Proceedings of Machine Learning Research vol 144:1–13, 2021

Learning Partially Observed Linear Dynamical Systems from Logarithmic Number of Samples

Salar Fattahi FATTAHI@UMICH.EDU
University of Michigan, Ann Arbor, MI

Abstract

In this work, we study the problem of learning partially observed linear dynamical systems from a single sample trajectory. A major practical challenge in the existing system identification methods is the undesirable dependency of their required sample size on the system dimension: roughly speaking, they presume and rely on sample sizes that scale linearly with the system dimension. Evidently, in high-dimensional regime where the system dimension is large, it may be costly, if not impossible, to collect as many samples from the unknown system. In this paper, we introduce an regularized estimator that can accurately estimate the Markov parameters of the system, provided that the number of samples scale poly-logarithmically with the system dimension. Our result significantly improves the sample complexity of learning partially observed linear dynamical systems: it shows that the Markov parameters of the system can be learned in the high-dimensional setting, where the number of samples is significantly smaller than the system dimension.

Keywords: Learning, System Identification, High-Dimensional Statistics

1. Introduction

Most of today's real-world systems are characterized by being large-scale, complex, and safety-critical. For instance, the nation-wide power grid is comprised of millions of active devices that interact according to uncertain dynamics and complex laws of physics (Blaabjerg et al., 2006; Amin and Wollenberg, 2005). As another example, the contemporary transportation systems are moving towards a spatially distributed, autonomous, and intelligent infrastructure with thousands of heterogeneous and dynamic components (Barbaresso et al., 2014; Krechmer et al., 2018). Other examples include aerospace systems (Kapila et al., 2000), decentralized wireless networks (Kubisch et al., 2003), and multi-agent robot networks (Nguyen-Tuong and Peters, 2011). A common feature of these systems is that they are comprised of massive networks of interconnected subsystems with complex and uncertain dynamics.

The unknown structure of the dynamics on the one hand, and the emergence of machine learning and reinforcement learning (RL) as powerful tools for solving sequential decision making problems on the other hand, strongly motivate the use of *data-driven* methods in the operation of unknown safety-critical systems. However, the applications of machine learning techniques in the safety-critical systems remain mostly limited due to fundamental challenges. First, to alleviate the so-called "curse of dimensionality" in these systems, any practical learning and control method must be data-, time-, and memory-efficient. Second, rather than being treated as "black-box" models, these systems must be governed via models that are interpretable by practitioners, and are amenable to well-established robust/optimal control methods.

With the goal of addressing the aforementioned challenges, this paper studies the efficient learning of partially observed linear systems from a single trajectory of input-output measurements. Despite a mature body of literature on the statistical learning and control of linear dynamical systems, their practicality remains limited for large-scale and safety-critical systems. A key challenge lies in the required sample sizes of these methods and their dependency on the system dimensions: for a system with dimension n, the best existing system identification techniques require sample sizes in the order of $\mathcal{O}(n)$ to $\mathcal{O}(n^4)$ to provide certifiable guarantees on their performance (Oymak and Ozay, 2019; Krauth et al., 2019; Dean et al., 2019a; Simchowitz et al., 2019; Sarkar et al., 2019). Such dependency may inevitably lead to exceedingly long interactions with the safety-critical system, where it is extremely costly or even impossible to collect nearly as many samples without jeopardizing its safety—consider sampling from a geographically distributed power grid with tens of millions of parameters, and this increasing difficulty becomes apparent.

Contributions: In this work, we show that the Markov parameters defining the input-output behavior of partially observed linear dynamical systems can be learned with a poly-logarithmic sample complexity, i.e., from a single sample trajectory whose length scales poly-logarithmically with the output dimension. Our result relies on the key assumption that the system is inherently stable, or alternatively, it is equipped with an initial stabilizing controller. We show that the inherent stability of the system is analogous to the notion of *weak sparsity* in the corresponding Markov parameters. We then show that this "prior knowledge" on the weak sparsity of the Markov parameters can be systematically captured and exploited via the so-called ℓ_1 regularization. Our results imply that the Markov parameters of a partially observed linear system can be learned with certifiable bounds in the high-dimensional settings, where the system dimension is significantly larger than the number of available samples, thereby paving the way towards the efficient learning of massive-scale safety-critical systems. Within the realm of statistics, the ℓ_1-regularized estimators have been traditionally used to promote (exact) sparsity in the unknown parameters. In this work, we show that a similar ℓ_1-regularized method can be used to estimate the Markov parameters of the system, *irrespective* of the true sparsity of the unknown system.

Paper organization: In Section 2, we provide a literature review on different system identification techniques, and explain their connection to our work. The problem is formally defined in Section 3, and the main results are presented in Section 4. We provide an empirical study of our method in Section 5, and end with conclusions in Section 6. To streamline the presentation, the proofs are deferred to the extended version of the paper (Fattahi, 2020).

