

Ainesh Bakshi ainesh@mit.edu Massachusetts Institute of Technology Cambridge, Massachusetts, USA

Ankur Moitra moitra@mit.edu Massachusetts Institute of Technology Cambridge, Massachusetts, USA

### ABSTRACT

Linear dynamical systems are the foundational statistical model upon which control theory is built. Both the celebrated Kalman filter and the linear quadratic regulator require knowledge of the system dynamics to provide analytic guarantees. Naturally, learning the dynamics of a linear dynamical system from linear measurements has been intensively studied since Rudolph Kalman's pioneering work in the 1960's. Towards these ends, we provide the first polynomial time algorithm for learning a linear dynamical system from a polynomial length trajectory up to polynomial error in the system parameters under essentially minimal assumptions; observability, controllability, and marginal stability. Our algorithm is built on a method of moments estimator to directly estimate Markov parameters from which the dynamics can be extracted. Furthermore we provide statistical lower bounds when our observability and controllability assumptions are violated.

#### **CCS CONCEPTS**

• Theory of computation  $\rightarrow$  Machine learning theory.

#### **KEYWORDS**

linear dynamical systems, marginal stability, method of moments, information-theoretically minimal, minimal assumptions, observability, controllability

#### **ACM Reference Format:**

Ainesh Bakshi, Allen Liu, Ankur Moitra, and Morris Yau. 2023. A New Approach to Learning Linear Dynamical Systems. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3564246.3585247

#### **1** INTRODUCTION

Linear dynamical systems are the canonical model for time series data. At each time step *t* there is an unknown hidden state  $x_t \in \mathbb{R}^n$  and a known exogenous input  $u_t \in \mathbb{R}^p$ . The transition dynamics

CC () BY

This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC '23, June 20–23, 2023, Orlando, FL, USA © 2023 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-9913-5/23/06. https://doi.org/10.1145/3564246.3585247 Allen Liu cliu568@mit.edu Massachusetts Institute of Technology Cambridge, Massachusetts, USA

Morris Yau morrisy@mit.edu Massachusetts Institute of Technology Cambridge, Massachusetts, USA

and observations  $y_t \in \mathbb{R}^m$  are generated according to the following rules:

$$x_{t+1} = Ax_t + Bu_t + w_t,$$
  
$$y_t = Cx_t + Du_t + z_t,$$

Here *A*, *B*, *C* and *D* are matrices of dimension  $n \times n$ ,  $n \times p$ ,  $m \times n$  and  $m \times p$  respectively. Moreover  $w_t$  and  $z_t$  are independent random variables and are called the process and observation noise respectively. It is standard to assume that they, along with the inputs and the initial state  $x_0$ , are all Gaussian, though we will work in a more general setting.

Linear dynamical systems have wide-ranging applications in control theory [26], computer vision [18], speech recognition [48], econometrics [2], healthcare [45] and neuroscience [61]. They are the de facto model of choice due to their mathematical simplicity and because, when the parameters are known, making predictions about subsequent observations and making inferences about the unknown state are both algorithmically tractable. In fact these algorithms are simple, practical and statistically optimal.

But what happens when the parameters are unknown? The problem of estimating A, B, C and D from input-output sequences is called system identification and has been intensively studied since Rudolph Kalman's pioneering work in the 1960's [35]. There is a well-developed theory that furnishes asymptotic guarantees [1, 46]. And more recently, many researchers have sought finite-sample guarantees both in the fully observed setting where C = I [13, 20, 58, 64] and in the partially observed setting [27, 51, 59, 63, 68]. Our focus here will be on obtaining running time and sample complexity bounds that are polynomial in the appropriate parameters and work under the most general conditions.

#### 1.1 Previous Work

In the fully observed setting, the maximum likelihood estimator can be computed by solving ordinary least squares. It is known to be statistically optimal and there are strong finite sample guarantees on its performance [13, 20, 58, 64]. The partially observed setting is significantly more challenging because the problem of computing the maximum likelihood estimator becomes nonconvex. The EM algorithm [24] is often used in practice but it can get stuck in bad local minima. Our main focus will be on algorithms for learning partially observed linear dynamical systems with provable guarantees. There is a vast literature on this and related prediction problems (see Section 2). But all existing algorithms need to make one or more of the following types of restrictive assumptions: STOC '23, June 20-23, 2023, Orlando, FL, USA

- (1) Assumptions about the characteristic polynomial q of A or the phases of its roots. Hardt, Ma and Recht [27] assumed that the image of the complex unit disk under q is contained in the cone of complex numbers whose real part is larger than the absolute value of its imaginary part. For example, this is satisfied if the  $\ell_1$ -norm of the coefficients of q is at most  $\sqrt{2}/2$ . Hazan et al. [28] studied the problem of predicting subsequent observations in a non-stochastic setting. Their bounds depend on the  $\ell_1$ -norm of the coefficients of a polynomial *p* that vanishes on the phases of the eigenvalues of A. In particular, when there are few distinct roots or they are pairwise separated, the  $\ell_1$ -norm of the coefficients of p can be much smaller than for q. This notion was further refined by Simchowitz et al. [63]. However it is not clear why one would expect these norm bounds to be small. In many settings, unless there is extreme cancellation, the coefficients of q would in fact be exponentially large.
- (2) Strict stability and mixing. Another popular assumption is called strict stability, which stipulates that the spectral radius  $\rho(A) < 1$ . Often the transition matrix A only satisfies  $\rho(A) \leq 1$ , which is called marginal stability. Consider a classic application in control theory, of tracking an object from radar measurements. The state of the object at some time step is its position, velocity and acceleration. The transition matrix is derived from Newton's laws and is upper triangular with ones along the diagonal, and so all of its eigenvalues are one. There are many other such examples, particularly in econometrics and coming from discretizations of ODEs. Algorithms that assume strict stability generally have bounds that depend on  $1/(1 - \rho(A))$  [27, 62]. Essentially, strict stability requires that the distribution of the  $y_t$ 's eventually converges and that there are no long-range correlations. So after about  $1/(1 - \rho(A))$  steps we essentially get fresh independent samples. Yet in many applications longrange correlations are an essential feature of the problem. Moreover getting around strict stability has many qualitative parallels with learning in graphical models without correlation decay [6], and learning in Gaussian graphical models without the restricted eigenvalue condition [38].
- (3) **Restrictions on the dimension, etc.** Some algorithms only work in the single-input single-output setting, i.e. when m = p = 1 [27]. Others have bounds that depend exponentially on the size of the largest Jordan block of *A*, or even treat the number of parameters of the linear dynamical system as a constant [63].

By now, there is a standard blueprint which works as follows: The first step is to estimate the Markov parameters, given by

$$\begin{bmatrix} D & CB & CAB & \cdots & CA^sB \end{bmatrix}$$

The second step is to apply the Ho-Kalman algorithm [30], which uses the Markov parameters to compute estimates  $\hat{A}$ ,  $\hat{B}$ ,  $\hat{C}$  and  $\hat{D}$  that are close to the true parameters in the appropriate metric. Oymak and Ozay [51] gave the first effective stability bounds for the Ho-Kalman algorithm. Thus the main issue is: How do you estimate the Markov parameters? Essentially all previous works use some form of linear regression. The analysis is based on expressing the observation  $y_t$  as a linear function of the previous inputs and some

Ainesh Bakshi, Allen Liu, Ankur Moitra, and Morris Yau

noise terms. The Markov parameters can then be extracted from the regressor. The noise terms are a function of observation and process noise and also the quantity  $A^s x_{t-s}$ , which captures how the state at some previous time step affects the current state. When A is strictly stable, this term decays exponentially. But when A is only marginally stable, controlling this error presents many challenges. Our main question is:

Are there efficient algorithms for learning high-dimensional linear dynamical systems whose running time and sample complexity are polynomial in the appropriate parameters, and whose assumptions are essentially optimal?

#### 1.2 Our Assumptions

It is important to draw a sharp distinction between the assumptions featured in the previous subsection and the more standard assumptions from control theory. In 1960, Rudolph Kalman [36] introduced the concepts of observability and controllability. Since then, it has been understood that they ought to in some sense govern what sorts of linear dynamical systems can be learned. In this subsection, we will review these assumptions and their natural quantitative counterparts.

*Observability and Controllability.* Consider the *observability matrix*: for an integer *s*, let

 $O_s = \begin{bmatrix} C^\top & (CA)^\top & \dots & \left(CA^{s-1}\right)^\top \end{bmatrix}^\top.$ 

A linear dynamical system is *observable* if for some *s*, the matrix  $O_s$  has full column rank. Intuitively, this condition ensures that there is no portion of the state space that we cannot observe eventually.

Now consider the *controllability matrix*: for an integer s, let

$$Q_s = \begin{bmatrix} B & AB & \dots & A^{s-1}B \end{bmatrix}$$

A linear dynamical system is *controllable* if the controllability matrix has full row rank. Intuitively this condition ensures that there is no portion of the state space that cannot be reached by the appropriate inputs. If either observability or controllability are violated, it is information-theoretically impossible to learn.

While the full rank conditions are enough to build an asymptotic theory, we will need natural quantitative counterparts to get finite sample guarantees. In particular we assume that  $O_s$  and  $Q_s$ have bounded condition number for some *s*. These assumptions are usually made in addition to the ones from the previous subsection, as they are needed in the stability bounds for the Ho-Kalman algorithm [51]. Furthermore we show that (see Theorem 1.4) they are information-theoretically necessary in order to learn a linear dynamical system from a polynomial length trajectory.

Finally, as is standard, we also assume that the system is *non-explosive*, i.e. the eigenvalues of *A* are bounded by 1 in magnitude. Note that assuming the eigenvalues are bounded is much weaker than assuming the singular values are bounded (e.g. consider the types of upper triangular matrices that arise in control theory, including *n*-dimensional integrators [57]).

*Relaxed Control and Noise.* In the literature, the standard assumption is that the initial state, the process and observation noise are all drawn from a Gaussian. But Gaussianity is not meant to literally be

true and it is often assumed for convenience. We show that we can dramatically relax this assumption to allow heavy-tailed distributions instead. In particular, for the control input,  $u_t$ , we only require the underlying distribution to have well-behaved fourth-moments:

**Definition 1.1** ((4,2)-Hypercontactivity). A distribution  $\mathcal{D}$  over  $\mathbb{R}^d$  is (4, 2)-hypercontractive if for all v,

$$\mathop{\mathbb{E}}_{x\sim\mathcal{D}}\left[\langle x,v\rangle^4\right] \leqslant O(1) \mathop{\mathbb{E}}_{x\sim\mathcal{D}}\left[\langle x,v\rangle^2\right]^2.$$

We note that several families of distributions are hypercontractive, including Gaussians, uniform distributions over the hypercube, sphere and other convex bodies, the Laplace, gamma, chi-squared, Wishart, Dirichlet and beta distributions, and in general, all logconcave distributions. Further, the set of hypercontractive distributions is closed under affine transformations, products and mixtures.

Finally, we only require that the distributions of the process noise,  $w_t$ , and observation noise,  $z_t$ , have bounded covariance. We state these assumptions formally in Section 4.

#### 1.3 Our Results

Our approach is based on the *method-of-moments* rather than least-squares regression. Our starting point is the following folklore observation: for any integers t > j,

$$\mathbb{E}_{u_t \sim \mathcal{D}_u} \left[ y_{t+j} u_t^\top \right] = \begin{cases} D & \text{if } j = 0\\ CA^{j-1}B & \text{otherwise} \end{cases}$$
(1)

However getting accurate estimates of the Markov parameters is a challenging task. For a fixed *j*, since the expectation of the estimator  $y_{t+j}u_t^{\top}$  does not depend on *t*, a natural approach to estimate  $\mathbb{E}\left[y_{t+j}u_t^{\top}\right]$  is to average over several control-observation pairs:  $\overline{CA^{j-1}B} = \frac{1}{T} \sum_{t \in [T]} y_{t+j}u_t^{\top}$  and hope that this estimator converges to its expectation. Unfortunately, this is just not true!

