

Learning Linearized Models from Nonlinear Systems with Finite Data

Lei Xin, George Chiu, Shreyas Sundaram

Abstract—Identifying a linear system model from data has wide applications in control theory. The existing work on finite sample analysis for linear system identification typically uses data from a single system trajectory under i.i.d random inputs, and assumes that the underlying dynamics is truly linear. In contrast, we consider the problem of identifying a linearized model when the true underlying dynamics is nonlinear. We provide a multiple trajectories-based deterministic data acquisition algorithm followed by a regularized least squares algorithm, and provide a finite sample error bound on the learned linearized dynamics. Our error bound demonstrates a trade-off between the error due to nonlinearity and the error due to noise, and shows that one can learn the linearized dynamics with arbitrarily small error given sufficiently many samples. We validate our results through experiments, where we also show the potential insufficiency of linear system identification using a single trajectory with i.i.d random inputs, when nonlinearity does exist.

I. INTRODUCTION

Learning good predictive models from data has wide applications, including in economics and machine learning [1], [2]. The problem of system identification is to learn a mathematical model of a dynamical system from data. System identification is an important problem in control theory since a good model can facilitate model-based control design [3]. Although physical systems are typically nonlinear, linear models are frequently used in practice due to their simplicity [4], and their ability to approximate nonlinear systems around a given reference point. Consequently, it is of interest to understand identification of appropriate linear models from data generated by nonlinear systems.

Classically, theories for system identification typically focus on asymptotic aspects [5], [6]. In recent years, however, finite sample analysis for system identification has been studied extensively. For linear system identification, existing works are either multiple trajectories-based or single trajectory-based. The multiple trajectories setup [7]–[10] requires the user to restart the system multiple times, but has a major advantage in its ability to handle unstable systems. In contrast, the single trajectory setup [11]–[16] performs system identification using data from a single experiment, i.e., the system does not need to reset, but has potential risks if the system is unstable. We note that when it comes to linear system identification, almost all existing works that have finite sample guarantees assume that the underlying system is

truly linear, except for [17]. Furthermore, Gaussian random inputs are typically applied to ensure persistent excitation.

The study on nonlinear system identification is less well-understood, in general, as compared to the case for linear system identification. Recent works on finite sample analysis for nonlinear system identification include [18]–[20]. It is worth noting that to obtain finite sample guarantees, the existing works on nonlinear system identification typically require that a certain model structure to be known in advance. However, when the specific model structure is unknown, a reasonable alternative goal is to learn a linearized model from the nonlinear system, due to the well-studied techniques on linear system control as discussed above.

There is a branch of recent research that focuses on learning a linear system representation that completely captures the behaviours of a nonlinear system using the Koopman Operator [21]. This approach typically requires carefully selected basis functions (e.g., leveraging neural networks [22]), and the analysis focuses on the noiseless setting. In contrast, our focus in this work is to learn a linearized system model, in the sense that the linear model captures the linear part of the nonlinear system after Taylor expansion, and to provide finite sample guarantees when the system has noise.

Most relevant to our work is the recent paper [17], which provides a finite sample error bound for learning linear models from systems that have unmodeled dynamics that could capture nonlinearities, using a single system trajectory. However, the method proposed in [17] assumes the system is “well-behaved” by requiring the unmodeled dynamics/nonlinear terms to be (globally) Lipschitz [23]. The method also requires the system to satisfy certain additional properties to ensure consistent estimation, supposing the inputs are carefully chosen. In contrast, we show in this work that one can learn a linearized system model from a nonlinear system with arbitrarily small error without the Lipschitzness assumption, given sufficiently many short trajectories, supposing that one has control over the initial conditions of the experiments.

In summary, our contributions are as follows.

- We provide a deterministic, multiple trajectories-based data acquisition algorithm that ensures persistent excitation under the constraint of being close to the reference point. Using this algorithm followed by a regularized least squares estimation algorithm, we provide a finite sample error bound of the identified linearized dynamics of a nonlinear system.
- Our bound shows that one can learn the linearized dynamics with arbitrarily small error, given sufficiently many experiments in the multiple trajectories setup, and

This work was partially supported by a contract from Saab Inc., under DARPA Award N65236-23-C-8012. Lei Xin and Shreyas Sundaram are with the Elmore Family School of Electrical and Computer Engineering, Purdue University. George Chiu is with the School of Mechanical Engineering, Purdue University. E-mails: {lxin, gchiu, sundara2}@purdue.edu.

demonstrates a trade-off between the error due to noise and the error due to nonlinearity. The bound further characterizes the benefits of using regularization. When the system is perfectly linear, we show a learning rate that matches the existing results on learning perfectly linear systems using random inputs.