Notation: Upper- and lower-case letters are used to denote matrices and vectors, respectively. For a matrix $M \in \mathcal{R}^{m\times n}$ the symbols $M_{:j}$ and $M_{j:}$ indicate the j^{th} column and row of M, respectively. Given a vector v and an index set $\mathcal{S}$, the notation $v_{\mathcal{S}}$ refers to a subvector of v whose indices are restricted to the set $\mathcal{S}$. For a vector v, $\|v\|_p$ corresponds to its ℓ_p-norm. For a matrix M, the notation $\|M\|_{p,q}$ is equivalent to $\| \begin{bmatrix} \|M_{1:}\|_p & \|M_{2:}\|_p & \dots & M_{m:} \end{bmatrix} \|_q$. Moreover, $\|M\|_q$ refers to the induced q-norm of the matrix M. The notation $\|M\|_F$ is used to denote the Frobenius norm, defined as $\|M\|_{2,2}$. Furthermore, $\rho(M)$ correspond to the spectral radius of M. Given the sequences $f(n)$ and $g(n)$ indexed by n, the notation $f(n) = \mathcal{O}(g(n))$ or $f(n) \lesssim g(n)$ implies that there exists a universal constant $C < \infty$, independent of n, that satisfies $f(n) \leq Cg(n)$. Moreover, $f(n) = \tilde{\mathcal{O}}(g(n))$ is used to denote $f(n) = \mathcal{O}(g(n))$, modulo logarithmic factors. Similarly, the notation $f(n) \asymp g(n)$ implies that there exist constants $C_1 > 0$ and $C_2 < \infty$, independent of n, that satisfy $C_1 g(n) \leq f(n) \leq C_2 g(n)$. Given two scalars a and b, the notation $a \vee b$ denotes

their maximum. We use $x \sim \mathcal{N}(\mu, \Sigma)$ to show that x is a multivariate random variable drawn from a Gaussian distribution with mean μ and covariance Σ. For two random variables x and y, the notation $x \sim y$ implies that they have the same distribution. $\mathbb{E}[x]$ denotes the expected value of the random variable x. For an event $\mathcal{X}$, the notation $\mathbb{P}(\mathcal{X})$ refers to its probability of occurrence. The scalar c denotes a universal constant throughout the paper.

2. Related Works

System identification: Estimating system models from repeated experiments has a well-developed theory dating back to the 1960s, particularly in the case of linear and time-invariant systems. Standard reference textbooks on the topic include (Åström and Eykhoff, 1971; Ljung, 1999; Chen and Guo, 2012; Goodwin and Payne, 1977), all focusing on establishing *asymptotic* consistency of the proposed estimators. On the other hand, contemporary results in statistical learning of dynamical systems seek to characterize *finite time and finite data* rates. For fully observed systems, (Dean et al., 2017) shows that a simple least-squares estimator can correctly recover the system matrices with multiple trajectories whose length scales linearly with the system dimension. This result was later generalized to fully observed systems with only a single sample trajectory, with sample complexities depending polynomially on the system dimension (Dean et al., 2018; Simchowitz et al., 2018; Dean et al., 2019a; Sarkar and Rakhlin, 2018). These results were later extended to partially observed stable (Oymak and Ozay, 2019; Sarkar et al., 2019; Tsiamis and Pappas, 2019; Simchowitz et al., 2019), and unstable (Zheng and Li, 2020) systems where it is shown that the system matrices can be learned with similar polynomial sample complexities.

Regularized estimation: To further reduce the sample complexity of the system identification, a recent line of works has focused on learning dynamical systems with prior information. The works by Fattahi et al. (2019); Fattahi and Sojoudi (2018a,b,c) employ ℓ_1- and ℓ_1/ℓ_∞-regularized estimators to learn fully observed sparse systems with sample complexities that scale polynomially with the number nonzero entries in different rows and columns of the system matrices, but only logarithmically with the system dimension. However, these methods are not applicable to partially observed systems with hidden states. Another line of works introduces a different regularization technique, where the nuclear norm of the Hankel matrix is minimized to learn inherently low-order systems (Sun et al., 2020; Wahlberg and Rojas, 2013; Cai et al., 2016).

Learning-based control: Complementary to the aforementioned results, a large body of works study adaptive (Dean et al., 2018; Abbasi-Yadkori and Szepesvári, 2011; Abbasi-Yadkori et al., 2019; Lale et al., 2020), robust (Dean et al., 2019a,b; Mania et al., 2019), or distributed (Fattahi et al., 2020; Furieri et al., 2020) control of unknown linear systems. These works, culminated under the umbrella of model-based RL, indicate that if a learned model is to be integrated into a safety-critical control loop, then the uncertainty associated with the learned model must be explicitly quantified, thereby pinpointing the importance of a reliable system identification technique.

3. Problem Statement

Consider the following linear time-invariant (LTI) dynamical system:

$$x_{t+1} = Ax_t + Bu_t + w_t \tag{1}$$

$$y_t = Cx_t + Du_t + v_t \tag{2}$$

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^p$, and $y_t \in \mathbb{R}^m$ are the state, input, and output of the system at time t. Moreover, the vectors $w_t \in \mathbb{R}^n$ and $v_t \in \mathbb{R}^m$ are the process (or disturbance) and measurement noises, respectively. Throughout the paper, we assume that both v_t and w_t have element-wise independent sub-Gaussian distributions with parameters σ_w and σ_v, respectively. Moreover, without loss of generality, we assume that $x_0 = 0$[1]. The parameters $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times p}$, $C \in \mathbb{R}^{m\times n}$, and $D \in \mathbb{R}^{m\times p}$ are the unknown system matrices, to be estimated from a single input-output sample trajectory $\{(u_t, y_t)\}_{t=0}^{N}$. Much of the progress on the system identification is devoted to learning different variants of *fully observed* systems, where $C = I$ and $v_t = 0$. While being theoretically important, the practicality of these results is limited, since the state in realistic dynamical systems is not directly observed, or it is corrupted with measurement noise.