The first issue is that samples of the form  $y_{t+j}u_t^{\top}$  are not independent for different values of t. The second issue is that, in the marginally stable setting, the variance of this statistic grows with t, even when the control and the noise are Gaussian (see the appendix of the full version for an example). Thus, directly using the empirical estimate can be highly inaccurate *no matter how long our trajectory is*. One of the key steps in our algorithm is to learn a transformation of the observations to a new time series  $\{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_T\}$  such that  $\mathbb{E}\left[\hat{y}_{t+j}u_t^{\top}\right] = CA^{j-1}B$  and the variance of our estimator is bounded. As a result, we obtain the following theorem:

THEOREM 1.2 (EFFICIENTLY LEARNING A LINEAR DYNAMICAL SYS-TEM, INFORMAL THEOREM 6.4). Given  $\epsilon > 0$ , a fixed polynomial length trajectory from a linear dynamical system satisfying mild nondegeneracy assumptions (see Subsection 1.2), there exists an algorithm that outputs estimates  $\hat{A}, \hat{B}, \hat{C}, \hat{D}$  such that with probability at least 9/10, there exists a similarity transform U satisfying

$$\left\|A-U^{-1}\hat{A}U\right\|\leqslant\epsilon, \left\|B-U^{-1}\hat{B}\right\|\leqslant\epsilon, \left\|C-\hat{C}U\right\|\leqslant\epsilon, \left\|D-\hat{D}\right\|\leqslant\epsilon.$$

Further, the algorithm runs in time that is a fixed polynomial in all the parameters.

**Remark 1.3.** Note that it is only possible to recover the system parameters up to some global transformation *U* since all such transformations lead to equivalent dynamics, see e.g. [51].

The main appeal of our algorithm is that it works in essentially the most general setting possible. In particular we show the following lower bound:

THEOREM 1.4 (SAMPLE COMPLEXITY LOWER BOUND FOR ILL-CON-DITIONED SYSTEMS). [Informal, see Theorem 9.4] If for an LDS, the observability matrix  $O_s$  has smallest singular value less than  $\delta$  for all orders s, then any algorithm that uses less than  $\sim 1/\sqrt{\delta}$  length trajectories incurs constant error in estimating A, B, C, D with constant probability. The same statement holds with the observability matrix  $O_s$  replaced by the controllability matrix  $Q_s$ . In particular, if  $\delta$  is exponentially small, then an exponential number of samples are required to learn the parameters.

It turns out that super-resolution [7, 17], namely the task of recovering a sparse signal from noisy low-frequency measurements, corresponds to a special case of learning linear dynamical systems. It is known that super-resolution exhibits a sharp phase transition, where the problem goes from having efficient algorithms with polynomial running time and sample complexity, to being information-theoretically impossible, unless the noise is exponentially small [49]. Thus there are some linear dynamical systems where it is impossible to learn the true parameters with bounded length trajectories. We refine this connection to show instancewise lower bounds for learning any linear dynamical system whose observability or controllability matrices are close to singular. Thus the assumptions our algorithm needs are qualitatively tight, and our results close the question of what linear dynamical systems can be efficiently learned.

#### 2 RELATED WORK

Linear Time Invariant Systems: Identification, Prediction, Estimation. There is a long history of identifying linear dynamical systems from measurements, see [23] for extensive references. A focus of these works is on the "pre-filtering" approach to handling long range correlations in learning dynamical systems, see [15] [70] [66] [63]. Recently, there is a flurry of work on prediction and estimation for LDS's through the framework of no regret learning both in the fully observable setting [65] [58] [19] and the partially observed setting [62] [27]. For a variety of assumptions on the dynamics matrix, such as diagonalizability, there is work on learning marginally stable LDS's [28] [29]. Many works take a regression approach to estimating the markov parameters of the LDS for strictly stable systems see [44] [60] [21] [16]. In these settings it is possible to take advantage of the decay of the coefficients of the associated regressors. Marginal stability can be handled with multiple trajectories see [71] [67]. Closed loop system identification has also been studied; see [43] [42].

Somewhat related to our work is the problem of prediction without system identification in marginally stable LDS's [68] [25] but with an assumption on the exponential decay of the kalman filter coefficients. In [56], the exponential stability of the Kalman filter assumption is removed via a procedure that builds a succinct bank of filters for the prediction task and with an additional assumption on the dynamics having real eigenvalues. For a survey of the area see [69]. *Relaxing distributional assumptions.* In recent years, there has been a tremendous amount of work on designing algorithms that do not rely on strong distributional assumptions, such as Gaussianity, and only require much milder conditions. In particular, hypercontractivity of linear functions and low-degree polynomials have been identified as key analytic conditions that admit efficient algorithms for numerous problems in high-dimensional algorithmic statistics. In particular, souped up variants of *hypercontractivity* are used for heavy-tailed mean and covariance estimation [11, 31, 47], robust moment estimation [41], robust regression [5, 10, 34, 39, 52, 53, 72], robustly clustering mixture models [4, 14, 32, 40] and list-decodable learning [3, 12, 33, 37, 54, 55]. Algorithms with relaxed distributional assumptions were also recently given for online regression [9] and Kalman filtering [8].

#### **3 TECHNICAL OVERVIEW**

In this section, we describe our key algorithmic ideas and the corresponding technical challenges involved.

#### 3.1 A Thought Experiment

Consider the setting where we already know the parameters, *A*, *B*, *C* and *D*, of the underlying linear dynamical system. While there is nothing left to learn in this setting, we can still ask whether there exists a transformation of the observations  $\{y_t\}_{t \in [0,T]}$  to a new time series  $\{\hat{y}_t\}_{t \in [0,T]}$ , such that the variance of the random variable  $\hat{y}_{t+i}u_t$  is bounded.

It is indeed possible to do so by considering a simple linear transformation of the observations: let  $\hat{y}_t = y_t - \sum_{j=1}^n c_j y_{t-j}$ , where the  $c_j$ 's are the coefficients of the characteristic polynomial of *A*. To see why this works, we recall that by the Cayley-Hamilton theorem (see Fact 5.4), the coefficients of the characteristic polynomial satisfy the following algebraic identity:

$$A^n - \sum_{j \in [n]} c_j A^{n-j} = 0 \tag{2}$$

Therefore, assuming (for the purposes of exposition) that  $w_t$ 's and  $z_t$ 's are bounded, and D = 0, we have

$$\hat{y}_{t} = y_{t} - \sum_{j=1}^{n} c_{j} y_{t-j}$$

$$= \underbrace{\sum_{i=1}^{n} \left( CA^{i-1}B - \sum_{j=1}^{i-1} c_{j}CA^{i-j-1}B \right) u_{t-i}}_{(a)}$$

$$+ \underbrace{\sum_{i=n+1}^{t} \left( CA^{i-n-1} \left( A^{n} - \sum_{j=1}^{n} c_{j}A^{n-j} \right) B \right) u_{t-i}}_{(b)}$$
(3)

We note that term (a) above only has *n* terms and does not grow as a function of *t*, and term (b) is in fact zero, since we can repeatedly apply the identity from Equation (2). A simple computation then implies that the estimator  $\hat{y}_{t+j}u_t^{\mathsf{T}}$  satisfies  $\mathbb{E}\left[\hat{y}_{t+j}u_t^{\mathsf{T}}\right] = CA^{j-1}B$ and has bounded variance. We replicate this thought experiment, by learning the coefficients that stabilize the variance of our estimator from the observations directly. We dedicate the rest of the technical overview to describe how we accomplish this task.

#### 3.2 Learning the Stabilizing Transform

For ease of exposition, we assume that *C* is a  $1 \times n$  matrix and therefore the resulting observations,  $y_t$ , are scalars. A natural approach is to then consider the following least-squares regression problem:

$$\min_{c_1, c_2, \dots, c_n} \sum_{t \in [T]} \left( y_t - \sum_{j \in [n]} c_j y_{t-j} \right)^2$$

We know that the coefficients of the characteristic polynomial are a feasible solution to this regression problem, and the resulting linear transformation of the observations results in an estimator with bounded variance. Such an approach also appears in [63] but they incur unspecified, potentially exponential dependencies on the system parameters due to the complexities of analyzing this regression problem directly. In particular, we do not have finegrained control over the solution returned by solving the regression problem, and apriori, the regression solution need not be close to the coefficients of the characteristic polynomial of A.

*Convex Program.* Instead, we take a more direct approach to stabilizing the variance and consider a different convex program, specifically designed to do so. In particular, we find a vector  $\alpha = (\alpha_1, \alpha_2, ..., \alpha_s)$  such that the following constraint system is feasible:

$$C_{\alpha} = \begin{cases} \forall j \in [s] & |\alpha_j|^2 \leq P_0 \\ \forall i \in [T] & |y_{i+k} - \sum_{j \in [s]} \alpha_j \cdot y_{i-j}|^2 \leq P_1 \end{cases}, \qquad (4)$$

where *s* is the integer satisfying the observability and controlability assumptions from Definition 4.4, and  $P_0$  and  $P_1$  are sufficiently large polynomials in the system parameters (see Algorithm 6.6 for details).

Intuitively, the first constraint posits that each coefficient,  $\alpha_j$  is bounded in magnitude. This is necessary since the process and observation noise scale proportional to the coefficients in the linear transformation, and we cannot afford to pay exponentially in these quantities. The second constraint posits that the resulting observations themselves are bounded, and tries to enforce a universal bound on the variance of each  $\hat{y}_t$  appearing in the estimator  $\frac{1}{T} \sum_{t \in [T]} \hat{y}_{t+k} u_t^{\mathsf{T}}$ .

*Feasibility.* Observe, in contrast to the characteristic polynomial, we are only taking a linear combination of the previous *s* (potentially  $\ll n$ ) observations. While this difference does not manifest itself when *C* is  $1 \times n$ , it becomes crucial when *C* is  $m \times n$  for m > 1 for obtaining guarantees that depend only on the observability and controllability matrix. Also, note that we are expressing  $y_{t+k}$ , rather than  $y_t$ , as a linear combination of  $y_{t-1}, \ldots, y_{y-s}$ . This difference is also crucial in the construction and analysis of our estimator.

In order to establish feasibility of  $C_{\alpha}$ , we invoke the observability assumption: since  $O_s$  has bounded condition number, there exists a vector  $\alpha^* = (\alpha_1^*, \ldots, \alpha_s^*)$  such that each  $\alpha$  is bounded and the following identity holds:

$$CA^{k+s} - \alpha_1 CA^{s-1} - \alpha_2 CA^{s-2} - \ldots - \alpha_s C = 0.$$

We then follow an argument similar to the one in Equation (3) to show that the magnitude of  $\hat{y}_t$ 's is bounded (see Lemma 7.2 for details). We also note that the above program is convex, and admits an efficient separation oracle, and therefore, we can find a feasible  $\alpha$  in polynomial time. Interestingly, the feasibility analysis only requires that the covariance of the control input, process noise and observation noise be bounded, and does not require strong assumptions such as sub-Gaussian tails.

The Anti-concentration Potential. Next, we show that any feasible solution to the constraint system actually yields a stabilized estimator. To accomplish this goal, we design a potential function that captures the variance of our estimator, and argue that if the potential is large, with high probability, some constraint in  $C_{\alpha}$  must be violated. In particular, for any vector  $\alpha$ , and integer l, we consider the potential

$$\mathcal{G}_{\alpha,l} = \sum_{i=0}^{l} \|F_{\alpha}(A)A^{i}B\|_{F}^{2}$$

where

$$F_{\alpha}(A) = CA^{k+s} - \alpha_1 CA^{s-1} - \alpha_2 CA^{s-2} - \ldots - \alpha_s C$$

We observe that the terms appearing in this potential are the trace of  $\mathbb{E}\left[\left(F_{\alpha}A^{i}Bu_{t-i}\right)^{\top}\left(F_{\alpha}A^{i}Bu_{t-i}\right)\right]$ , which captures how large intermediate terms are, as a function of  $\alpha$ . In particular, if  $\alpha = \alpha^{*}$ , the trace would be 0. We make this intuition precise in Lemma 7.4.

Next, we show that we can split up the terms appearing in  $\hat{y}_{t+k}$  into three parts as follows:

$$\hat{y}_{t+k} = X_t + V_t + W_t,$$

where  $X_t$  is polynomially bounded in the system parameters,  $V_t$  is a random variable such that the covariance matrix of  $V_t$ , denoted by  $\Sigma_{V_t}$  satisfies  $\text{Tr}(\Sigma_{V_t}) = \mathcal{G}_{\alpha,l}$ , and  $W_t$  is a random variable that we do not have control over, and may potentially be unbounded. This presents obstacle since  $W_t$  can wipe out the information contained in  $V_t$ , and  $\mathcal{G}_{\alpha,l}$  may be large without violating any constraint in  $\mathcal{C}_{\alpha}$ .

