{T}} + \frac{m}{T} \right) + \sigma_u^2 \left(\sqrt{\frac{\log\left(\frac{c_2}{\delta}\right)}{Tc_3}} + \frac{\log\left(\frac{c_2}{\delta}\right)}{Tc_3} \right) = \mathcal{O}(1/\sqrt{T})$. For (e), we can apply the same concentration inequality by defining an augmented random vector $\boldsymbol{v}_t = [\boldsymbol{u}_t^\top \ \boldsymbol{w}_t^\top]^\top$ since

$$\frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=1}^{T}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right]-\sum_{t=1}^{T}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right\| \leq \frac{1}{T}\left\|\mathbb{E}\left[\sum_{t=1}^{T}\boldsymbol{v}_t\boldsymbol{v}_t^\top\right]-\sum_{t=1}^{T}\boldsymbol{v}_t\boldsymbol{v}_t^\top\right\|.$$

Therefore, there exists universal constants $\bar{c}_1, \bar{c}_2, \bar{c}_3 > 0$ such that

$$\mathbb{P}\left[\left\|\mathbb{E}\left[\sum_{t=1}^{T}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right]-\sum_{t=1}^{T}\boldsymbol{w}_t\boldsymbol{u}_t^\top\right\| \leq \epsilon_{(d)}(\delta)\right] \geq 1-\delta, \tag{34}$$

where $\epsilon_{(d)}(\delta) = \max\left(\sigma_u^2, \sigma_w^2\right)\cdot \bar{c}_1\left(\sqrt{\frac{n+m}{T}}+\frac{n+m}{T}\right) + \max\left(\sigma_u^2, \sigma_w^2\right)\left(\sqrt{\frac{\log\left(\frac{\bar{c}_2}{\delta}\right)}{T\bar{c}_3}}+\frac{\log\left(\frac{\bar{c}_2}{\delta}\right)}{T\bar{c}_3}\right) = \mathcal{O}(1/\sqrt{T})$.

Finally, combining (a)-(e) yields the claim. □

We make several comments about Proposition 1 before discussing the system identification errors. Although the upper bound $\epsilon(\delta)$ in (16) can be made more explicit by identifying upper bounds on the universal constants, it would be too conservative and thus lack practical usage by itself. However, the rate $\widetilde{\mathcal{O}}(1/\sqrt{T})$ provides guidance on how to adjust the regularization parameter of our robust LSE as more samples are available. In our numerical experiment, we utilize the cross-validation procedure to determine the initial regularization parameter which provides good performance in various applications.

Subsequently, leveraging the results in Proposition 1, we present finite-sample system identification errors for both the robust LSE and the regularized LSE, given their equivalence.

Theorem 2. *(System identification errors). Suppose that $\epsilon(\delta)$ is the upper bound in* (16)*. Then, for any significance level $\delta \in (0,1]$, as long as*

$$T \geq T(\delta) = \left(\frac{400}{3}\right)\left(\log\left(\frac{1}{\delta}\right) + 2(n+m)\log\left(\frac{200}{3}\right) + \operatorname{logdet}\left(\widetilde{\boldsymbol{\Gamma}}_1\overline{\boldsymbol{\Gamma}}_1^{-1}\right)\right)$$

where

$$\widetilde{\boldsymbol{\Gamma}}_1 = \begin{bmatrix} \boldsymbol{\Gamma}_1(\boldsymbol{\theta}^\star) & \boldsymbol{0} \\ \boldsymbol{0} & \sigma_u^2\mathbb{I}_m \end{bmatrix} \text{ and } \overline{\boldsymbol{\Gamma}}_1 = \frac{n+m}{\delta}\mathbb{E}\left[\boldsymbol{z}_1\boldsymbol{z}_1^\top\right],$$

we have the following system identification errors

$$\mathbb{P}\left[\|\boldsymbol{\theta}^\star - \widehat{\boldsymbol{\theta}}_T\| \leq \frac{\epsilon(\delta)\sqrt{\min\{n,m\}}}{\widehat{\alpha}}\left(2+2\|\boldsymbol{\theta}^\star\| + \left\|\nabla_{\boldsymbol{\theta}}\|\boldsymbol{A}(\boldsymbol{\theta}^\star)\|_q\right\|\right)\right] \geq 1-\delta, \tag{35}$$

where $\widehat{\alpha} = \frac{1}{24}\left(\frac{3}{20}\right)^2\min\left\{\sigma_w^2, \sigma_u^2\right\}$.