On the other hand, the lack of "intermediate" states x_t in partially observed systems gives rise to a mapping from u_k to y_k that is highly nonlinear in terms of the system parameters:

$$y_t = \underbrace{Du_t + \sum_{\tau=1}^{T-1} CA^{\tau-1}Bu_{t-\tau}}_{\text{Effect of the last } T \text{ inputs}} + \underbrace{\sum_{\tau=1}^{T-1} CA^{\tau-1}w_{t-\tau} + v_t}_{\text{Effect of noise}} + \underbrace{CA^{T-1}x_{t-T+1}}_{\text{Effect of the state at time } t-T+1}, \tag{3}$$

where $t \geq T-1$. The first term in (3) captures the effect of the past T inputs on y_t, while the second term corresponds to the effect of the unknown disturbance and measurement noises on y_t. Finally, the third term controls the contribution of the unknown state x_{t-T+1} on y_t, whose effect diminishes exponentially fast with T, provided that A is stable. A closer look at the first term reveals that the relationship between y_t and $\{u_t, u_{t-1}, \ldots, u_{t-T+1}\}$ becomes linear in terms of the *Markov matrix*

$$G = \begin{bmatrix} D & G_0 & G_1 \ldots G_{T-2} \end{bmatrix} = \begin{bmatrix} D & CB & CAB \ldots CA^{T-2}B \end{bmatrix} \in \mathbb{R}^{m\times Tp}, \tag{4}$$

whose components are commonly known as Markov parameters of the system. One of the main goals of this paper is to obtain an accurate estimate of G given a single input-output trajectory. The Markov parameters can be used to directly estimate the outputs of the system from the past input. Moreover, as will be shown later, a good estimation of the Markov parameters can be translated into an accurate estimate of the Hankel matrix, which in turn can be used in the model reduction and $\mathcal{H}_\infty$ methods in control theory (Antoulas, 2005; Zhou and Doyle, 1998). Finally, given the estimated Markov matrix G, one can estimate the system matrices. Note that it is only possible to extract the system parameters up to a nonsingular transformation: given any nonsingular matrix S, the system matrices (A, B, C, D) and $(S^{-1}AS, SB, CS^{-1}, D)$ correspond to the same Markov matrix. Therefore, a common approach is to first construct the associated Hankel matrix, and then extract a *realization* of the system parameters from the Hankel matrix using the *Ho-Kalman* method (Ho and Kálmán, 1966; Ljung, 1999). In fact, it has been recently shown in (Oymak and Ozay, 2019; Sarkar et al., 2019) that the Ho-Kalman method can robustly obtain a balanced realization of the system matrices, provided that the estimated Markov matrix enjoys a small estimation error.

Proposition 1 (Oymak and Ozay (2019), informal) *Suppose that the true system is controllable and observable. Given an estimate $\widehat{G}$ of G, the Ho-Kalman method outputs system matrices $(\widehat{A}, \widehat{B},$*

1. Our results can be readily extended to scenarios where x_0 is randomly drawn from a sub-Gaussian distribution.

$\widehat{C}, \widehat{D}$) *that satisfy*

$$\|B - \mathcal{U}\widehat{B}\|_F \lesssim \sqrt{T}\|G - \widehat{G}\|_F \tag{5}$$

$$\|C - \widehat{C}\mathcal{U}^\top\|_F \lesssim \sqrt{T}\|G - \widehat{G}\|_F \tag{6}$$

$$\|A - \mathcal{U}\widehat{A}\mathcal{U}^\top\|_F \lesssim T\|G\|_2\|G - \widehat{G}\|_F \tag{7}$$

for some unitary matrix $\mathcal{U}$, provided that $\widehat{G}$ is sufficiently close to G.

Therefore, without loss of generality, our focus will be devoted to obtaining accurate estimates of the Markov and Hankel matrices. To streamline the presentation, the concatenated input and process noise vectors are defined as:

$$\bar{u}_t = \begin{bmatrix} u_t^\top & u_{t-1}^\top & \dots & u_{t-T+1}^\top \end{bmatrix}^\top \in \mathbb{R}^{Tp}, \quad \bar{w}_t = \begin{bmatrix} w_t^\top & w_{t-1}^\top & \dots & w_{t-T+1}^\top \end{bmatrix}^\top \in \mathbb{R}^{Tn}, \tag{8}$$

Moreover, the following concatenated matrix will be used throughout the paper:

$$F = \begin{bmatrix} 0 & C & CA & \dots & CA^{T-2} \end{bmatrix} \in \mathbb{R}^{m \times Tn} \tag{9}$$

Based on the above definitions, the input-output relation (3) can be written compactly as $y_t = G\bar{u}_t + F\bar{w}_t + e_t + v_t$ where $e_t = CA^{T-1}x_{t-T+1}$. To estimate the Markov matrix G, the work by Oymak and Ozay (2019) proposes the following least-squares estimator:

$$\widetilde{G} = \arg\min_X \sum_{t=T-1}^{N+T-2} \|y_t - X\bar{u}_t\|_2^2 \tag{10}$$

Define $q = p + n + m$ as the system dimension, and σ_e^2 as the effective variance of e_t, as in

$$\sigma_e = \Phi(A)\|CA^{T-1}\|\sqrt{\frac{T\|\Gamma_\infty\|}{1 - \rho(A)^{2T}}} \tag{11}$$

where

$$\Phi(A) = \sup_{\tau \geq 0} \frac{\|A^\tau\|}{\rho(A)^\tau}, \quad \Gamma_\infty = \sum_{i=0}^{\infty} \sigma_w^2 A^i (A^\top)^i + \sigma_u^2 A^i B B^\top (A^\top)^i \tag{12}$$

Oymak and Ozay (2019) characterizes the non-asymptotic behavior of this least-squares estimator.

Theorem 2 (Oymak and Ozay (2019)) *Suppose that $u_t \sim \mathcal{N}(0, \sigma_u^2 I)$ for every $t = 0, \dots, T + N - 2$, and $N \gtrsim Tq\log^2(Tq)\log^2(Nq)$. Then, with overwhelming probability, the following inequalities hold:*

$$\|\widetilde{G} - G\|_2 \lesssim \frac{\sigma_v + \sigma_e + \sigma_w\|F\|_2}{\sigma_u}\sqrt{\frac{Tq\log^2(Tq)\log^2(Nq)}{N}}, \tag{13}$$

$$\|\widetilde{G} - G\|_F \lesssim \frac{(\sigma_v + \sigma_e)\sqrt{m} + \sigma_w\|F\|_2}{\sigma_u}\sqrt{\frac{Tq\log^2(Tq)\log^2(Nq)}{N}} \tag{14}$$

where σ_u^2, σ_v^2, σ_w^2 are the variances of the random input, disturbance noise, and the measurement noise, respectively.