Here, we observe that such an event can be avoided precisely when the random variable  $V_t$  is *anti-concentrated*, i.e. the probability that  $V_t$  lands in a ball of small radius is small. Perhaps counterintuitively, we show that if the 4-th moment of  $V_t$  concentrates, then it already possesses the *anti-concentration* properties we require. We make this precise in Lemma 5.9, where we establish a Paley-Zigmund style inequality, showing that if a random variable is (4, 2)-hypercontractive (see Definition 5.7), then the probability it lands in any interval that is a constant fraction of its variance is bounded by a constant. Finally, we show that  $V_t$  is (4, 2)hypercontractive, and therefore, significantly relax the Gaussianity assumption.

To summarize, we show that if the potential is large, the magnitude of  $V_t$  is large, and since  $V_t$  is anti-concentrated,  $W_t$  cannot wash away this information. Therefore,  $|\hat{y}_{t+k}|$  must be large, for some *t*, which is a contradiction to the feasibility of  $C_{\alpha}$ .

Dependent Random Variables and Decoupling. We then establish that if the potential  $\mathcal{G}_{\alpha}$  is small, for a fixed setting of  $\alpha$ , the resulting

estimator has bounded variance:

$$\mathbb{E}\left[\left\|\frac{1}{L}\sum_{t\in[L]}\hat{y}_{t+j}u_t^{\top} - CA^{j-1}B\right\|^2\right] \leqslant \frac{P_1}{L}\left(\|\alpha\|^2 + G_{\alpha,L}\right), \quad (5)$$

for some fixed polynomial  $P_1$  in the system parameters. We treat *L* as a sufficiently large polynomial in the system parameters and  $1/\varepsilon$  (where  $\varepsilon$  is the desired accuracy).

This argument is fairly involved and heavily uses the independence of the  $u_t$ ,  $w_t$  and  $z_t$ 's. We refer the reader to Lemma 7.8 for a complete proof. While the above inequality holds for a fixed setting of  $\alpha$ , we note that the  $\alpha$ 's output by solving the constraint  $C_{\alpha}$  themselves depend on the randomness in the control input and the noise non-trivially.

To overcome this issue, we decouple the  $\alpha$ 's from  $u_t$ 's and establish a symbolic matrix inequality, where the matrices only depend on the  $u_t$ 's. Here, we treat the vector  $v_{\alpha} = (1, \alpha_1, ..., \alpha_s)$  as a formal variable, and write the potential as a quadratic form in the vector  $v_{\alpha}$ :

$$\mathcal{G}_{\alpha,L} = v_{\alpha}^{\top} G_L v_{\alpha},$$

where  $G_L$  is a PSD matrix. We note that such a representation always exists and is unique since  $\mathcal{G}_{\alpha,L}$  is a sum-of-squares in  $\alpha$ . Similarly, we observe that the variance we want to bound admits a similar decomposition: let  $M_j$  be the PSD matrix such that

$$\mathbb{E}\left[\left\|\frac{1}{L}\sum_{t\in[L]}\hat{y}_{t+j}u_t^{\top}-CA^{j-1}B\right\|^2\right]=v_{\alpha}^{\top}M_jv_{\alpha}.$$

Observe, the matrices  $G_L$  and  $M_j$  are independent of the formal variables  $\alpha$ , and in Corollary 7.10 we establish the following inequality:

$$\mathbb{E}_{l,\mathbf{w},z}\left[M_j\right] \le \frac{P_1}{L} \left(I + G_L\right) \tag{6}$$

Since the above inequality holds for all quadratic forms simultaneously, one natural way to proceed would be to consider an  $\epsilon$ -net over the  $\alpha$ 's and union bound over each vector in the net satisfying the quadratic form in Equation (6). To execute this, we need a net that is fine enough to account for how much error we accumulate in a term of the from  $\hat{y}_{t+j}u_t^{\mathsf{T}}$ . Note, the largest terms we need to account for are roughly of the form  $\alpha_j A^L$ , which naively requires a  $1/||A||^L$ -net. Unfortunately, this net is too fine and we cannot afford to union bound over all the vectors in this net because in some sense we only have L samples.

Bounded Eigenvalues to Smaller Nets. To address the issue above, we show that for any  $n \times n$  matrix A with complex entries, if the eigenvalues of A are bounded by 1 in magnitude, the operator norm of  $A^L$  in fact grows as  $L^n$ , instead of exponentially in L (see Lemma 5.6 for a precise statement). Here, we crucially note that we only assume the eigenvalues, instead of the singular values (which would make this statement trivial but would rule out several important families of linear dynamical systems), are bounded. With this insight, we can then union bound over all vectors  $\alpha$  in a  $1/L^n$ -net, and establish equation (5) for the  $\alpha$  output by solving the constraint system  $C_{\alpha}$ . To conclude, we have shown that the variance of our estimator is bounded, and therefore, we can estimate each block of the Markov Parameter matrix.

STOC '23, June 20-23, 2023, Orlando, FL, USA

#### 3.3 Lower Bound for Ill-Conditioned LDS's

In light of our results on learning LDS's, a natural question is whether the assumptions on observability and controllability are necessary. We exhibit information-theoretic lower bounds on the sample complexity of learning a linear dynamical systems with Gaussian noise and Gaussian inputs, when the observability or controllability matrices are exponentially ill conditioned. For this section, we discuss the case when the observability matrix is ill conditioned. The case when the controllability matrix is ill conditioned follows essentially the same argument.

**Definition 3.1.** We say that an LDS  $\mathcal{L}(A, B, C, D)$  is  $(\delta, v)$ -unobservable if v is a unit vector such that for all integers  $s \ge 0$ ,

 $\|CA^s v\| \leq \delta.$ 

Now the key to proving an information-theoretic lower bound is the observation that when the input and noise distributions are Gaussian, the system measurements and control inputs

$$(u_0, ..., u_T, y_0, ..., y_T)$$

are a Gaussian process. Therefore, the joint distribution is uniquely determined by its covariance matrix. On the other hand, we can explicitly compute the covariance matrix in terms of the system parameters *A*, *B*, *C*, *D*. While there are several terms in the expression, the main point is that essentially all of the terms look like  $CA^{j}B$ . Now, when the system is  $(\delta, v)$ -unobservable, we can replace *B* with  $B + vu^{\top}$  for an arbitrary vector  $u \in \mathbb{R}^{p}$  while only changing expressions of the form  $CA^{j}B$  by a little bit.

Overall, we can show that the pair of LDS's  $\mathcal{L} = \mathcal{L}(A, B, C, D)$ and  $\mathcal{L}' = \mathcal{L}(A, B + vu^{\top}, C, D)$  are statistically close up to time  $T \ll 1/\delta$ . This means that no algorithm using  $\ll 1/\delta$  length trajectories can distinguish the two systems and since their parameters are not close to equivalent up to similarity (for generic choices of parameters), the algorithm must incur large error.

#### **4 FORMAL SETUP**

In this section, we formally state the linear dynamical system model, and our assumptions.

MODEL 4.1 (LINEAR DYNAMICAL SYSTEM). Let  $A \in \mathbb{C}^{n \times n}$ ,  $B \in \mathbb{C}^{n \times p}$ ,  $C \in \mathbb{C}^{m \times n}$ , and  $D \in \mathbb{C}^{m \times p}$  be complex valued matrices. Let  $\mathcal{D}_0$ ,  $\mathcal{D}_u$ ,  $\mathcal{D}_w$ ,  $\mathcal{D}_z$  be distributions with mean zero. Then, a Linear Dynamical System,  $\mathcal{L}(A, B, C, D)$ , is defined as follows:

$$x_{t+1} = Ax_t + Bu_t + w_t,$$
  
$$y_t = Cx_t + Du_t + z_t,$$

where  $x_0 \sim \mathcal{D}_0$ , and for all  $t \in \mathbb{N}$ ,  $u_t \sim \mathcal{D}_u$ ,  $w_t \sim \mathcal{D}_w$  and  $z_t \sim \mathcal{D}_z$ .

We see only the sequence of observations  $y_1, y_2, \ldots, y_T$  up to some time *T* and our goal is to learn the parameters of the system *A*, *B*, *C*, *D*. We need some assumptions about the parameters and also the input and noise distributions which we discuss below (as otherwise the system may be degenerate and it may be informationtheoretically impossible to learn, see Section 9).

#### 4.1 Assumptions on the System Parameters

We begin by ensuring that the linear dynamical system at hand is not degenerate. This notion can be made precise by considering the *observability matrix*: **Definition 4.2** (Observability Matrix). For an integer *s*, define the matrix  $O_s \in \mathbb{R}^{sm \times n}$  as

$$O_s = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{s-1} \end{bmatrix}.$$

A LDS is *observable* if for some *s*, the matrix  $O_s$  has full columnrank. Similarly, we need to ensure that the control input is not degenerate, and only acts in a subspace that is not spanned by *A*. This is made precise by considering the *controllability matrix*:

**Definition 4.3** (Controllability Matrix). For an integer *s*, define the matrix  $Q_s \in \mathbb{R}^{sp \times n}$  as

$$Q_s = \begin{bmatrix} B & AB & \dots & A^{s-1}B \end{bmatrix}$$

A LDS is *controllable* is the controllability matrix has full rowrank. We note that we assume a quantitative strengthening of these two assumptions to  $O_s$  and  $Q_s$  having bounded condition number (and this is necessary, recall Theorem 1.4).

**Definition 4.4** (Well-Behaved Linear Dynamical System). We say a linear dynamical system  $\mathcal{L}(A, B, C, D)$  is *well-behaved* if the following assumptions hold:

(1) **Non-trivial Controller.** The matrix *B* satisfies  $||B|| \ge 1$ .

tude at most 1.

- (2) Non-trivial Measurement. The matrix C satisfies ||C|| ≥ 1.
  (3) Non-exposive System. All eigenvalues of A have magni-
- (4) **Bounded Condition Number.**  $O_s$  has full column-rank,  $Q_s$  has full row-rank and for some integer *s*, and parameter  $\kappa \ge 1$ ,

$$\sigma_{\max}(O_{2s})/\sigma_{\min}(O_s) \leq \kappa, \sigma_{\max}(Q_{2s})/\sigma_{\min}(Q_s) \leq \kappa.$$

**Remark 4.5.** Crucially, the above assumption is on the eigenvalues and not the singular values of *A*, which would be a far stronger assumption, as discussed in the introduction.

**Remark 4.6.** While the bounded condition number assumption is standard in the literature [22, 50], in Section 9, we show that a polynomial bound on the condition number is necessary in the sample complexity, even information-theoretically. As a consequence, it is impossible to learn an exponentially ill-conditioned system with polynomially bounded observations. Finally, note that up to polynomial factors, it suffices to have a bound on  $\sigma_{\max}(O_t)/\sigma_{\min}(O_s)$  as long as t/s > 1 + c for some positive constant c (see Claim 5.16). For simplicity, we wrote the above condition for t = 2s.

# 4.2 Assumptions on the Distribution of the Control and Noise

We consider the following assumptions over the control input, system and process noise distributions:

**Definition 4.7** (Distributional Assumptions). For all  $t \in [T]$ , we assume that  $u_t \sim \mathcal{D}_u$ ,  $w_t \sim \mathcal{D}_w$  and  $z_t \sim \mathcal{D}_z$  are each sampled independently from the corresponding distributions. Additionally,  $x_0 \sim \mathcal{D}_0$ . Then,

• Mean Zero:  $\mathcal{D}_u$ ,  $\mathcal{D}_w$ ,  $\mathcal{D}_z$  and  $\mathcal{D}_0$  are all mean 0 distribution.

- Isotropic and Hypercontractive Control: The covariance of  $\mathcal{D}_u$ ,  $\Sigma_{\mathcal{D}_u} = I$ , and  $\mathcal{D}_u$  is (4, 2, K)-hypercontractive for a fixed constant  $K \ge 3$  (see Definition 5.7).
- Bounded Variance Noise: For σ<sub>w</sub>, σ<sub>z</sub> ≥ 1, the covariances of D<sub>w</sub> and D<sub>z</sub> satisfy Σ<sub>D<sub>w</sub></sub> ≤ σ<sub>w</sub>I and Σ<sub>D<sub>z</sub></sub> ≤ σ<sub>z</sub>I.
- Starting Point: The distribution  $\mathcal{D}_0$  has covariance  $\Sigma_{\mathcal{D}_0} \leq \sigma_0 I$

**Remark 4.8.** We don't actually need that the  $u_i$ ,  $w_i$ ,  $z_i$  are all drawn from the same distribution across different time-steps. We only need that they are independent. In other words, all of our results still hold if we allow for there to be different distributions  $\mathcal{D}_{u,t}$ ,  $\mathcal{D}_{w,t}$ ,  $\mathcal{D}_{z,t}$  at each time-step that all satisfy the above assumptions.