- We provide numerical experiments to validate our results and insights, and show the potential insufficiency of linear system identification using random inputs from a single trajectory when nonlinearity does exist.

Our paper is organized as follows. Section II introduces relevant mathematical notation. Section III introduces the system identification problem and the algorithms we use. In Section IV, we present our theoretical results. We present numerical examples in Section V to validate our results, and conclude in Section VI.

II. NOTATION

Vectors are taken to be column vectors unless indicated otherwise. Let \mathbb{R} and \mathbb{Z} denote the set of real numbers and integers, respectively. Let $\lambda_{max}(\cdot)$ and $\lambda_{min}(\cdot)$ be the largest and the smallest eigenvalue in magnitude, respectively, of a given matrix. For a given matrix A , we use A' to denote its conjugate transpose. We use $\|A\|$, $\|A\|_1$ and $\|A\|_F$ to denote the spectral norm, 1-norm, and Frobenius norm, respectively, of matrix A . We use I_n to denote the identity matrix with dimension n . We use the symbol \bmod to denote the modulo operation. The union of sets is denoted as \cup . The open l_1 ball in d -dimensional space with center at x_0 and radius r is denoted by $\mathcal{B}_d(x_0, r) \triangleq \{x \in \mathbb{R}^d : \|x - x_0\|_1 < r\}$. We denote e_i^d as a d -dimensional vector with the i -th component equal to 1 and all other components equal to 0. The symbols $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are used to denote the floor and ceiling functions, respectively. We use $\mathbf{0}$ to denote a zero vector with dimension that is clear from the context. The symbol $\sigma(\cdot)$ is used to denote the sigma field generated by the corresponding random vectors. The symbol \mathcal{S}^{n-1} is used to denote the unit sphere in n -dimensional space.

III. PROBLEM FORMULATION AND SYSTEM IDENTIFICATION ALGORITHM

Consider the following discrete time nonlinear time invariant system

$$x_{k+1} = f(z_k) + w_k, \quad (1)$$

where $f : \mathbb{R}^{n+p} \rightarrow \mathbb{R}^n$, $z_k = [x_k' \ u_k']' \in \mathbb{R}^{n+p}$, $x_k \in \mathbb{R}^n$, $u_k \in \mathbb{R}^p$, and $w_k \in \mathbb{R}^n$. Here, x_k , u_k and w_k are the state, input, and process noise, respectively. The noise terms w_k are assumed to be independent sub-Gaussian random vectors with parameter σ_w^2 , where the definition is given below [24].

Definition 1: A real-valued random variable w is called sub-Gaussian with parameter σ^2 if we have

$$\forall \alpha \in \mathbb{R}, \mathbb{E}[\exp(\alpha w)] \leq \exp\left(\frac{\alpha^2 \sigma^2}{2}\right).$$

A random vector $x \in \mathbb{R}^n$ is called σ^2 sub-Gaussian if for all unit vectors $v \in \mathcal{S}^{n-1}$ the random variable $v'x$ is σ^2 sub-Gaussian. \square

Note that sub-Gaussian distributions are commonly used to model noise processes [16]. In particular, every (zero-mean) Gaussian random vector is sub-Gaussian.

Assume that for each component function of f , all second order partial derivatives exist and are continuous on \mathbb{R}^{n+p} . From Taylor's theorem [25], system (1) using reference point $z_k = \mathbf{0}$ can be rewritten as