The above theorem shows that the spectral norm of the estimation error for the Markov parameters via least-squares method is in the order of $\tilde{\mathcal{O}}\left(\sqrt{T(n+m+p)/N}\right)$, provided that $N = \tilde{\mathcal{O}}(T(n+m+p))$. Moreover, Oymak and Ozay (2019) shows that the number of samples N can be reduced to $\tilde{\mathcal{O}}(Tp)$ (without improving the spectral norm error). Such dependency on the system dimension is unavoidable if one does not exploit any prior information on the structure of G. In particular, it is easy to see that the Markov parameter G has Tmp unknown parameters, and hence, one needs to collect at least Tp outputs (each with size m) to obtain a well-defined least-squares estimator. Evidently, such dependency on the system dimension may be prohibitive for large-scale and safety-critical systems, where it is expensive to collect as many output samples. Motivated by this shortcoming of the existing methods, we aim to address the following open question:

Question: *Can partially observed linear systems be learned with a sublinear sample complexity?*

4. Main Results

In this section, we provide an affirmative answer to the aforementioned question. At a high-level, we will use the fact that, due to the stability of A, the Markov parameters decay *exponentially fast*, which in turn implies that the rows of the extended matrix G exhibit a bounded ℓ_1-norm (also known as weak sparsity (Wainwright, 2019)). This observation strongly motivates the use of the following regularized estimator:

$$\widehat{G} = \arg\min_{X} \left(\frac{1}{2N} \sum_{t=T-1}^{N+T-2} \|y_t - X\bar{u}_t\|_2^2 \right) + \lambda \|X\|_{1,1} \tag{15}$$

Due to the stability of A, there exist scalars $C_{\text{sys}} \geq 1$ and $\rho < 1$ such that $\|A^\tau\|_1 \leq C_{\text{sys}}\rho^\tau$. Without loss of generality and to simplify the notation, we assume that $\max\{\|B\|_1, \|C\|_1, \|D\|_1\} \leq C_{\text{sys}}$. Finally, define the *effective variance* of the disturbance noise as $\bar{\sigma}_w = \left(\frac{C_{\text{sys}}^2}{1-\rho}\right)\sigma_w$. The main result of the paper is the following theorem.

Theorem 3 *Suppose that $u_t \sim \mathcal{N}(0, \sigma_u^2 I)$ for every $t = 0, \ldots, T+N-2$. Moreover, suppose that N and T satisfy the following inequalities:*

$$N \gtrsim \log^2(Tp), \qquad T \gtrsim T_0 = \frac{\log\log(Nn+Tp) + \log\left(\frac{C_{\text{sys}}}{1-\rho}\right) + \log(\sigma_w + \sigma_v) + \log\left(\frac{1}{\epsilon}\right)}{1-\rho} \tag{16}$$

for an arbitrary $\epsilon > 0$. Finally, assume that λ is chosen such that

$$\lambda \asymp \sigma_u \left(\bar{\sigma}_w + \sigma_v\right) \sqrt{\frac{\log(Tpn)}{N}} + \epsilon \tag{17}$$

Then, with overwhelming probability, the following inequalities hold:

$$\|G - \widehat{G}\|_{2,\infty} \lesssim \mathcal{E}_1 \vee \mathcal{E}_2, \quad \|G - \widehat{G}\|_F \lesssim \sqrt{m}(\mathcal{E}_1 \vee \mathcal{E}_2) \tag{18}$$

where

$$\mathcal{E}_1 = \sqrt{\frac{C_{\text{sys}}^3}{1-\rho}} \left(\sqrt{\frac{\bar{\sigma}_w + \sigma_v}{\sigma_u^3}} \left(\frac{\log(Tpn)}{N}\right)^{1/4} + \frac{\epsilon}{\sigma_u^2} \right), \quad \mathcal{E}_2 = \frac{C_{\text{sys}}^3}{1-\rho} \left(\frac{\log(Tp)}{N}\right)^{1/4} \tag{19}$$

The above theorem can be used to provide estimation error bounds on the higher order Markov parameters and Hankel matrices (which can be used to recover a realization of the system parameters (A, B, C, D), as delineated in Proposition 1). Similar to Oymak and Ozay (2019), define the true and estimated K^{th} order (where $K \geq T$) Markov parameters as

$$G^{(K)} = \begin{bmatrix} D & CB & CAB & \dots & CA^{K-2}B \end{bmatrix} \in \mathbb{R}^{m \times Kp}, \quad \widehat{G}^{(K)} = \begin{bmatrix} \widehat{G} & \mathbf{0}_{m \times (K-T)p} \end{bmatrix} \in \mathbb{R}^{m \times Kp} \tag{20}$$

Moreover, define the true and estimated K^{th} order Hankel matrices as

$$H^{(K)} = \begin{bmatrix} D & CB & \dots & CA^{K-2}B \\ CB & CAB & \dots & CA^{K-1}B \\ & & \vdots & \\ CA^{K-2}B & CA^{K-1}B & \dots & CA^{2K-3}B \end{bmatrix} \in \mathbb{R}^{Km \times Kp}, \tag{21}$$

$$\widehat{H}^{(K)} = \begin{bmatrix} \widehat{D} & \widehat{G}_0 & \dots & \widehat{G}_{T-3} & \widehat{G}_{T-2} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} \\ \widehat{G}_0 & \widehat{G}_1 & \dots & \widehat{G}_{T-2} & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} \\ & & \vdots & & & & & \\ \widehat{G}_{T-2} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} \\ \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} \\ & & \vdots & & & & & \\ \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \mathbf{0}_{m\times n} & \dots & \mathbf{0}_{m\times n} \end{bmatrix} \in \mathbb{R}^{Km \times Kp} \tag{22}$$

Our next corollary provides estimation error bounds for the K^{th} order Markov and Hankel matrices.