#### **5 PRELIMINARIES**

We begin with some notation and basic facts from linear algebra and probability. For a matrix  $A \in \mathbb{C}^{n \times n}$ , we use  $||A|| = ||A||_{\text{op}} = \max_{\|u\|=1} ||Au\|_2$  and  $||A||_F = \sqrt{\sum_{i,j \in [n]} |A_{i,j}|^2}$ . We use the notation  $A^{\top}$  to denote the transpose when A only has real entries. Further, for  $A \in \mathbb{C}^{n \times m}$  such that  $n \ge m$ , let  $\text{SVD}(A) = U\Sigma V^{\top}$  denote the singular value decomposition of A, where  $U \in \mathbb{C}^{n \times m}$  and  $V^{\top} \in \mathbb{C}^{m \times m}$  are unitary matrices (see Definition 5.1, and  $\Sigma$  is a diagonal matrix, with the singular values denoted by  $\sigma_1 \ge \sigma_2 \ge \ldots \sigma_m \ge 0$ .

#### 5.1 Linear Algebra Background

**Definition 5.1** (Unitary Matrices). Given a symmetric matrix  $U \in \mathbb{C}^{n \times n}$  we say U is a unitary matrix if  $U^{\top}U = UU^{\top} = I$ .

FACT 5.2 (OPERATOR NORM OF UNITARY MATRICES). If  $Q \in \mathbb{C}^{n \times n}$  is a unitary matrix,  $\|Q\|_{op} = 1$ .

FACT 5.3 (SUB-MULTIPLICATIVITY OF OPERATOR NORMS). Given matrices  $A \in \mathbb{C}^{n \times d}$ ,  $B \in \mathbb{C}^{d \times m}$ ,  $||AB||_{op} \leq ||A||_{op} \cdot ||B||_{op}$ .

FACT 5.4 (CAYLEY-HAMILTON THEOREM). Given a square matrix  $A \in \mathbb{C}^{n \times n}$ , the characteristic polynomial of A is defined as  $p_A(\lambda) = det(\lambda I - A) = \lambda^n + c_{n-1}\lambda^{n-1} + \ldots c_1\lambda + c_0$ , where the coefficients,  $c_i$ , are scalar. Then, consider the matrix valued polynomial  $p_A(A) = A^n + c_{n-1}A^{n-1} + \ldots c_1A + c_0I$ . The Cayley-Hamilton theorem states  $p_A(A) = 0$ .

FACT 5.5. For any matrix  $A \in \mathbb{C}^{n \times n}$ , there is a unitary matrix Q such that  $Q^{-1}AQ$  is upper triangular.

PROOF. A must have some eigenvector, say v. Also normalize v so that it is a unit vector. Let  $Q_0$  be a matrix whose first column is v and whose columns form an orthonormal basis of  $\mathbb{C}^n$ . Then  $A' = Q_0^{-1}AQ_0$  has all entries in the first column equal to 0 except possibly the first entry. Now it suffices to compute a unitary matrix in  $\mathbb{C}^{(n-1)\times(n-1)}$  that transforms the  $(n-1)\times(n-1)$  submatrix of A' (excluding the first row and column) into an upper triangular matrix but this can be done by induction.

We also establish the following key lemma to upper bound the operator norm of matrix polynomials, when the underlying matrix has bounded eigenvalues.

LEMMA 5.6 (OPERATOR NORM OF A MATRIX WITH BOUNDED EIGEN-VALUES ). Let  $A \in \mathbb{C}^{n \times n}$  be a matrix and assume that all eigenvalues of A have magnitude at most 1. Then for any integer L,

$$||A^{L}|| \leq n \cdot (2(1 + ||A||)L)^{n}.$$

**PROOF.** By Fact 5.5 and Fact 5.2, without loss of generality we can assume that *A* is upper triangular. Then, all of its diagonal entries are eigenvalues so all of its diagonal entries have magnitude at most 1. Now we bound the magnitude of all entries of  $A^L$ . Consider the entry indexed by *i*, *j*. Clearly  $i \leq j$  or the corresponding entry is 0. Next, by definition of  $(A^L)_{i,j}$ ,

$$|(A^{L})_{ij}| = \sum_{\substack{i_{1},...,i_{L} \\ i \leqslant i_{1} \leqslant ... \leqslant i_{L-1} \leqslant j}} A_{ii_{1}} \cdots A_{i_{L-1}j}$$

$$\leqslant \sum_{\substack{i_{1},...,i_{L} \\ i \leqslant i_{1} \leqslant ... \leqslant i_{L-1} \leqslant j}} |A_{ii_{1}}| \cdots |A_{i_{L-1}j}|$$

$$\leqslant \binom{L-1+n}{n} (1+||A||)^{n}$$

$$\leqslant (2(1+||A||)L)^{n}.$$
(7)

where the second inequality counts the number of paths and uses that each path can contain at most n entries strictly above the diagonal.

#### 5.2 Probability Background

Next, we recall the definition of  $\ell_{4\rightarrow 2}$ -hypercontractivity for distributions.

**Definition 5.7** (Hypercontractivity). We say a distribution  $\mathcal{D}$  on  $\mathbb{R}^n$  is (4, 2, K)-hypercontractive if for any vector  $v \in \mathbb{R}^n$ , we have  $\mathbb{E}_{x \sim \mathcal{D}}[\langle v, x \rangle^4] \leq K \cdot \mathbb{E}_{x \sim \mathcal{D}}[\langle v, x \rangle^2]^2$ .

We note that *hypercontractivity* is a very mild assumption on the concentration behavior of 4-th moments of a distribution, and several well-studied families of distributions, including all sub-Gaussian, sub-Exponential and log-concave distributions satisfy this assumption.

LEMMA 5.8 (LINEAR TRANSFORM OF A HYPERCONTRACTIVE DISTRI-BUTION). Let  $\{u_i\}_{i \in [t]}$  be t iid samples from a distribution  $\mathcal{D}$  that is (4, 2, K)-hypercontractive (see Definition 5.7). Then, for any matrices  $M_1, \ldots, M_t \in \mathbb{R}^{m \times p}$ , the random variable  $\sum_{i \in [t]} M_i u_i + \cdots + M_t u_t$ is (4, 2, K)-hypercontractive.

Next, we obtain a weak anti-concentration bound via a Paley Zygmund like inequality:

LEMMA 5.9 (WEAK ANTI-CONCENTRATION VIA HYPERCONTRAC-TIVITY). Let z be a real-valued random variable such that  $\mathbb{E}[z] = 0$ ,  $\mathbb{E}[z^2] \ge 1$ ,  $\mathbb{E}[z^4] \le K$  for some constant  $K \ge 3$ . Then for any real number  $\beta$ ,

$$\Pr[|z-\beta| \leq 0.1] \leq 1 - \frac{1}{10K}.$$

PROOF. Clearly we must have  $K \ge 1$ . Assume for the sake of contradiction that the desired inequality is false. Without loss of generality we have  $\beta \ge 0$ . First consider the case where  $\beta \ge 0.3$ .

STOC '23, June 20-23, 2023, Orlando, FL, USA

Let *p* be the probability that  $z \le 0.2$ . We must have  $p \le 1/(10K)$ . Furthermore, since  $\mathbb{E}[z] = 0$ , we must have

$$0 = \mathbb{E}[z] = p\mathbb{E}[z|z \le 0.2] + (1-p)\mathbb{E}[z|z > 0.2]$$
  
$$\ge p\mathbb{E}[z|z \le 0.2] + 0.2(1-p)$$

which rearranges as

$$\mathbb{E}[z | z \leq 0.2] \leq \frac{-0.2(1-p)}{p}$$

Thus, by Jensen's inequality (since  $z^4$  is convex), this implies that

$$\mathbb{E}[z^4] \ge p\mathbb{E}[z^4|z \le 0.2] \ge p\left(\frac{0.2(1-p)}{p}\right)^4 > K$$

which is a contradiction. Now it remains to consider the case where  $\beta \le 0.3$ . Then let *q* be the probability that  $|z - \beta| \le 0.1$ . We have

$$\begin{split} \mathbb{E}[z^2] &= q \mathbb{E}\left[z^2 \middle| |z - \beta| \leqslant 0.1\right] + (1 - q) \mathbb{E}\left[z^2 \middle| |z - \beta| > 0.1\right] \\ &\leqslant q(\beta + 0.1)^2 + (1 - q) \mathbb{E}\left[z^2 \middle| |z - \beta| > 0.1\right] \,. \end{split}$$

Thus, since  $\mathbb{E}[z^2] \ge 1$ , we must have

$$\mathbb{E}\left[z^2\right||z-\beta| > 0.1\right] \ge \frac{0.8}{1-q}$$

Thus, by convexity, we must have

$$\mathbb{E}[z^4] \ge (1-q) \cdot \left(\frac{0.8}{1-q}\right)^2 = \frac{0.64}{1-q}$$

and combining with the fact that  $\mathbb{E}[z^4] \leq K$ , we deduce that  $q \leq 1 - \frac{1}{10K}$  and we are done.

#### 5.3 Linear Dynamical Systems Background

Next, we establish some basic definitions and identities that we utilize throughout. We begin with the definition of the Markov parameters of a LDS.

**Definition 5.10** (Markov Parameters). Given a linear dynamical system,  $\mathcal{L}(A, B, C, D)$ , and an integer  $T \ge 1$ , the Markov Parameter matrix  $G \in \mathbb{R}^{m \times Tp}$  is defined as the following block matrix:

$$G = \begin{bmatrix} D & CB & CAB & \dots & CA^{T-2}B \end{bmatrix}.$$

It will be important to consider linear combinations of the observations  $y_t$ . In particular, for different integers k, we will consider linear combinations of the form  $\hat{y}_{t+k} = y_{t+k} - \sum_{j=1}^{s} \alpha_j y_{t-j}$  where  $\alpha_1, \ldots, \alpha_s \in \mathbb{R}^{m \times m}$  are  $m \times m$  matrices. To ease notation, it will be useful to consider the following matrix polynomial.

**Definition 5.11** (Matrix Polynomial). For  $\alpha = (\alpha_1, ..., \alpha_s)$  where  $\alpha_1, ..., \alpha_s \in \mathbb{R}^{m \times m}$  are matrices, define the matrix polynomial  $F_{\alpha,k} : \mathbb{R}^{n \times n} \to \mathbb{R}^{m \times n}$ 

$$F_{\alpha,k}(X) = CX^{s+k} - \alpha_1 CX^{s-1} - \dots - \alpha_s C.$$

Also for i = 0, 1, ..., k + s, let

$$F_{\alpha,k}^{(i)}(X) = \begin{cases} CX^i & 0 \leq i \leq k \\ CX^i - \sum_{j=1}^{i-k} \alpha_j CX^{i-k-j} & k+1 \leq i \leq k+s \end{cases}$$

We now have the following identities.

FACT 5.12 (ALGEBRAIC IDENTITIES FOR LDS'S). Let  $\mathcal{L}(A, B, C, D)$  be a Linear Dynamical System (see Definition 4.1). Then, for any  $t \in \mathbb{N}$ ,

$$y_t = \sum_{i=1}^t \left( CA^{i-1} Bu_{t-i} + CA^{i-1} w_{t-i} \right) + CA^t x_0 + Du_t + z_t,$$

Further, given  $k \in \mathbb{N}$  and  $\alpha_1, \alpha_2, \ldots, \alpha_s \in \mathbb{R}^{m \times m}$ , let  $\hat{y}_{t+k} = y_{t+k} - \sum_{i=1}^{s} \alpha_i y_{t-i}$ . Then,

$$\begin{split} \widehat{y}_{t+k} = & \left( z_{t+k} - \sum_{j=1}^{s} \alpha_j z_{t-j} \right) + \left( Du_{t+k} - \sum_{j=1}^{s} \alpha_j Du_{t-j} \right) \\ & + \left( F_{\alpha,k}(A) A^{t-s} \right) x_0 + \sum_{i=1}^{k+s} F_{\alpha,k}^{(i-1)}(A) (Bu_{t+k-i} + w_{t+k-i}) \\ & + \sum_{i=k+s+1}^{t+k} F_{\alpha,k}(A) A^{i-(k+s+1)} (Bu_{t+k-i} + w_{t+k-i}). \end{split}$$

**PROOF.** The proof is direct computation (see full version).  $\Box$ 

Next, we observe that the *cross-covariance* between the control input and the observation is an unbiased estimator of the Markov parameters.