Proof. The guarantee (15) in Proposition 1 implies that the following holds:

$$\mathbb{P}\Bigg[\min_{\boldsymbol{\theta}} \underbrace{\operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\boldsymbol{\Omega}_T^\star\right)}_{=f(\boldsymbol{\theta})} \leq \min_{\boldsymbol{\theta}} \underbrace{\operatorname{tr}\left(\boldsymbol{G}(\boldsymbol{\theta})\widehat{\boldsymbol{\Omega}}_T\right) + \epsilon(\delta)\|\boldsymbol{G}(\boldsymbol{\theta})\|_q}_{=g(\boldsymbol{\theta})}\Bigg] \geq 1-\delta.$$

Let $f(\boldsymbol{\theta})$ and $g(\boldsymbol{\theta})$ be the objective function of the true and robust LSE problems, respectively. First, we show that $g(\boldsymbol{\theta})$ is an α-strongly convex function w.h.p. Following the definition of strong convexity, showing strong convexity amounts to showing that $g(\boldsymbol{\theta})$ can be rewritten as $g(\boldsymbol{\theta}) = g'(\boldsymbol{\theta}) + \alpha\|\boldsymbol{\theta}\|_F^2$ where $g'(\boldsymbol{\theta})$ is a convex function and $\alpha > 0$. Note that $\operatorname{tr}(\boldsymbol{G}(\boldsymbol{\theta})\widehat{\boldsymbol{\Omega}}_T)$ contains the convex quadratic function, i.e., $\operatorname{tr}(1/T\sum_{t=0}^{T}\boldsymbol{z}_t\boldsymbol{z}_t^\top\boldsymbol{\theta}^\top\boldsymbol{\theta})$. As shown in [18], under i.i.d. exploration noise, the stochastic process of $\boldsymbol{z}_t = [\boldsymbol{x}_t^\top \ \boldsymbol{u}_t^\top]^\top$ satisfies the block martingale small ball (BMSB) condition with parameters $(k, \widetilde{\boldsymbol{\Gamma}}_{\lfloor k/2\rfloor}, 3/20)$ where parameter k can be set to a positive integer and

$$\widetilde{\boldsymbol{\Gamma}}_{\lfloor k/2\rfloor} = \begin{bmatrix} \boldsymbol{\Gamma}_{\lfloor k/2\rfloor}(\boldsymbol{\theta}^\star) & \mathbf{0} \\ \mathbf{0} & \sigma_u^2\mathbb{I}_m \end{bmatrix} \text{ is the covariance matrix of } \boldsymbol{z}_{\lfloor k/2\rfloor}.$$

It can be shown that the BMSB condition can guarantee the persistent excitation w.h.p. (see Proposition 2.5 in [2]). Therefore, by setting $k=2$, we can establish the following persistent excitation of the stochastic process $\boldsymbol{z}_t$ for $T \geq T(\delta)$ (defined earlier):

$$\mathbb{P}\left[\frac{1}{T}\sum_{t=0}^{T}\boldsymbol{z}_t\boldsymbol{z}_t^\top \succeq \widehat{\alpha}\mathbb{I}_{(n+m)}\right] \geq 1-\delta, \text{ where } \widehat{\alpha} = \frac{1}{16}\left(\frac{3}{20}\right)^2\left(\frac{2}{3}\right)\min\left\{\sigma_w^2, \sigma_u^2\right\}. \tag{36}$$

Hence, we can claim that, for any significance level $\delta \in (0,1]$, $g(\boldsymbol{\theta})$ is $\widehat{\alpha}$-strongly convex w.p. at least $1-\delta$ when T is sufficiently large. Suppose $g(\boldsymbol{\theta})$ is indeed an $\widehat{\alpha}$-strongly convex function. Then, we can upper-bound the system identification errors as follows:

$$\|\boldsymbol{\theta}^\star - \widehat{\boldsymbol{\theta}}_T\|_F \leq \frac{2}{\widehat{\alpha}}\left\|\nabla_{\boldsymbol{\theta}} g\left(\boldsymbol{\theta}^\star\right)\right\|_F \leq \frac{2\sqrt{\min\{n,m\}}}{\widehat{\alpha}}\left\|\nabla_{\boldsymbol{\theta}} g\left(\boldsymbol{\theta}^\star\right)\right\|. \tag{37}$$