Corollary 4 *Suppose that $u_t \sim \mathcal{N}(0, \sigma_u^2 I)$ for every $t = 0, \dots, T+N-2$, and N and λ satisfy (16) and (17), respectively. Moreover, assume that $T \gtrsim T_0 \vee (\log(\|C\|_\infty) + \log(1/\tilde{\epsilon}))/(1-\rho)$ for an arbitrary $\tilde{\epsilon} > 0$. Then, for any $K \geq T$ (including $K = \infty$), the following inequalities hold with overwhelming probability:*

$$\|G^{(K)} - \widehat{G}^{(K)}\|_{2,\infty} \lesssim \mathcal{E}_1 \vee \mathcal{E}_2 + \tilde{\epsilon}, \quad \|G^{(K)} - \widehat{G}^{(K)}\|_F \lesssim \sqrt{m}(\mathcal{E}_1 \vee \mathcal{E}_2 + \tilde{\epsilon}) \tag{23}$$

$$\|H^{(K)} - \widehat{H}^{(K)}\|_{2,\infty} \lesssim \mathcal{E}_1 \vee \mathcal{E}_2 + \tilde{\epsilon}, \quad \|H^{(K)} - \widehat{H}^{(K)}\|_F \lesssim \sqrt{Tm}(\mathcal{E}_1 \vee \mathcal{E}_2 + \tilde{\epsilon}) \tag{24}$$

Next, we will explain the implications of Theorem 3 and Corollary 4.

Sample complexity: According to Theorem 3 and Corollary 4, the required number of samples N for estimating the Markov parameters and the Hankel matrix scales poly-logarithmically with the system dimension, making it particularly well-suited to massive-scale dynamical systems, where the system dimension surpasses the number of available input-output samples. In contrast, the existing methods for learning partially observed linear systems do not provide any guarantee on the estimation error under such "high-dimension/low-sampling" regime. Moreover, the imposed lower bound on T scales double-logarithmically with respect to the system dimension[2], which can be treated as a constant number for all practical purposes.[3]

2. The imposed lower bound on T is to simplify the derived bounds, and hence, can be relaxed at the expense of less intuitive estimation bounds.
3. It is easy to verify that $\log\log(s) \leq 5$ for any $s \leq 10^{50}$!

Estimation error: The estimation error bounds in Theorem 3 and Corollary 4 are in terms of the row-wise ℓ_2 and Frobenius norms. In contrast, most of the existing methods provide upper bounds on the spectral norm of the estimation error. An important benefit of the provided row-wise bound is that it provides a finer control over the element-wise estimation error, which in turn can be used in the recovery of the special sparsity patterns in the Hankel matrices (Jin and Ye, 2017). We note that although the provided bound on the Frobenius norm of the estimation error can be readily applied to its spectral norm, we believe that it can be strengthened. Moreover, the provided estimation error bound reduces at the rate $N^{-1/4}$, which is slower than the rate $N^{-1/2}$ for the simple least-squares estimator (see Theorem 2). However, a more careful scrutiny of (18) and (14) reveals that our proposed estimator outperforms the least-squares even in a high-sampling regime where

$$\frac{N}{T^2} \lesssim \frac{q^2 \log^4(Tq) \log^4(Nq)}{\log(Tpn)} = \tilde{O}\left((n+m+p)^2\right) \tag{25}$$

In fact, a stronger statement can be made on the ratio between the estimation errors:

Corollary 5 *Denote the right hand sides of* (14) *and* (18) *as* $\mathcal{E}_F^{LS}$ *and* $\mathcal{E}_F^{\ell_1}$*, respectively. Suppose that* $\sigma_u \vee \sigma_w \vee \sigma_v \vee \frac{C_{\text{sys}}}{1-\rho} \vee \Phi(A) \vee \|B\|_2 \vee \|C\|_2 = \mathcal{O}(1)$*, and* $T \gtrsim T_0 + \log(n+m+p)$*. Then, we have* $\lim_{n,m,p\to\infty} \frac{\mathcal{E}_F^{\ell_1}}{\mathcal{E}_F^{LS}} = 0$ *provided that* T *and* N *satisfy* $\lim_{n,m,p\to\infty} \frac{N \log(Tpn)}{T^2(n+m+p)^2} = 0$

Method	Sample Complexity	Error Bound ($\|\cdot\|_F$)	Additional Notes
proposed method	$\mathcal{O}(\log^2(Tp))$	$\mathcal{O}\left(\sqrt{m}\left(\frac{\log(Tnp)}{N}\right)^{1/4}\right)$	Single trajectory
Oymak and Ozay (2019)	$\tilde{\mathcal{O}}(Tq)$	$\tilde{\mathcal{O}}\left(\sqrt{m}\left(\frac{Tq}{N}\right)^{1/2}\right)$	Single trajectory
Sarkar et al. (2019)	$\tilde{\mathcal{O}}(n^2)$	$\tilde{\mathcal{O}}\left(\sqrt{m}\left(\frac{pn^2}{N}\right)^{1/2}\right)$	Single trajectory, Suitable for systems with unknown order
Zheng and Li (2020)	$\tilde{O}(mT+q)$	$\tilde{\mathcal{O}}\left(\sqrt{m}\left(\frac{T^3q}{N}\right)^{1/2}\right)$	Multiple trajectories, Stable and unstable systems
Sun et al. (2020)	$\tilde{\mathcal{O}}(pR)$	$\tilde{\mathcal{O}}\left(\left(\frac{Rnp}{N}\right)^{1/2}\right)$	Multiple trajectories, MISO ($m=1$)
Tu et al. (2017)	$\tilde{\mathcal{O}}(r)$	$\tilde{\mathcal{O}}\left(\left(\frac{r}{T}\right)^{1/2}\right)$	Multiple trajectories, SISO ($p=m=1$)