FACT 5.13 (CROSS-COVARIANCE OF CONTROL AND OBSERVATION). For any  $t, k \in \mathbb{N}$ , and any  $0 \le j \le k$ , we have

$$\mathbb{E}\left[\hat{y}_{t+j}u_t^{\top}\right] = \begin{cases} D \ if j = 0\\ CA^{j-1}B \ otherwise \end{cases}$$
(8)

**PROOF.** We use the formula in Fact 5.12 and the independence of the  $u_i$ ,  $w_i$ ,  $z_i$ . When j = 0, we immediately have

$$\mathbb{E}[\hat{y}_t u_t^{\top}] = \mathbb{E}[Du_t u_t^{\top}] = D.$$

When j > 0, we have

$$\mathbb{E}[\hat{y}_{t+j}u_t^{\top}] = \mathbb{E}[F_{\alpha,k}^{(j-1)}(A)Bu_tu_t^{\top}] = F_{\alpha,k}^{(j-1)}(A)B.$$

Next, since  $j \leq k$ , by definition we have  $F_{\alpha,k}^{(j-1)}(A) = CA^{j-1}$  and we are done.

In light of the above, we make the following definition.

**Definition 5.14.** For an integer  $j \ge 0$ , we define the matrix  $X_j$  as

$$X_j = \begin{cases} D \text{ if } j = 0\\ CA^{j-1}B \text{ otherwise} \end{cases}$$

Of course, we have  $X_j = \mathbb{E}[\hat{y}_{t+j}u_t^{\top}]$  by the previous fact.

We also require the following straight-forward consequences of the bounded condition number assumption on the observability and controlability matrices from Section 4:

**Claim 5.15.** Given an integer *s*, let  $O_s$ ,  $Q_s$  be the observability and controlability matrices from definition 4.2 and 4.3. Then,  $\sigma_{\min}(O_s) \leq \sqrt{s} ||C||$  and  $\sigma_{\min}(Q_s) \leq \sqrt{s} ||B||$ .

PROOF. Let *v* be an eigenvector of *A* and say  $Av = \lambda v$ . Note that  $|\lambda| \leq 1$ . Thus,

$$||O_s v|| \leq \sqrt{||Cv||^2(1+\lambda^2+\lambda^4+\cdots+\lambda^{2s-2})} \leq \sqrt{s}||C|| ||v||.$$

A similar argument works for  $Q_s$ .

**Claim 5.16.** Given integers t > s, let  $\sigma_{\max}(O_t) / \sigma_{\min}(O_s) \leq \kappa$ . Then for any integer k > 0,

$$\|A^k\|_F \leq (\sqrt{n\kappa})^{k/(t-s)}$$

The same holds with O replaced with Q.

**PROOF.** Note that  $||O_s A^{t-s}||_F \ge \sigma_{\min}(O_s) ||A^{t-s}||_F$ . On the other hand,

$$\|O_s A^{t-s}\|_F \leq \|O_t\|_F \leq \sqrt{n} \|O_t\| \leq \sqrt{n} \kappa \sigma_{\min}(O_s).$$

Thus,  $||A^{t-s}||_F \leq \sqrt{n\kappa}$  and this immediately implies the desired inequality. 

#### 6 ALGORITHM

Our algorithm follows the outline described in Section 3. We first learn a stabilizing transformation that transforms  $y_t \rightarrow \hat{y}_t$  (see Algorithm 6.6). Note that in the general case, when *C* is  $m \times n$  (as opposed to  $1 \times n$  as described in Section 3), the coefficients  $\alpha_1, \ldots, \alpha_s$ are  $m \times m$  matrices. Once we have these matrices, we empirically estimate the Markov parameters (recall Fact 5.13). Finally, once we obtain estimates for the Markov parameters, we extract the system parameters via the Ho-Kalman algorithm.

Before we describe our algorithm formally, we introduce a few definitions and notational simplifications. Throughout this section, we will use  $\varepsilon$  to denote the desired accuracy and  $\delta$  to be a parameter for the failure probability. We will also make the simplification that  $x_0 = 0$ . This is allowed because we can absorb the distribution of  $x_0$  into the distribution for  $w_0$  i.e.  $\mathcal{D}_w \leftarrow A\mathcal{D}_0 + \mathcal{D}_w$  which is still polynomially bounded in the system parameters. As mentioned in Remark 4.8, we do not need the  $w_i$  to have identical distributions, just that they are independent and bounded variance.

Algorithm 6.1 (Learning a Linear Dynamical System).

- **Input:** *T* observations  $\{y_1, y_2, \ldots, y_T\}$  and the corresponding control inputs  $\{u_1, u_2, \ldots, u_T\}$  generated from a Linear Dynamical System  $\mathcal{L}(A, B, C, D)$  satisfying the assumptions in Section 4.
- **Input:**  $s \in \mathbb{N}$  such that the observability and controllability matrices satisfy the condition number bounds in Definition 4.4, accuracy parameter  $0 < \epsilon < 1$ , and failure probability parameter  $0 < \delta < 1$ .

**Operation:** 

- (1) Stabilizing the System: Run Algorithm 6.6 to obtain coefficient matrices  $\alpha_1, \ldots, \alpha_s$ .
- (2) Estimating the Markov Parameters: Set k = 10s. For all i = 0, 1, ..., k
  - (a) For all  $t \in [T]$  with t > k + s, compute  $\hat{y}_t =$  $y_t - \sum_{i=1}^s \alpha_i y_{t-k-i}$ .

b) Compute 
$$\hat{X}_i = \frac{1}{\pi i} \sum_{i=1}^{T} \sum_{i=1}^{T} \hat{u}_i u_i$$

(b) Compute  $\hat{X}_j = \frac{1}{T-k-s} \sum_{t=k+s+1}^T \hat{y}_t u_{t-j}^\top$ . (3) **Robust Ho-Kalman**: Run Ho-Kalman on

$$\hat{G} = [\hat{X}_0, \hat{X}_1, \dots, \hat{X}_{2s}]$$

to obtain estimates  $\hat{A}, \hat{B}, \hat{C}, \hat{D}$ **Output:**  $\hat{A}, \hat{B}, \hat{C}, \hat{D}$  satisfying guarantees of Theorem 6.4 STOC '23, June 20-23, 2023, Orlando, FL, USA

In the definition below, we condense all of the system parameters into a set S so that we don't need to list out the full set of parameters in future computations. Recall Definition 4.7 where K is the hypercontractivity parameter for the input distribution (and will be defined as such throughout this section) and  $\sigma_w, \sigma_z$  are upper bounds on the variances of the noise distributions.

#### **Definition 6.2.** Let *S* denote the set of parameters

 $\{ \|A\|, \|B\|, \|C\|, \|D\|, m, n, p, s, \kappa, K, 1/\delta, \sigma_w, \sigma_z \}.$ 

We will write poly(S) for a quantity that depends polynomially on these parameters.

Note that we will be more explicit about dependencies on the accuracy  $\varepsilon$  so it is not included in the definition of S.

Now we define a constraint system for a convex program that is at the core of our algorithm. Let  $\varepsilon$  be the desired accuracy, and k = 10s. Throughout this section, we will treat k as fixed. Choose sufficiently large polynomials  $P_0, P_1, P_2$  in terms of the parameters in S such that  $P_0 \ll P_1 \ll P_2$  and let  $L = P_2 \log^2(1/\varepsilon)/\varepsilon^2$ .

**Definition 6.3.** We define the constraint system  $C_{\alpha}$  for matrices of coefficients  $\alpha_1, \ldots, \alpha_s \in \mathbb{R}^{m \times m}$  as follows. We define  $\hat{y}_{t+k} = y_{t+k}$  –  $\alpha_1 y_{t-1} - \cdots - \alpha_s y_{t-s}$ . We also enforce the following constraints

- (1)  $\|\alpha_i\|_F \leq P_0$  for all i = 1, 2, ..., s
- (2) For all  $i \in [100 \text{ sn} m^2 K \log L]$ , we have

 $\|\hat{y}_{iL+k}\| \leq P_1 \log(1/\varepsilon)$ 

Our main theorem is that Algorithm 6.1 runs with polynomial time and sample complexity and outputs estimates  $\hat{A}$ ,  $\hat{B}$ ,  $\hat{C}$ ,  $\hat{D}$  that are  $\varepsilon$ -close to the true system parameters up to a global similarity transformation (which is always necessary).

THEOREM 6.4 (LEARNING A LINEAR DYNAMICAL SYSTEM). Given  $0 < \epsilon, \delta < 1$ , an integer s, and trajectory length

$$T = \Omega\left(\operatorname{poly}(\mathcal{S}) \cdot \frac{\log^3(\frac{1}{\epsilon})}{\epsilon^4}\right),\,$$

and the corresponding observations and inputs  $\{y_i, u_i\}_{i \in [T]}$ , from a linear dynamical system  $\mathcal{L}(A, B, C, D)$ , satisfying the assumptions in Section 4, Algorithm 6.1 outputs estimates  $\hat{A}$ ,  $\hat{B}$ ,  $\hat{C}$ ,  $\hat{D}$  such that with probability at least  $1 - \delta$ , there exists a similarity tranform U satisfying

$$\max\left(\left\|A - U^{-1}\hat{A}U\right\|, \left\|C - \hat{C}U\right\|, \left\|B - U^{-1}\hat{B}\right\|, \left\|D - \hat{D}\right\|\right) \le \epsilon,$$

Further, Algorithm 6.1 runs in poly  $(S, \frac{1}{c})$  time.

The overall structure of our approach is to first estimate the Markov parameters (see Definition 5.10). We then recover the estimates of the system parameters by setting up a Generalized eigenvalue problem and using the Ho-Kalman Algorithm (see for example Theorem 5.3 in [51]).

The key techincal theorem we obtain for learning each block matrix in the Markov Parameter matrix is as follows:

THEOREM 6.5 (LEARNING THE MARKOV PARAMETERS). Given 0 <  $\epsilon, \delta < 1$ , an integer s, and trajectory length

$$T = \Omega\left(\operatorname{poly}(\mathcal{S}) \cdot \frac{\log^3(\frac{1}{\epsilon})}{\epsilon^2}\right),$$

with the corresponding observations and inputs  $\{y_i, u_i\}_{i \in [T]}$ , from a linear dynamical system  $\mathcal{L}$  (A, B, C, D), satisfying the assumptions in Section 4, Algorithm 6.6 outputs  $\alpha_1, \alpha_2, \ldots, \alpha_s$  in time poly( $\mathcal{S}, 1/\epsilon$ ) such that with probability at least  $1 - \delta$ , for all  $0 \leq j \leq k$ 

$$\left\|\frac{1}{T-k-s}\sum_{t=k+s+1}^T \left(y_t-\sum_{i\in[s]}\alpha_i y_{t-k-i}\right)u_{t-j}^\top-X_j\right\|_F^2\leqslant\epsilon^2$$

where recall  $X_i$  is defined in Definition 5.14.

Once we have proven Theorem 6.5, we can combine it with a result from [51] for estimating the system parameters from the Markov Parameters using Ho-Kalman. See Section 8 for details. The main technical work of this paper is in proving Theorem 6.5, which we focus on in Section 7.

Algorithm 6.6 (Stabilizing the System).

- **Input:** *T* observations  $\{y_1, y_2, ..., y_T\}$  from a single trajectory and the corresponding inputs  $\{u_1, u_2, ..., u_T\}$  generated from a Linear Dynamical System  $\mathcal{L}(A, B, C, D)$  satisfying the assumptions in Section 4
- **Input:**  $s \in \mathbb{N}$  such that the observability and controllability matrices satisfy the condition number bounds in Definition 4.4, accuracy parameter  $0 < \epsilon < 1$ , and failure probability parameter  $0 < \delta < 1$ .