The first inequality follows from the properties of strong convexity. The second inequality holds due to the equivalence of norms. To ease the notation, we define the following block matrix notations for $\boldsymbol{\Omega}_T^\star$ and $\widehat{\boldsymbol{\Omega}}_T$:

$$\boldsymbol{\Omega}_T^\star = \begin{bmatrix} \boldsymbol{Q}^\star & \boldsymbol{W}^\star \\ \boldsymbol{W}^{\star\top} & \boldsymbol{E}^\star \end{bmatrix} \text{ and } \widehat{\boldsymbol{\Omega}}_T = \begin{bmatrix} \widehat{\boldsymbol{Q}} & \widehat{\boldsymbol{W}} \\ \widehat{\boldsymbol{W}}^\top & \widehat{\boldsymbol{E}} \end{bmatrix}. \tag{38}$$

Then, we can write the gradient in (37) as $\nabla_{\boldsymbol{\theta}} g\left(\boldsymbol{\theta}^\star\right) = -2\widehat{\boldsymbol{W}} + 2\boldsymbol{\theta}^\star\widehat{\boldsymbol{E}}^\top + \epsilon(\delta)\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q$. Subsequently, we can establish the following inequalities:

$$\begin{aligned}
\left\|\nabla_{\boldsymbol{\theta}} g\left(\boldsymbol{\theta}^\star\right)\right\| &= \left\|-2\widehat{\boldsymbol{W}} + 2\boldsymbol{\theta}^\star\widehat{\boldsymbol{E}}^\top + \epsilon(\delta)\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\right\| \\
&\leq \sup_{\left\|\begin{bmatrix}\Delta\boldsymbol{Q} & \Delta\boldsymbol{W} \\ \Delta\boldsymbol{W}^\top & \Delta\boldsymbol{E}\end{bmatrix}\right\| \leq \epsilon(\delta)} \left\|-2\left(\boldsymbol{W}^\star - \Delta\boldsymbol{W}\right) + 2\boldsymbol{\theta}^\star\left(\boldsymbol{E}^\star - \Delta\boldsymbol{E}\right)^\top + \epsilon(\delta)\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\right\| && (39) \\
&= \sup_{\left\|\begin{bmatrix}\Delta\boldsymbol{Q} & \Delta\boldsymbol{W} \\ \Delta\boldsymbol{W}^\top & \Delta\boldsymbol{E}\end{bmatrix}\right\| \leq \epsilon(\delta)} \left\|2\Delta\boldsymbol{W} - 2\boldsymbol{\theta}^\star\Delta\boldsymbol{E}^\top + \epsilon(\delta)\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\right\| && (40) \\
&\leq \sup_{\left\|\begin{bmatrix}\Delta\boldsymbol{Q} & \Delta\boldsymbol{W} \\ \Delta\boldsymbol{W}^\top & \Delta\boldsymbol{E}\end{bmatrix}\right\| \leq \epsilon(\delta)} 2\left\|\Delta\boldsymbol{W}\right\| + 2\|\boldsymbol{\theta}^\star\|\|\Delta\boldsymbol{E}\| + \epsilon(\delta)\|\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\| \\
&\leq 2\epsilon(\delta) + 2\|\boldsymbol{\theta}^\star\|\epsilon(\delta) + \epsilon(\delta)\|\nabla_{\boldsymbol{\theta}}\|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\|
\end{aligned}$$

$$= \epsilon(\delta)(2 + 2\|\boldsymbol{\theta}^\star\| + \left\|\nabla_{\boldsymbol{\theta}} \|\boldsymbol{G}(\boldsymbol{\theta}^\star)\|_q\right\|) \tag{41}$$

The first inequality (39) holds due to our guarantee in Proposition 1. In the next equality (40), we cancel out the terms $\boldsymbol{W}^\star$ and $\boldsymbol{E}^\star$ using the optimality condition for the true LSE problem, namely, $\nabla_{\boldsymbol{\theta}} f\left(\boldsymbol{\theta}^\star\right) = \boldsymbol{0} \Rightarrow \boldsymbol{W}^\star = \boldsymbol{\theta}^\star \boldsymbol{E}^{\star\top}$. Combining (36) and (37) ($\|\nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}^\star)\|$ in (37) replaced by (41)) using union bound yields the claim. □

5 Numerical Experiments

In this section, we conduct numerical experiments to assess the performance of our proposed method. The proposed method and the benchmarks presented are implemented in Python 3.7. Specifically, the optimization problem (7) is modeled using the CVXPY [19] interface and solved with the commercial solver MOSEK [20] on a laptop with a 6-core, 2.3 GHz Intel Core i7 CPU and 16 GB of RAM

We compare our robust LSE with the standard (i.e., unregularized) LSE for system identification tasks. Additionally, we consider an adaptive control task where we combine our robust LSE with the existing adaptive linear quadratic (LQ) control algorithms. We then compare the regret of different algorithms to demonstrate how improved performance in system identification can be translated into better regret.