Table 1: Sample complexity and error bounds on the estimated Markov parameters for different methods. The parameters $R \leq n$ and r are respectively the order of the system and the length of the FIR impulse response; see Sun et al. (2020) and Tu et al. (2017) for more information.

The above proposition implies that in the regime where N is not significantly larger than the system dimension, the derived upper bound on the estimation error of the regularized estimator becomes arbitrarily smaller than that of the least-squares method. Our numerical analysis also reveals the superior performance of the proposed estimator, even when $N \gg Tp$. Finally, we point out that similar error bounds have been derived for the classical linear regression with weakly sparse structures. In particular, Negahban et al. (2012) considers a "simpler" linear model where the samples/outputs are assumed to be independent, and shows that a ℓ_1-regularized estimator achieves an error bound in the order of $\mathcal{O}\left((\log(d)/N)^{1/4}\right)$, where d is the dimension of the unknown regression

vector. Theorem 3 reveals that the same rate can be achieved in the context of system identification with a single input-output trajectory, despite the fact that the samples are correlated. Table 1 compares the performance of the proposed estimator with other state-of-the-art methods.

Role of signal-to-noise ratio: Intuitively, the estimation error should improve with an increasing signal-to-noise (SNR) ratio (in our problem, the SNR ratio is defined as $\sigma_u^3/(\sigma_w+\sigma_v)$); this behavior is also observed in the related works (Oymak and Ozay, 2019; Sun et al., 2020; Tu et al., 2017). In contrast, our provided bound is the maximum of two terms, one of which is independent of the SNR ratio. In other words, an increasing SNR ratio can only shrink the estimation error down to a certain positive threshold. The reason behind this behavior lies in the statistical behavior of the random input matrix U, which is explained as follows. For two different vectors ζ and $\tilde{\zeta}$, the quantity $U(\zeta-\tilde{\zeta})$ measures how *distinguishable* these vectors are under the considered linear model. For the cases where $N \gtrsim Tp$, it is easy to see that these two vectors are easily distinguishable, since $\|U(\zeta-\tilde{\zeta})\|_2^2 \geq \kappa\sigma_u^2\|\zeta-\tilde{\zeta}\|_2^2$ holds with high probability, for some strictly positive κ (see, e.g., Oymak and Ozay (2019)). However, in the high-dimensional setting, the matrix U will inevitably have zero singular values, and $\|U(\zeta-\tilde{\zeta})\|_2^2 \geq \kappa\sigma_u^2\|\zeta-\tilde{\zeta}\|_2^2$ may no longer hold for specific choices of ζ and $\tilde{\zeta}$. Under such circumstances, we show that the *relaxed* inequality $\|U(\zeta-\tilde{\zeta})\|_2^2 \geq \kappa\sigma_u^2\|\zeta-\tilde{\zeta}\|_2^2 - \sigma_u^2 f(\zeta-\tilde{\zeta})$ holds for any ζ and $\tilde{\zeta}$, where $f(\cdot)$ is a function defined in (Fattahi, 2020). Upon replacing $\zeta-\tilde{\zeta}$ with $G_{i:}-\widehat{G}_{i:}$ for some row index i, it is easy to see that this lower bound becomes nontrivial only if $\|G_{i:}-\widehat{G}_{i:}\|_2^2 > f(G_{i:}-\widehat{G}_{i:})/\kappa$, which is *independent* of the SNR ratio. This intuition is formalized in the extended version of the paper (Fattahi, 2020). In particular, it is shown that: (1) the threshold $f(G_{i:}-\widehat{G}_{i:})/\kappa$ is small, i.e., it is upper bounded by $\mathcal{E}_2^2$; (2) whenever $\|G_{i:}-\widehat{G}_{i:}\|_2^2$ is larger than $\mathcal{E}_2^2$, it can be upper bounded by $\mathcal{E}_1^2$.