**Operation:** 

- (1) Let  $P_0, P_1, P_2$  be sufficiently large polynomials in the parameters  $m, n, p, s, k, \kappa$  and K such that  $P_0 \ll P_1 \ll P_2$ . Let  $L = P_2 \log^2(1/\epsilon)/\epsilon^2$ . Let  $S = \{1, 2, ...\}$ .
- (2) Solve the following system in the matrix variables  $\alpha_1, \alpha_2, \ldots \alpha_s \in \mathbb{R}^{m \times m}$ :

$$C_{\alpha} = \begin{cases} \forall j \in [s] \\ \|\alpha_{j}\|_{F}^{2} \leq P_{0} \\ \forall i \in [100snm^{2}K\log(L)] \\ \|y_{iL+k} - \sum_{j \in [s]} \alpha_{j} \cdot y_{iL-j}\|_{2}^{2} \leq P_{1}\log(1/\epsilon) \end{cases} \end{cases}$$
  
**Output:** Matrices  $\alpha_{1}, \alpha_{2}, \dots \alpha_{s}$  obtained above.

#### 7 ANALYSIS OF ALGORITHM 6.6

In this section, we analyze Algorithm 6.6 and prove Theorem 6.5. Since we treat k as fixed throughout this section, we will write  $F_{\alpha}(A)$  for  $F_{\alpha,k}(A)$  (recall Definition 5.11). First, we need the following basic observation.

**Claim 7.1** (Uniform bounds on the control and noise). Let  $P_0$ ,  $P_1$ ,  $P_2$  be sufficiently large polynomials in the parameters in S such that  $P_0 \ll P_1 \ll P_2$ . Let  $L = P_2 \log^2(1/\epsilon)/\epsilon^2$ . With probability  $1 - 0.1\delta$ , the following events all hold:

(1) For all  $i \in [100 snm^2 K \log L]$  and integers  $-10(k + s) \le c \le 10(k + s)$ , we have

$$||u_{iL+c}||, ||w_{iL+c}||, ||z_{iL+c}|| \leq P_0 \log(1/\varepsilon)$$

(2) For all 
$$t \leq P_2^2 \log^3(1/\varepsilon)/\varepsilon^2$$
, we have  
 $\|u_t\|, \|w_t\|, \|z_t\| \leq LP_0$ .

**PROOF.** Note that from the covariance bounds on  $u_t$ ,  $w_t$ ,  $z_t$  (and using Markov's inequality), we have that for a fixed index t and any parameter  $\alpha$ ,

$$\Pr[\|u_t\| \ge \alpha] \le \frac{n}{\alpha^2}$$
$$\Pr[\|w_t\| \ge \alpha] \le \frac{n\sigma_w^2}{\alpha^2}$$
$$\Pr[\|z_t\| \ge \alpha] \le \frac{n\sigma_z^2}{\alpha^2}.$$

Now to prove the first statement, note that we only need to union bound over

$$O(100(k+s)snm^2K\log L) = \text{poly}(S)\log(1/\varepsilon)$$

variables so as long as we choose  $P_0$  sufficiently large we get that the statement holds with at least  $1 - 0.05\delta$  probability. The proof of the second statement is similar except we union bound over more variables.

We begin by establishing the feasibility of the constraint system,  $C_{\alpha}$ , as defined in Algorithm 6.6:

LEMMA 7.2 (FEASIBILITY OF THE CONSTRAINT SYSTEM). Assume we are given  $0 < \epsilon, \delta < 1$ , and  $T = \Omega\left(\operatorname{poly}(\mathcal{S}) \cdot \frac{\log^3(\frac{1}{\epsilon})}{\epsilon^2}\right)$  observations  $\{y_i, u_i\}_{i \in [T]}$  from a single trajectory of a linear dynamical system  $\mathcal{L}$  (A, B, C, D), satisfying the assumptions in Section 4. Then, as long as the events in Claim 7.1 hold, the constraint system  $C_{\alpha}$  is feasible.

Next, we argue that any  $\alpha$  that is feasible for  $C_{\alpha}$  must actually be useful for stabilizing the system. To do this, we introduce the following potential.

**Definition 7.3** (Anti-Concentration Potential). For an integer *l* and coefficients  $\alpha = (\alpha_1, ..., \alpha_s)$ , define the function

$$G_{\alpha,l} = \sum_{i=0}^{l} \|F_{\alpha}(A)A^{i}B\|_{F}^{2}$$

 $G_{\alpha,L}$  is a potential measuring the variance of  $\hat{y}$ . We will show that with high probability over the randomness of the  $u_t, w_t, z_t$ , any  $\alpha$  that is feasible for  $C_{\alpha}$  must have  $G_{\alpha,l}$  be small.

First, we express  $G_{\alpha,l}$  as the variance of a random variable that naturally arises when computing  $\hat{y}_l$  (using the formula in Fact 5.12).

LEMMA 7.4 (POTENTIAL CAPTURES VARIANCE). Given  $l < t \in \mathbb{N}$ , consider the random variable

$$\gamma_{t,l} = \sum_{i=0}^{l} F_{\alpha}(A) A^{i} B u_{t-i}$$

Then we have

$$Tr\left[\mathbb{E}[\gamma_{t,l}\gamma_{t,l}^{\mathsf{T}}]\right] = G_{\alpha,l}$$

PROOF. Since the  $u_i$  are independent  $\mathbb{E}[u_i u_j] = 0$  for  $i \neq j$  and  $\mathbb{E}[u_i u_i^{\top}] = I$  for all *i*, we have

$$\mathbb{E}[\gamma_{t,l}\gamma_{t,l}^{\top}] = \mathbb{E}\left[\left(\sum_{i=0}^{l} F_{\alpha}(A)A^{i}Bu_{t-i}\right)\left(\sum_{i=0}^{l} F_{\alpha}(A)A^{i}Bu_{t-i}\right)^{\top}\right]$$
$$= \mathbb{E}\left[\sum_{i=0}^{l} \left(F_{\alpha}(A)A^{i}Bu_{t-i}\right)\left(F_{\alpha}(A)A^{i}Bu_{t-i}\right)^{\top}\right]$$
$$= \sum_{i=0}^{l} \left(\left(F_{\alpha}(A)A^{i}B\right)\left(\left(F_{\alpha}(A)A^{i}B\right)^{\top}\right)\right)$$

Now taking the trace of both sides, and using the linearity of trace,

$$\operatorname{Tr}\left[\mathbb{E}\left[\gamma_{t,l}\gamma_{t,l}^{\top}\right]\right] = \sum_{i=0}^{l} \operatorname{Tr}\left[\left((F_{\alpha}(A)A^{i}B)((F_{\alpha}(A)A^{i}B)^{\top}\right)\right]$$
$$= \sum_{i=0}^{l} ||F_{\alpha}(A)A^{i}B||_{F}^{2} = G_{\alpha,l}$$

as desired.

Next, we show that if the potential  $G_{\alpha,L}$  is large, with high probability, there must be a violated constraint in  $C_{\alpha}$ .

LEMMA 7.5. Let  $\alpha = (\alpha_1, ..., \alpha_s)$  be a fixed sequence with  $\|\alpha_i\|_F \leq P_0$  and  $G_{\alpha,L} > m(100P_1 \log(1/\varepsilon))^2$ . Now, consider a trajectory of length T, and the corresponding observations and inputs,  $\{y_i, u_i\}_{i \in [T]}$ , from the LDS  $\mathcal{L}(A, B, C, D)$  after fixing this sequence and let  $\zeta$  be the event that Lemma 7.1 holds. Conditioned on  $\zeta$ , with  $1 - (1/L)^{10snm^2}$  probability, there exists some integer  $1 \leq i \leq 10^2 snm^2 K \log L$  such that

$$\|\hat{y}_{iL+k}\| \ge 2P_1 \log(1/\varepsilon) \,.$$

PROOF. See full version.

Next, we relate the potential function  $G_{\alpha,L}$  to the variance of our actual estimators for the Markov parameters. First, we will need to define a similar-looking potential and relate it to  $G_{\alpha,L}$ .

**Definition 7.6.** For an integer *l* and matrix-valued coefficients  $\alpha = (\alpha_1, ..., \alpha_s)$ , define the function

$$H_{\alpha,l} = \sum_{i=0}^{l} \|F_{\alpha}(A)A^i\|_F^2$$

Intuitively, this potential simply does not include the matrix *B* and thus is polynomially related to  $G_{\alpha,l}$ , since the condition number of the controllability matrix is bounded. We make this precise as follows:

LEMMA 7.7 (POTENTIAL WITHOUT B). For a fixed sequences of matrices  $\alpha = (\alpha_1, \alpha_2, ..., \alpha_s)$ , we have

$$H_{\alpha,l} \leq \kappa^2 s G_{\alpha,l+s}.$$

PROOF. See full version.

STOC '23, June 20-23, 2023, Orlando, FL, USA

• •

Next, for any fixed sequence of  $\alpha$ 's we show that the variance of our estimator can be bounded in terms of the norm of the  $\alpha_i$ 's and the two potentials we defined above. In particular,

LEMMA 7.8 (POTENTIAL TO VARIANCE BOUND). For any fixed  $\alpha_1, \ldots, \alpha_s \in \mathbb{R}^{m \times m}$  and any  $0 \leq j \leq k$ , we have

$$\mathbb{E}\left[\left\|\frac{1}{L}\sum_{t=s+1}^{s+L}\hat{y}_{t+k}u_{t+k-j}^{\top} - X_{j}\right\|_{F}^{2}\right]$$
  
$$\leq \frac{P_{1}}{L}\left(1 + \|\alpha_{1}\|_{F}^{2} + \dots + \|\alpha_{s}\|_{F}^{2} + G_{\alpha,L} + H_{\alpha,L}\right).$$

where recall  $X_j$  is defined as in Definition 5.14 and the expectation is over the draws of  $u_t$ ,  $w_t$ ,  $z_t$ .

PROOF. See full version.

Unfortunately, we cannot use Lemma 7.8 directly because the choice of  $\alpha = (\alpha_1, \ldots, \alpha_s)$  that we compute using the program  $C_{\alpha}$  already depends on the realizations of the  $u_t, w_t, z_t$  meaning that there is no more fresh randomness. In order to use the randomness over  $u_t, w_t$  and  $z_t$  to bound the variance of our estimator, we need decouple the  $\alpha$ 's out of the variance expression.

To circumvent this, we will derive a symbolic inequality from Lemma 7.8 that holds simultaneously for all choices of  $\alpha$  and thus can be applied even if  $\alpha$  depends on the realizations of  $u_t, w_t, z_t$ .

Note that  $G_{\alpha,l}$  and  $H_{\alpha,l}$  are both quadratic expressions in the  $\alpha_i$ . It will be useful to extract out the matrix of coefficients, which we do in the following definition.

**Definition 7.9** (Coefficient Matrix of a quadratic polynomial). Define  $v_{\alpha} = (1, \alpha_1, ..., \alpha_s)$  where we view the  $\alpha_i$  as formal variables and flatten each of the matrices  $\alpha_i$  into a vector and concatenate them so that  $v_{\alpha}$  has length  $sm^2 + 1$ . Let  $G_L$  (respectively  $H_L$ ) be the unique symmetric  $(sm^2 + 1) \times (sm^2 + 1)$  matrix such that

$$v_{\alpha}^T G_L v_{\alpha} = G_{\alpha,L} \,.$$

Note that matrices defined above are unique because we force them to be symmetric so the coefficients of the monomials in the  $\alpha_i$  uniquely determine the entries of the matrices. Also, the entries of  $G_l$ ,  $H_l$  are purely functions of the system parameters A, B, C, D. Now we can prove a symbolic version of Lemma 7.8. Intuitively, this symbolic version is a way to decouple the vector valued random variables  $u_t$ ,  $w_t$  and  $z_t$  from the matrix random variables  $\alpha$ , and then use properties of the input distribution on the expressions that are independent of  $\alpha$ .

COROLLARY 7.10 (SYMBOLIC MATRIX INEQUALITY). Consider the vector of formal variables  $v_{\alpha} = (1, \alpha_1, ..., \alpha_s)$ . For an integer j with  $0 \le j \le k$ , define  $M_j$  to be the unique symmetric matrix such that

$$\left\|\frac{1}{L}\sum_{t=s+1}^{s+L}\hat{y}_{t+k}u_{t+k-j}^{\top}-X_j\right\|_F^2=v_{\alpha}^{\top}M_jv_{\alpha},$$

where recall  $X_i$  is defined in Definition 5.14. Then we have

$$\mathbb{E}[M_j] \leq \frac{P_1}{L}(I + G_L + H_L),$$

where the expectation is over the randomness of the realizations of the  $u_i$ ,  $w_i$ ,  $z_i$ .

П

**PROOF.** We simply use the fact that Lemma 7.8 holds over all choices of  $\alpha = (\alpha_1, ..., \alpha_s)$ . See full version for details.

Now we can complete the analysis of our algorithm for learning the Markov parameters.