Instead of focusing on a system from a particular example, we randomly generated four different sets of 500 synthetic stable systems $\boldsymbol{\theta}^\star = [\boldsymbol{A}^\star \ \boldsymbol{B}^\star] \in \mathbb{R}^{5\times 10}$, each set having the same spectral radius $\rho(\boldsymbol{A}^\star)$ ranging from 0.1 to 0.8. We compared the system identification errors of the robust LSE and the standard LSE as we collected more samples over time. We observed that the robust LSE showed greater performance improvement with smaller $\rho(\boldsymbol{A}^\star)$. Figure 1 shows the mean system identification errors in view of the operator norm for the set of 500 synthetic systems with $\rho(\boldsymbol{A}^\star) = 0.8$.

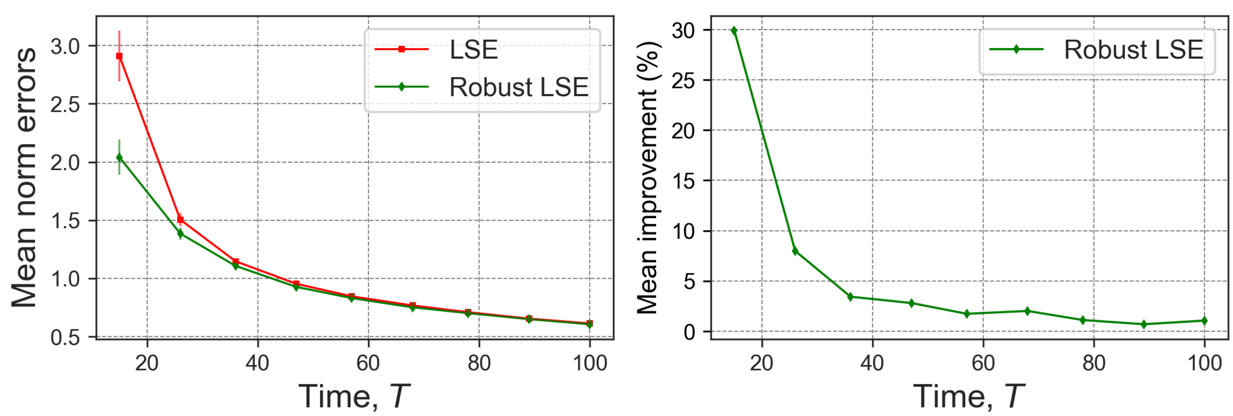

Figure 1. Mean system identification errors over 500 synthetic systems with $\rho(\boldsymbol{A}^\star) = 0.8$: mean errors in the operator norm, with the first and third quartiles represented by a bar (left) and percentage improvement of the robust LSE over the standard LSE (right).

As commented earlier, while the rate $\widetilde{\mathcal{O}}(1/\sqrt{T})$ derived in Proposition 1 is still useful, choosing the regularization parameter directly from the theoretical upper bound $\epsilon(\delta)$ leads to a too conservative estimate $\widehat{\boldsymbol{\theta}}_T$. In fact, a similar argument is made in [10]. Instead of adopting the theoretical guarantee, the authors use the standard bootstrap method to obtain an empirical upper bound $\bar{\epsilon}(\delta)$ on the system identification errors of the standard LSE, i.e., $\|\boldsymbol{\theta}^\star - \overline{\boldsymbol{\theta}}_T\| \leq \bar{\epsilon}(\delta)$. For the robust LSE, we use a 3-fold cross-validation procedure to determine the regularization parameter, as follows. We split the samples into three equal-sized subsets where two of the three subsets are put together to learn the robust estimate. The resulting estimate is then tested on the remaining set for all $\epsilon = (a \cdot 10^b)/\sqrt{T}$ where $a \in \{1, 3, 5, 7, 9\}$ and $b \in \{-3, \ldots, 3\}$. This process is repeated three times for different partitions of the samples to choose the ϵ that performs best overall.