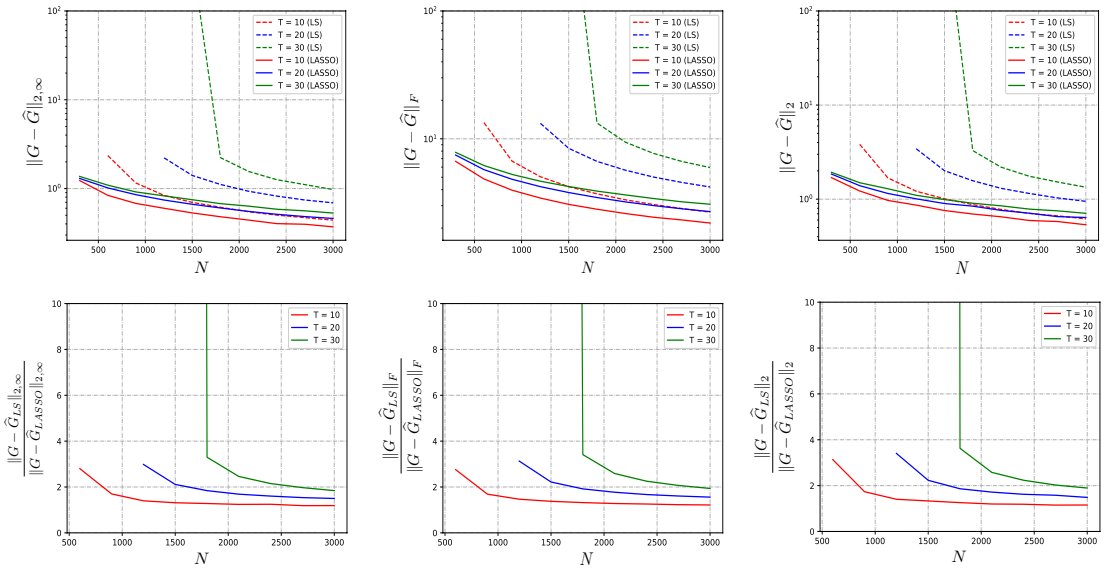

Figure 1: The estimation error of the Markov parameters for LASSO (denoted as $\widehat{G}_{LASSO}$) and LS (denoted as $\widehat{G}_{LS}$) with respect to the sample size, with $\sigma_w^2 = \sigma_v^2 = 0.1$ and varying T. When $N < Tp$, LASSO achieves small estimation error, while LS is not well-defined. Moreover, LASSO significantly outperforms LS when $N \geq Tp$. The y-axis in all figures are clipped to better illustrate the differences in the curves.

5. Simulations

In this section, we showcase the performance of the proposed regularized estimator. In particular, we will provide an empirical comparison between our method and the least-squares approach introduced by Oymak and Ozay (2019).[4] In all of our simulations, we set $n = 200$, $m = p = 50$, and $D = 0$. A is chosen as a banded matrix, with the bandwidth equal to 5. Moreover, each nonzero entry of A is selected uniformly from $[-0.5, 0.5]$. To ensure the stability of the system, A is normalized to ensure that $\rho(A) = 0.8$. The $(i, j)^{\text{th}}$ entry of B is set to 1 if $i = 4j$, and it is set to 0 otherwise, for every $(i, j) \in \{1, \dots, n\} \times \{1, \dots, p\}$. Finally, C is chosen as a Gaussian matrix, with entries selected from $\mathcal{N}(0, 1/m)$. Note that, despite the sparse nature of A and B, the Markov parameters of the system are fully dense, due to the dense nature of C. Throughout our simulations, σ_u is set of 1, and the values of σ_w and σ_v are changed to examine the effect of SNR ratio on the quality of our estimates. Moreover, in all of our simulations, we set the regularization parameter to

$$\lambda = 0.2(\sigma_w + \sigma_v)\sqrt{\frac{\log(Tpn)}{N}} + 0.02 \times 0.8^T \tag{26}$$

Note that the above choice of the regularization parameter does not require any further fine-tuning, and it is in line with Theorem 3, after replacing ϵ with 0.02×0.8^T in (17). The exponential decay in ϵ correctly captures the diminishing effect of the unknown initial state x_{t-T+1} on the output y_t with T (see (3)). We point out that a better choice of λ may be possible via cross-validation (Shao, 1993). Figure 1 shows the estimation error of the proposed method compared to the least-squares estimator (referred to as `LASSO` and `LS`, respectively) for $\sigma_w^2 = \sigma_v^2 = 0.1$ (averaged over 10 independent trials). It can be seen that `LASSO` significantly outperforms `LS` for all values of N and T. In the high-dimensional setting, where $N < Tp$, `LS` is not well-defined, while `LASSO` results in small estimation errors. Moreover, when $N \geq Tp$, the incurred estimation error of `LASSO` is 1.2 to 1077 times smaller than that of `LS`. Although the main strength of `LASSO` is in the high-dimensional regime, it still outperforms `LS` when $N \gg Tp$.

6. Conclusions

In this paper, we propose a method for learning partially observed linear systems from a single sample trajectory in high-dimensional settings, i.e., when the number of samples is less than the system dimension. Most of the existing inference methods presume and rely on the availability of prohibitively large number of samples collected from the unknown system. In this work, we address this issue by reducing the sample complexity of estimating the Markov parameters of partially observed systems via an ℓ_1-regularized estimator. We show that, when the system is inherently stable, the required number of samples for a reliable estimation of the Markov parameters scales poly-logarithmically with the dimension of the system.

Acknowledgments

We would like to thank Necmiye Ozay, Nikolai Matni, Yang Zheng, and Kamyar Azizzadenesheli for their insightful suggestions and constructive comments.

4. It has been recently verified in Sun et al. (2020) that the method proposed by Oymak and Ozay Oymak and Ozay (2019) outperforms that of Sarkar *et. al.* Sarkar et al. (2019). Therefore, we only focus on the former.

References

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, 2011.

Yasin Abbasi-Yadkori, Nevena Lazic, and Csaba Szepesvári. Model-free linear quadratic control via reduction to expert prediction. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3108–3117, 2019.

S Massoud Amin and Bruce F Wollenberg. Toward a smart grid: power delivery for the 21st century. *IEEE power and energy magazine*, 3(5):34–41, 2005.

Athanasios C Antoulas. *Approximation of large-scale dynamical systems*. SIAM, 2005.

Karl Johan Åström and Peter Eykhoff. System identification—a survey. *Automatica*, 7(2):123–162, 1971.

Jim Barbaresso, Gustave Cordahi, Dominie Garcia, Christopher Hill, Alex Jendzejec, Karissa Wright, and Booz Allen Hamilton. Usdot's intelligent transportation systems (its) its strategic plan, 2015-2019. Technical report, United States. Department of Transportation. Intelligent Transportation, 2014.