PROOF OF THEOREM 6.5. With probability at least  $1 - 0.1\delta$ , the event in Claim 7.1 holds and we condition on it. Now we solve the program in Definition 6.3. By Lemma 7.2, it is feasible. Let  $\tilde{\alpha} = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_s)$  be a feasible solution.

Now construct a  $\gamma$ -net, denoted by  $\mathcal{T}$ , over the matrices  $(\alpha_1, \ldots, \alpha_s)$ in Frobenius norm with  $\gamma = 1/L^{4n}$ . Then, observe for any  $(\alpha_1, \ldots, \alpha_s)$ with  $\|\alpha_i\| \leq P_0$ , there exists  $(\alpha'_1, \ldots, \alpha'_s) \in \mathcal{T}$  such that

$$\sqrt{\|\alpha_1 - \alpha_1'\|_F^2 + \dots + \|\alpha_s - \alpha_s'\|_F^2} \leq \frac{1}{L^{4n}}$$

It is clear that such a net exists with  $|\mathcal{T}| \leq L^{4nsm^2}$ . Now, we can union bound over all the events such that with probability at least  $1 - 0.2\delta$ , simultaneously, for all  $\alpha$ , for all  $i \in [100snm^2K \log(L)]$ , it follows from Lemma 7.5 that  $\|\hat{y}_{iL+k}\|^2 \leq 2P_1 \log(1/\epsilon)$ .

Next, we show that the solution  $\hat{\alpha}$  we obtain must satisfy

$$G_{\hat{\alpha},L} \leq m(200P_1\log(1/\varepsilon))^2 \,. \tag{9}$$

To see this, assume for the sake of contradiction that the above doesn't hold. Now round our solution  $\hat{\alpha}$  to the nearest  $\alpha' = (\alpha'_1, \ldots, \alpha'_s)$  in the net. By Lemma 5.6, as long as  $P_2$  (and recall  $L = P_2 \log^2(1/\varepsilon)/\varepsilon^2$ ) is chosen sufficiently large, we have

$$G_{\alpha',L} \ge m(100P_1\log(1/\varepsilon))^2$$

However, Lemma 7.5 implies that there is some integer  $1 \le i \le 10^2 snm^2 K \log L$  such that

$$\|y_{iL+k} - \alpha'_1 y_{iL-1} - \cdots - \alpha'_s y_{iL-s}\| \ge 2P_1 \log(1/\varepsilon).$$

However, using the assumptions in Lemma 7.1, the formula for  $y_t$  in Fact 5.12 and the bounds in Lemma 5.6, and the properties of the net, the above implies that

$$\|y_{iL+k} - \hat{\alpha}_1 y_{iL-1} - \dots - \hat{\alpha}_s y_{iL-s}\| \ge P_1 \log(1/\varepsilon)$$

which contradicts the fact that  $(\hat{\alpha}_1, \dots, \hat{\alpha}_s)$  is a feasible solution. Thus, we actually must have (9). Now by Lemma 7.7, we have  $H_{\alpha,L-s} \leq \kappa^2 sm(200P_1 \log(1/\varepsilon))^2$ . Now let L' = L - s and set

$$\hat{X}_{j} = \frac{1}{L'} \sum_{t=s+1}^{s+L'} \hat{y}_{t+k} u_{t+k-j}^{\top}$$

for all  $0 \le j \le k$ . Let  $M_j$  be defined as in Corollary 7.10 (with L replaced by L'). Let

$$\widetilde{M_j} = \frac{P_1}{L'} (I + G_{L'} + H_{L'}) \,. \label{eq:mass_matrix}$$

Note that  $\widetilde{M_j}$  is clearly PSD. Also  $M_j$  is always PSD regardless of the realizations of the  $u_i, w_i, z_i$ , so by Markov's inequality and Corollary 7.10, we have that with  $1 - 0.1\delta/k$  probability

$$\operatorname{Tr}\left(\widetilde{M_{j}}^{-1/2}M_{j}\widetilde{M_{j}}^{-1/2}\right) \leq \frac{10k}{\delta} \mathbb{E}\left[\operatorname{Tr}\left(\widetilde{M_{j}}^{-1/2}M_{j}\widetilde{M_{j}}^{-1/2}\right)\right]$$
$$\leq \frac{10k}{\delta}(sm^{2}+1).$$

Ainesh Bakshi, Allen Liu, Ankur Moitra, and Morris Yau

This means that with  $1 - 0.1\delta/k$  probability,

$$\widetilde{M_j}^{-1/2} M_j \widetilde{M_j}^{-1/2} \le \frac{10k(sm^2 + 1)}{\delta}$$

which then implies

$$M_j \leq \frac{10k(sm^2+1)}{\delta}\widetilde{M_j} \,.$$

Assuming that this happens, we have that

$$\begin{aligned} \|\hat{X}_{j} - X_{j}\|_{F}^{2} &= v_{\alpha}^{\top} M_{j} v_{\alpha} \\ &\leq \frac{10k(sm^{2} + 1)}{\delta} \frac{P_{1}}{L'} (1 + \|\alpha_{1}\|_{F}^{2} + \dots + \|\alpha_{s}\|_{F}^{2} + G_{\alpha,L'} + H_{\alpha,L'}) \\ &\leq \varepsilon^{2} \end{aligned}$$

$$(10)$$

where the last inequality uses that  $L = P_2(\log(1/\varepsilon))^2/\varepsilon^2$  and  $P_2$  is chosen sufficiently large. Finally, union bounding the above over all choices of j = 0, 1, ..., k completes the proof for the guarantees of the estimator.

Finally, note that our algorithm runs in polynomial time in all the parameters because the convex program  $C_{\alpha}$  admits an efficient separation oracle. To see this, note that there are only polynomially many constraints in  $C_{\alpha}$  and each one is either a linear constraint or an ellipsoid constraint both of which admit an efficient separation oracle.

#### 8 FROM MARKOV PARAMETERS TO SYSTEM PARAMETERS

Note that Theorem 6.5 guarantees that we can get good estimates for the markov parameters. To complete the proof of our full learning result, Theorem 6.4, we apply the Ho-Kalman algorithm black box to extract the system matrices  $\{A, B, C, D\}$  where  $\{A, B, C\}$  are recovered up to a similarity transformation. Recall that linear dynamical systems are specified only up to similarity transformation (see [51] for a discussion on this point).

The following lemma from [51] establishes error guarantees for the Ho-Kalman algorithm given operator norm bounds on estimating G.

LEMMA 8.1 ([51]). For observability and controllability matrices that are rank n, the Ho-Kalman algorithm applied to  $\hat{G}$  produces estimates  $\hat{A}$ ,  $\hat{B}$ , and  $\hat{C}$  such that there exists similarity transform  $T \in \mathbb{R}^{n \times n}$  such that

 $\max\{\|C - \hat{C}T\|_F, \|B - T^{-1}\hat{B}\|_F\} \le 5\sqrt{n\|G - \hat{G}\|}$ 

$$||A - T^{-1}\hat{A}T||_F \leq \frac{\sqrt{n||G - \hat{G}||}||H||}{\sigma_{min}^{3/2}(H^-)}$$

and

and

$$\left\|D - \hat{D}\right\|_{F} \leq \sqrt{n} \left\|G - \hat{G}\right\|$$

where in the above

$$G = [D, CB, CAB, \dots, CA^{2s-1}B]$$

A straightforward application of this lemma allows us to complete the proof of Theorem 6.4.

PROOF OF THEOREM 6.4. By Theorem 6.5, in Algorithm 6.1, the input  $\hat{G}$  to the Ho-Kalman algorithm satisfies

$$\|G - \hat{G}\|_F \leq \sqrt{2s + 1}\varepsilon$$

with probability at least  $1 - \delta$ . Now we apply Lemma 8.1. The two things we need to do are upper bound ||H|| and lower bound  $\sigma_{min}(H^-)$ . We have an upper bound on  $||H|| \leq \sigma_{max}(O_s)\sigma_{max}(Q_s) \leq \kappa^2 s ||B|| ||C||$  where we use Claim 5.15. We also have  $\sigma_{min}(H^-) \geq \sigma_{min}(O_s)\sigma_{min}(Q_s) \geq ||B|| ||C|| \geq 1$ . Therefore we conclude that there is a similarity transform *T* such that

$$\max\{\|A - T^{-1}\hat{A}T\|_{F}, \|B - \hat{B}T\|_{F}, \|C - T^{-1}\hat{C}\|_{F}\} \le \operatorname{poly}(n, \kappa, s, \|B\|, \|C\|)\sqrt{\epsilon}.$$

Redefining  $\varepsilon$  appropriately immediately gives the desired result.  $\Box$ 

#### 9 SAMPLE COMPLEXITY LOWER BOUND FOR ILL-CONDITIONAL LDS

In this section, we prove a lower bound, that when the observability or controllability matrix of an LDS is close to singular, then it is information-theoretically impossible to learn. We consider the case where the distributions  $\mathcal{D}_u = \mathcal{D}_0 = N(0, I)$  and  $\mathcal{D}_w = N(0, \Sigma_w)$ where  $\Sigma_w$  will be set later. For simplicity, we also set  $x_0 = 0$  and also D = 0.

**Definition 9.1.** We say that an LDS  $\mathcal{L}(A, B, C, D)$  is  $(\delta, v)$ -unobservable if v is a unit vector such that for all integers  $s \ge 0$ ,

 $\|CA^sv\| \leq \delta.$ 

Note that the above condition depends only on *A*, *C* so we will sometimes talk about a pair of matrices *A*, *C* being  $(\delta, v)$ -unobservable.

**Definition 9.2.** We say that an LDS  $\mathcal{L}(A, B, C, D)$  is  $(\delta, v)$ -uncontrollable if v is a unit vector such that for all integers  $s \ge 0$ ,

 $\|(A^s B)^\top v\| \leq \delta.$ 

Note that the above condition depends only on A, B so we will sometimes talk about a pair of matrices A, B being  $(\delta, v)$ -unobservable.

For our formal lower bound, we need a minor assumption that *A* is generic, in particular, we need that it is not too close to a multiple of the identity plus a rank-1 perturbation. Essentially all matrices satisfy this assumption as long as  $n \ge 3$ .

**Definition 9.3.** We say a matrix  $A \in \mathbb{R}^{n \times n}$  and vector v are *c*-generic if  $||A|| \ge c$  and there are unit vectors u, w such that

THEOREM 9.4. Let A, C be matrices that are  $(\delta, v)$ -unobservable for some  $0 < \delta < 0.1$  and unit vector v. Assume that (A, v) are cgeneric for some constant c. Then any algorithm that is given an LDS  $\mathcal{L}(A, B, C, D)$  that uses at most

$$o\left(\frac{1}{\sqrt{\delta(m+p)}}\right)$$

samples has probability at least 0.4 of outputting  $\hat{A}$ ,  $\hat{B}$ ,  $\hat{C}$ ,  $\hat{D}$  such that there is no invertible matrix U with

$$||A - U^{-1}\hat{A}U||_F, ||B - U^{-1}\hat{B}||_F, ||C - \hat{C}U||_F \le 0.1c^2$$

STOC '23, June 20-23, 2023, Orlando, FL, USA

Similarly, the same holds if A, B are matrices that are  $(\delta, v)$  uncontrollable.

PROOF. See full version.

#### ACKNOWLEDGMENTS

Ainesh Bakshi was supported by Ankur Moitra's ONR grant. Allen Liu was supported by an NSF Graduate Research Fellowship and a Fannie and John Hertz Foundation Fellowship. Ankur Moitra was supported by an ONR grant and a David and Lucile Packard Fellowship.