To showcase how our robust LSE can be used in practice, we performed adaptive LQ control tasks using standard examples in the literature: i) a marginally unstable Laplacian system from [10], ii) the longitudinal flight control of Boeing 747 from [21], and iii) UAV in a 2D plane from [22]. We consider several adaptive LQ algorithms proposed in recent years: **1) OFULQ** from [3], **2) STABL** from [23], **3) ARBMLE** from [24]. Broadly speaking, these algorithms conduct two main tasks: identifying the system and deriving the best control input. In particular, OFULQ and STABL utilize the standard LSE for their system identification task. Hence, we can replace the standard LSE with the robust LSE which we referred to as **4) R-OFULQ** and **5) R-STABL**. Moreover, we can vary our adaptive control algorithms by choosing different values of parameter q as in Schatten q-norm of the SDP (7).

For each of the algorithms **1)-5)**, we ran 500 simulations over time horizon $T = 1000$ and recorded the mean regrets in Figure 2. Every algorithm presented in our experiments requires several parameters. We adopted the parameter setups suggested by the corresponding papers. However, we acknowledge that their setups are not identical to each other. For example, some papers start recording regret after $t = 50$, while others assume a tight upper bound on $\|\boldsymbol{\theta}^\star - \overline{\boldsymbol{\theta}}_t\|$ is available at each time step t. Irrespective of the choice of the Schatten norm parameter q, our algorithms offer significant advantages over other benchmark algorithms. The results demonstrate not only that the robust LSE can be utilized for various adaptive control algorithms, but also that optimizing the regularization parameter in real-time (i.e., with respect to T) for both the robust LSE and the regularized LSE is indeed advantageous.

6 Concluding Remarks

We have presented a general framework for robust system identification by leveraging robust optimization methodology to immunize the standard LSE against small sample estimation errors. We derive finite-sample guarantees on the system identification errors by analyzing the concentration of a single sample trajectory of states. Surprisingly, robustifying the estimation comes at no cost as the error rate matches the best one from non-robust LSE. While our analysis is based on a single trajectory of states, we emphasize that our

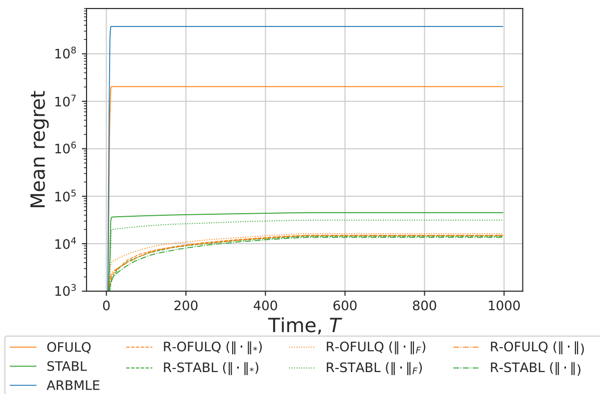

(a) Marginally unstable Laplacian system

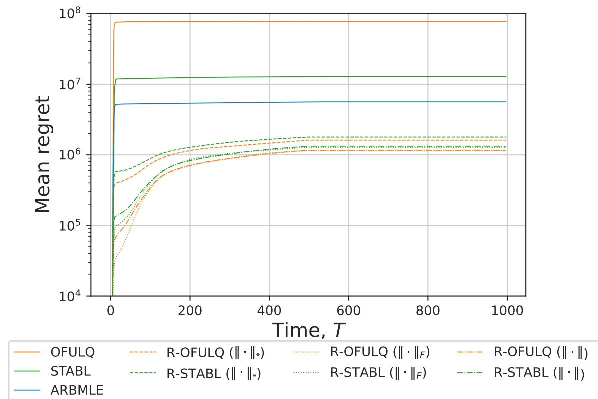

(b) Boeing 747

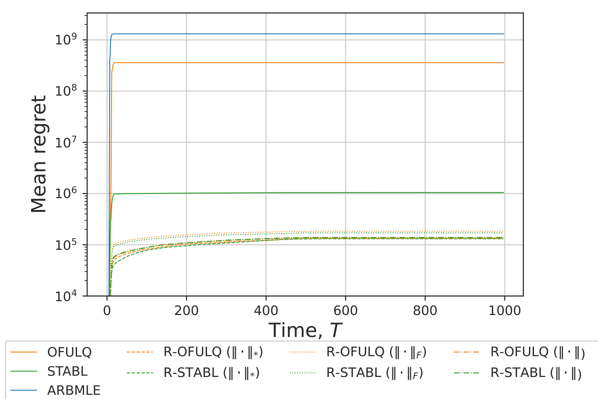

(c) UAV

Figure 2. Mean regret over 500 replications

framework can be applied straightforwardly to the simpler setting where multiple trajectories are available.