Frede Blaabjerg, Remus Teodorescu, Marco Liserre, and Adrian V Timbus. Overview of control and grid synchronization for distributed power generation systems. *IEEE Transactions on industrial electronics*, 53(5):1398–1409, 2006.

Jian-Feng Cai, Xiaobo Qu, Weiyu Xu, and Gui-Bo Ye. Robust recovery of complex exponential signals from random gaussian projections via low rank hankel matrix reconstruction. *Applied and computational harmonic analysis*, 41(2):470–490, 2016.

Han-Fu Chen and Lei Guo. *Identification and stochastic adaptive control*. Springer Science & Business Media, 2012.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019a.

Sarah Dean, Stephen Tu, Nikolai Matni, and Benjamin Recht. Safely learning to control the constrained linear quadratic regulator. In *2019 American Control Conference (ACC)*, pages 5582–5588. IEEE, 2019b.

Salar Fattahi. Learning partially observed linear dynamical systems from logarithmic number of samples. *arXiv preprint arXiv:2010.04015*, 2020.

Salar Fattahi and Somayeh Sojoudi. Data-driven sparse system identification. In *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 462–469. IEEE, 2018a.

Salar Fattahi and Somayeh Sojoudi. Non-asymptotic analysis of block-regularized regression problem. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 27–34. IEEE, 2018b.

Salar Fattahi and Somayeh Sojoudi. Sample complexity of sparse system identification problem. *arXiv preprint arXiv:1803.07753*, 2018c.

Salar Fattahi, Nikolai Matni, and Somayeh Sojoudi. Learning sparse dynamical systems from a single sample trajectory. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2682–2689. IEEE, 2019.

Salar Fattahi, Nikolai Matni, and Somayeh Sojoudi. Efficient learning of distributed linear-quadratic control policies. *SIAM Journal on Control and Optimization*, 58(5):2927–2951, 2020.

Luca Furieri, Yang Zheng, and Maryam Kamgarpour. Learning the globally optimal distributed lq regulator. In *Learning for Dynamics and Control*, pages 287–297, 2020.

Graham Clifford Goodwin and RL Payne. Dynamic system identification. experiment design and data analysis. 1977.

BL Ho and Rudolf E Kálmán. Effective construction of linear state-variable models from input/output functions. *at-Automatisierungstechnik*, 14(1-12):545–548, 1966.

Kyong Hwan Jin and Jong Chul Ye. Sparse and low-rank decomposition of a hankel structured matrix for impulse noise removal. *IEEE Transactions on Image Processing*, 27(3):1448–1461, 2017.

Vikram Kapila, Andrew G Sparks, James M Buffington, and Qiguo Yan. Spacecraft formation flying: Dynamics and control. *Journal of Guidance, Control, and Dynamics*, 23(3):561–564, 2000.

Karl Krauth, Stephen Tu, and Benjamin Recht. Finite-time analysis of approximate policy iteration for the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 8512–8522, 2019.

Daniel Krechmer, Erin Flanigan, Aldo Rivadeneyra, Katherine Blizzard, Sam Van Hecke, and Robert Rausch. Effects on intelligent transportation systems planning and deployment in a connected vehicle environment. Technical report, United States. Federal Highway Administration. Office of Operations, 2018.

Martin Kubisch, Holger Karl, Adam Wolisz, Lizhi Charlie Zhong, and Jan Rabaey. Distributed algorithms for transmission power control in wireless sensor networks. In *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003.*, volume 1, pages 558–563. IEEE, 2003.

Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Regret bound of adaptive control in linear quadratic gaussian (lqg) systems. *arXiv preprint arXiv:2003.05999*, 2020.

Lennart Ljung. System identification. *Wiley encyclopedia of electrical and electronics engineering*, pages 1–19, 1999.

Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalent control of lqr is efficient. *arXiv preprint arXiv:1902.07826*, 2019.

Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, Bin Yu, et al. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical science*, 27(4):538–557, 2012.

Duy Nguyen-Tuong and Jan Peters. Model learning for robot control: a survey. *Cognitive processing*, 12(4):319–340, 2011.

Samet Oymak and Necmiye Ozay. Non-asymptotic identification of lti systems from a single trajectory. In *2019 American Control Conference (ACC)*, pages 5655–5661. IEEE, 2019.

Tuhin Sarkar and Alexander Rakhlin. How fast can linear dynamical systems be learned? *arXiv preprint arXiv:1812.01251*, 2018.

Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. Finite-time system identification for partially observed lti systems of unknown order. *arXiv preprint arXiv:1902.01848*, 2019.

Jun Shao. Linear model selection by cross-validation. *Journal of the American statistical Association*, 88(422):486–494, 1993.

Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. *arXiv preprint arXiv:1802.08334*, 2018.

Max Simchowitz, Ross Boczar, and Benjamin Recht. Learning linear dynamical systems with semi-parametric least squares. *arXiv preprint arXiv:1902.00768*, 2019.

Yue Sun, Samet Oymak, and Maryam Fazel. Finite sample system identification: Improved rates and the role of regularization. 2020.

Anastasios Tsiamis and George J Pappas. Finite sample analysis of stochastic system identification. *arXiv preprint arXiv:1903.09122*, 2019.

Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht. Non-asymptotic analysis of robust control from coarse-grained identification. *arXiv preprint arXiv:1707.04791*, 2017.

Bo Wahlberg and Cristian Rojas. Matrix rank optimization problems in system identification via nuclear norm mimization. In *(Tutorial) at the European Control Conference*, 2013.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Yang Zheng and Na Li. Non-asymptotic identification of linear dynamical systems using multiple trajectories. *arXiv preprint arXiv:2009.00739*, 2020.

Kemin Zhou and John Comstock Doyle. *Essentials of robust control*, volume 104. Prentice hall Upper Saddle River, NJ, 1998.