#### REFERENCES

- Karl Johan Åström and Peter Eykhoff. 1971. System identification—a survey. Automatica 7, 2 (1971), 123–162.
- [2] Michael Athans. 1974. The importance of Kalman filtering methods for economic systems. In Annals of Economic and Social Measurement, Volume 3, number 1. NBER, 49–64.
- [3] Ainesh Bakshi and Pravesh Kothari. 2020. List-Decodable Subspace Recovery via Sum-of-Squares. arXiv preprint arXiv:2002.05139 (2020).
- [4] Ainesh Bakshi and Pravesh Kothari. 2020. Outlier-Robust Clustering of Non-Spherical Mixtures. arXiv preprint arXiv:2005.02970 (2020).
- [5] Ainesh Bakshi and Adarsh Prasad. 2021. Robust linear regression: Optimal rates in polynomial time. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. 102–115.
- [6] Guy Bresler. 2015. Efficiently learning Ising models on arbitrary graphs. In Proceedings of the forty-seventh annual ACM symposium on Theory of computing. 771–782.
- [7] Emmanuel J Candès and Carlos Fernandez-Granda. 2014. Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics* 67, 6 (2014), 906–956.
- [8] Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. 2022. Kalman filtering with adversarial corruptions. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. 832–845.
- [9] Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. 2022. Online and distribution-free robustness: Regression and contextual bandits with huber contamination. In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 684–695.
- [10] Yeshwanth Cherapanamjeri, Efe Aras, Nilesh Tripuraneni, Michael I Jordan, Nicolas Flammarion, and Peter L Bartlett. 2020. Optimal robust linear regression in nearly linear time. arXiv preprint arXiv:2007.08137 (2020).
- [11] Yeshwanth Cherapanamjeri, Samuel B Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesh Tripuraneni. 2020. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. 601–609.
- [12] Yeshwanth Cherapanamjeri, Sidhanth Mohanty, and Morris Yau. 2020. List Decodable Mean Estimation in Nearly Linear Time. arXiv preprint arXiv:2005.09796 (2020).
- [13] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. 2020. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics* 20, 4 (2020), 633–679.
- [14] Ilias Diakonikolas, Samuel B Hopkins, Daniel Kane, and Sushrut Karmalkar. 2020. Robustly Learning any Clusterable Mixture of Gaussians. arXiv preprint arXiv:2005.06417 (2020).
- [15] Feng Ding. 2013. Two-stage least squares based iterative estimation algorithm for CARARMA system modeling. *Applied Mathematical Modelling* 37, 7 (2013), 4798–4808. https://doi.org/10.1016/j.apm.2012.10.014
- [16] Boualem Djehiche and Othmane Mazhar. 2022. Efficient learning of hidden state LTI state space models of unknown order. https://doi.org/10.48550/ARXIV.2202. 01625
- [17] David L Donoho. 1992. Superresolution via sparsity constraints. SIAM journal on mathematical analysis 23, 5 (1992), 1309–1331.
- [18] Gianfranco Doretto, Alessandro Chiuso, Ying Nian Wu, and Stefano Soatto. 2003. Dynamic textures. International Journal of Computer Vision 51, 2 (2003), 91–109.
- [19] Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. 2017. Finite Time Identification in Unstable Linear Systems. *CoRR* abs/1710.01852 (2017). arXiv:1710.01852 http://arxiv.org/abs/1710.01852
- [20] Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. 2018. Finite time identification in unstable linear systems. *Automatica* 96 (2018), 342–353.
- [21] Salar Fattahi. 2020. Learning Partially Observed Linear Dynamical Systems from Logarithmic Number of Samples. CoRR abs/2010.04015 (2020). arXiv:2010.04015 https://arxiv.org/abs/2010.04015

- [22] Erwin Enrique Fetzer and PM Anderson. 1975. Observability in the state estimation of power systems. *IEEE transactions on power Apparatus and Systems* 94, 6 (1975), 1981–1988.
- [23] Miguel Galrinho. 2016. Least Squares Methods for System Identification of Structured Models.
- [24] Zoubin Ghahramani and Geoffrey E Hinton. 1996. Parameter estimation for linear dynamical systems. (1996).
- [25] Udaya Ghai, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. 2020. No-Regret Prediction in Marginally Stable Systems. CoRR abs/2002.02064 (2020). arXiv:2002.02064 https://arxiv.org/abs/2002.02064
- [26] Mohinder S Grewal and Angus P Andrews. 2010. Applications of Kalman filtering in aerospace 1960 to the present [historical perspectives]. *IEEE Control Systems Magazine* 30, 3 (2010), 69–78.
- [27] Moritz Hardt, Tengyu Ma, and Benjamin Recht. 2018. Gradient Descent Learns Linear Dynamical Systems. Journal of Machine Learning Research 19 (2018), 1–44.
- [28] Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. 2018. Spectral filtering for general linear dynamical systems. Advances in Neural Information Processing Systems 31 (2018).
- [29] Elad Hazan, Karan Singh, and Cyril Zhang. 2017. Learning Linear Dynamical Systems via Spectral Filtering. CoRR abs/1711.00946 (2017). arXiv:1711.00946 http://arxiv.org/abs/1711.00946
- [30] BL HO and Rudolf E Kálmán. 1966. Effective construction of linear state-variable models from input/output functions. at-Automatisierungstechnik 14, 1-12 (1966), 545-548.
- [31] Samuel B Hopkins. 2018. Sub-Gaussian Mean Estimation in Polynomial Time. arXiv preprint arXiv:1809.07425 (2018).
- [32] Samuel B Hopkins and Jerry Li. 2018. Mixture models, robustness, and sum of squares proofs. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing. 1021–1034.
- [33] Misha Ivkov and Pravesh K Kothari. 2022. List-decodable covariance estimation. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. 1276–1283.
- [34] Arun Jambulapati, Jerry Li, Tselil Schramm, and Kevin Tian. 2021. Robust regression revisited: Acceleration and improved estimation rates. Advances in Neural Information Processing Systems 34 (2021), 4475–4488.
- [35] Rudolph Emil Kalman. 1960. A new approach to linear filtering and prediction problems. (1960).
- [36] Rudolf E Kalman. 1960. On the general theory of control systems. In Proceedings First International Conference on Automatic Control, Moscow, USSR. 481–492.
- [37] Sushrut Karmalkar, Adam Klivans, and Pravesh Kothari. 2019. List-decodable linear regression. In Advances in Neural Information Processing Systems. 7423– 7432.
- [38] Jonathan Kelner, Frederic Koehler, Raghu Meka, and Ankur Moitra. 2020. Learning some popular gaussian graphical models without condition number bounds. Advances in Neural Information Processing Systems 33 (2020), 10986–10998.
- [39] Adam Klivans, Pravesh K Kothari, and Raghu Meka. 2018. Efficient Algorithms for Outlier-Robust Regression. arXiv preprint arXiv:1803.03241 (2018).
- [40] Pravesh K Kothari, Jacob Steinhardt, and David Steurer. 2018. Robust moment estimation and improved clustering via sum of squares. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing. ACM, 1035–1046.
- [41] Pravesh K Kothari and David Steurer. 2017. Outlier-robust moment-estimation via sum-of-squares. arXiv preprint arXiv:1711.11581 (2017).
- [42] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. 2020. Regret Bound of Adaptive Control in Linear Quadratic Gaussian (LQG) Systems. *CoRR* abs/2003.05999 (2020). arXiv:2003.05999 https://arxiv.org/abs/ 2003.05999
- [43] Bruce Lee and Andrew Lamperski. 2019. Non-asymptotic Closed-Loop System Identification using Autoregressive Processes and Hankel Model Reduction. https://doi.org/10.48550/ARXIV.1909.02192
- [44] Holden Lee. 2020. Improved rates for identification of partially observed linear dynamical systems. *CoRR* abs/2011.10006 (2020). arXiv:2011.10006 https://arxiv. org/abs/2011.10006
- [45] Qiao Li, Roger G Mark, and Gari D Clifford. 2007. Robust heart rate estimation from multiple asynchronous noisy sources using signal quality indices and a Kalman filter. *Physiological measurement* 29, 1 (2007), 15.
- [46] Lennart Ljung. 1998. System identification. In Signal analysis and prediction. Springer, 163–173.
- [47] Gábor Lugosi and Shahar Mendelson. 2019. Sub-Gaussian estimators of the mean of a random vector. *The annals of statistics* 47, 2 (2019), 783–794.
- [48] Bertrand Mesot and David Barber. 2007. Switching linear dynamical systems for noise robust speech recognition. *IEEE Transactions on Audio, Speech, and Language Processing* 15, 6 (2007), 1850–1858.
- [49] Ankur Moitra. 2015. Super-resolution, extremal functions and the condition number of Vandermonde matrices. In Proceedings of the forty-seventh annual

ACM symposium on Theory of computing. 821–830.

- [50] PC Müller and HI Weber. 1972. Analysis and optimization of certain qualities of controllability and observability for linear dynamical systems. *Automatica* 8, 3 (1972), 237–246.
- [51] Samet Oymak and Necmiye Ozay. 2019. Non-asymptotic identification of lti systems from a single trajectory. In 2019 American control conference (ACC). IEEE, 5655–5661.
- [52] Ankit Pensia, Varun Jog, and Po-Ling Loh. 2020. Robust regression with covariate filtering: Heavy tails and adversarial contamination. arXiv preprint arXiv:2009.12976 (2020).
- [53] Adarsh Prasad, Arun Sai Suggala, Sivaraman Balakrishnan, and Pradeep Ravikumar. 2020. Robust estimation via robust gradient estimation. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 82, 3 (2020), 601–627. https://doi.org/10.1111/rssb.12364 arXiv:https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/rssb.12364
- [54] Prasad Raghavendra and Morris Yau. 2020. List decodable learning via sum of squares. In Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, 161–180.
- [55] Prasad Raghavendra and Morris Yau. 2020. List decodable subspace recovery. In Conference on Learning Theory. PMLR, 3206–3226.
- [56] Paria Rashidinejad, Jiantao Jiao, and Stuart Russell. 2020. SLIP: Learning to predict in unknown dynamical systems with long-term memory. In Advances in Neural Information Processing Systems, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 5716–5728. https://proceedings. neurips.cc/paper/2020/file/3e91970f771a2c473ae36b60d1146068-Paper.pdf
- [57] Derek Rowell. 2002. State-space representation of LTI systems. URL: http://web. mit. edu/2.14/www/Handouts/StateSpace. pdf (2002), 1–18.
- [58] Tuhin Sarkar and Alexander Rakhlin. 2019. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*. PMLR, 5610–5618.
- [59] Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. 2019. Nonparametric finite time LTI system identification. arXiv preprint arXiv:1902.01848 (2019).
- [60] Tuhin Sarkar, Alexander Rakhlin, and Munther A. Dahleh. 2022. Finite Time LTI System Identification. J. Mach. Learn. Res. 22, 1, Article 26 (jul 2022), 61 pages.
- [61] Steven J Schiff. 2009. Kalman meets neuron: the emerging intersection of control theory with neuroscience. In 2009 annual international conference of the IEEE engineering in medicine and biology society. IEEE, 3318–3321.
- [62] Parikshit Shah, Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. 2012. Linear system identification via atomic norm regularization. In 2012 IEEE 51st IEEE conference on decision and control (CDC). IEEE, 6265–6270.
- [63] Max Simchowitz, Ross Boczar, and Benjamin Recht. 2019. Learning linear dynamical systems with semi-parametric least squares. In *Conference on Learning Theory*. PMLR, 2714–2802.
- [64] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. 2018. Learning without mixing: Towards a sharp analysis of linear system identification. In Conference On Learning Theory. PMLR, 439–473.
- [65] Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. 2018. Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification. *CoRR* abs/1802.08334 (2018). arXiv:1802.08334 http: //arxiv.org/abs/1802.08334
- [66] W. Spinelli, L. Piroddi, and M. Lovera. 2005. On the role of prefiltering in nonlinear system identification. *IEEE Trans. Automat. Control* 50, 10 (2005), 1597–1602. https://doi.org/10.1109/TAC.2005.856655
- [67] Yue Sun, Samet Oymak, and Maryam Fazel. 2022. System Identification via Nuclear Norm Regularization. https://doi.org/10.48550/ARXIV.2203.16673
- [68] Anastasios Tsiamis and George J Pappas. 2019. Finite sample analysis of stochastic system identification. In 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 3648–3654.
- [69] Anastasios Tsiamis, Ingvar Ziemann, Nikolai Matni, and George J. Pappas. 2022. Statistical Learning Theory for Control: A Finite Sample Perspective. https: //doi.org/10.48550/ARXIV.2209.05423
- [70] Yong Zhang. 2011. Unbiased identification of a class of multi-input singleoutput systems with correlated disturbances using bias compensation methods. Mathematical and Computer Modelling 53, 9 (2011), 1810–1819. https: //doi.org/10.1016/j.mcm.2010.12.059
- [71] Yang Zheng and Na Li. 2021. Non-Asymptotic Identification of Linear Dynamical Systems Using Multiple Trajectories. *IEEE Control Systems Letters* 5, 5 (nov 2021), 1693–1698. https://doi.org/10.1109/lcsys.2020.3042924
- [72] Banghua Zhu, Jiantao Jiao, and Jacob Steinhardt. 2020. Robust estimation via generalized quasi-gradients. arXiv preprint arXiv:2005.14073 (2020).

Received 2022-11-07; accepted 2023-02-06