Our proposed formulation constitutes a simple semidefinite program, which is easy to implement using standard off-the-shelf solvers. In the special case when the ∞-norm is used in the uncertainty set, the formulation reduces to an efficiently solvable quadratic program. The experimental results on standard examples showcase the significant advantage of our robust model as it achieved unprecedented performance, particularly when the sample size is small. When further deployed in adaptive LQ control algorithms, the robust system estimates yield substantially lower regret than the non-robust estimates, demonstrating the practical advantage of our scheme.

In terms of limitations, our current work focuses on fully observable systems. So, future work will concentrate on developing a robust optimization framework to identify non-observable systems with performance guarantees. Another important future direction is to extend the framework to identify nonlinear systems, e.g., bilinear systems, nonlinear systems with linear parametrization, etc.

References

[1] Yassir Jedra and Alexandre Proutiere. Finite-time identification of stable linear systems: Optimality of the least-squares estimator. *arXiv preprint arXiv:2003.07937*, 2020.

[2] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.

[3] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26. JMLR Workshop and Conference Proceedings, 2011.

[4] Yue Sun, Samet Oymak, and Maryam Fazel. System identification via nuclear norm regularization. *arXiv preprint arXiv:2203.16673*, 2022.

[5] Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*, pages 5610–5618. PMLR, 2019.

[6] Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite-time adaptive stabilization of linear systems. *IEEE Transactions on Automatic Control*, 64(8):3498–3505, 2018.

[7] Dylan Foster, Tuhin Sarkar, and Alexander Rakhlin. Learning nonlinear dynamical systems from a single trajectory. In *Learning for Dynamics and Control*, pages 851–861. PMLR, 2020.

[8] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. *Advances in Neural Information Processing Systems*, 32, 2019.

[9] Yingying Li, Tianpeng Zhang, Subhro Das, Jeff Shamma, and Na Li. Non-asymptotic system identification for linear systems with nonlinear policies. *IFAC-PapersOnLine*, 56(2):1672–1679, 2023.

[10] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020.

[11] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009.

[12] Laurent El Ghaoui and Hervé Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on matrix analysis and applications*, 18(4):1035–1064, 1997.

[13] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Online least squares estimation with self-normalized processes: An application to bandit problems. *arXiv preprint arXiv:1102.2670*, 2011.

[14] David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.

[15] Areesh Mittal, Can Gokalp, and Grani A Hanasusanto. Robust quadratic programming with mixed-integer uncertainty. *INFORMS Journal on Computing*, 32(2):201–218, 2020.

[16] David Lee Hanson and Farroll Tim Wright. A bound on tail probabilities for quadratic forms in independent random variables. *The Annals of Mathematical Statistics*, 42(3):1079–1083, 1971.

[17] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.

[18] Anastasios Tsiamis, Ingvar Ziemann, Nikolai Matni, and George J Pappas. Statistical learning theory for control: A finite-sample perspective. *IEEE Control Systems Magazine*, 43(6):67–97, 2023.

[19] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 2016. To appear.

[20] MOSEK ApS. *The MOSEK optimization toolbox for Python manual. Version 10.2.0*, 2024.

[21] Tadashi Ishihara, Hai-Jiao Guo, and Hiroshi Takeda. A design of discrete-time integral controllers with computation delays via loop transfer recovery. *Automatica*, 28(3):599–603, 1992.

[22] Feiran Zhao, Keyou You, and Tamer Başar. Infinite-horizon risk-constrained linear quadratic regulator with average cost. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 390–395. IEEE, 2021.

[23] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Animashree Anandkumar. Reinforcement learning with fast stabilization in linear dynamical systems. In *International Conference on Artificial Intelligence and Statistics*, pages 5354–5390. PMLR, 2022.

[24] Akshay Mete, Rahul Singh, and PR Kumar. Augmented rbmle-ucb approach for adaptive control of linear quadratic systems. *Advances in Neural Information Processing Systems*, 35:9302–9314, 2022.