$$x_{k+1} = Ax_k + Bu_k + o + w_k + r_k, \quad (2)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times p}$, are system matrices that capture the linear part of $f(z_k)$, $o = f(\mathbf{0}) \in \mathbb{R}^n$, and $r_k = h(z_k) \in \mathbb{R}^n$ is a remainder vector that contains higher order terms that are state/input dependent, where $h : \mathbb{R}^{n+p} \rightarrow \mathbb{R}^n$. The above model is less studied in the literature on finite sample analysis for system identification, and we will consider this model in the sequel. Note that we assume o is possibly non-zero to capture scenarios where the equilibrium points of the system are unknown. When the system is perfectly linear, we have $o = r_k = \mathbf{0}$, which is the commonly used model in the literature. In this paper, we assume that both the state x_k and input u_k can be perfectly measured. Suppose that we can restart the system multiple times from an arbitrary initial state x_0 using arbitrary input u_0 , and obtain multiple length 1 trajectories (i.e., state-input pairs obtained by running the system for a single time step, as will be explained next). Using superscript to denote the trajectory index, we denote the set of samples we have as $\{(x_1^i, x_0^i, u_0^i) : 1 \leq i \leq N\}$. Our goal is to learn the linear approximation system matrices $\Theta \triangleq [A \ B \ o] \in \mathbb{R}^{n \times (n+p+1)}$ in system (2) from the set of samples available to us.

Our result will leverage the following mild assumption on the remainder vector $r_k = h(z_k)$ in system (2).

Assumption 1: Let $r_{i,k}$ denote the i -th component of r_k . There exist $c > 0$ and $\beta = \beta(c)$ such that $|r_{i,k}| \leq \beta \|z_k\|_1^2$ for all $i \in \{1, \dots, n\}$ and all $z_k \in \mathcal{B}_{n+p}(\mathbf{0}, c)$.

Remark 1: The above assumption is, in fact, a direct result of assuming that each component function of the original nonlinear dynamics f has all second order partial derivatives being continuous on \mathbb{R}^{n+p} , due to Taylor's theorem for multivariable functions from [26, Corollary 1]. Intuitively, this assumption says that the higher order terms are dominated by the second order terms, if the arguments of the function are sufficiently close to the origin. Note that it does not require the function h to be Lipschitz (which is the assumption used in [17]). As an example, consider a scalar system with the dynamics given by $f(z_k) = x_k + u_k + x_k^2 + x_k^3$. Here $r_k = x_k^2 + x_k^3$ satisfies Assumption 1 for $c = 1$ and $\beta = 2$ since $|x_k^2 + x_k^3| \leq |x_k^2| + |x_k^3| \leq 2|x_k|^2 \leq 2\|z_k\|_1^2$ for all $z_k \in \mathcal{B}_2(\mathbf{0}, 1)$, but the corresponding function h is not Lipschitz. In general, a larger c may lead to a larger β .

Let $q > 0$ be a (small) design parameter that constrains the magnitude of the initial conditions z_0 , and N be the number of experiments to perform. We deploy a data collection scheme specified in Algorithm 1.

Algorithm 1 Data Acquisition

Input Norm constraint parameter $q > 0$, number of experiments $N > 0$

```
1: Initialize  $s_1 = 1$ 
2: for  $i = 1, \dots, N$  do
3:   if  $i \bmod (n+p) \neq 0$  then
4:     Set  $z_0^i = [x_0^{i'} \ u_0^{i'}]^\top = s_i \times q e_{i \bmod (n+p)}^{n+p}$ 
5:     Collect  $x_1^i$ , where  $x_1^i = Ax_0^i + Bu_0^i + w_0^i + o + r_0^i$ 
6:     Set  $s_{i+1} = s_i$ 
7:   else
8:     Set  $z_0^i = [x_0^{i'} \ u_0^{i'}]^\top = s_i \times q e_{n+p}^{n+p}$ 
9:     Collect  $x_1^i$ , where  $x_1^i = Ax_0^i + Bu_0^i + w_0^i + o + r_0^i$ 
10:    Set  $s_{i+1} = -s_i$ 
11:   end if
12: end for
13: Output  $\{(x_1^i, x_0^i, u_0^i) : 1 \leq i \leq N\}$ 
```

Remark 2: Intuitively, we want the initial conditions to stay as close to the reference point (in this case, the origin) as possible, to avoid excessive bias from the higher order terms. Hence, the reason of using of multiple length 1 trajectories is to prevent the noise from driving the system too far from the reference point, and amplifying the effects from r_k . The key idea of Algorithm 1 is to ensure persistent excitation (i.e., the smallest eigenvalue of the sample covariance matrix becomes larger as one gets more data), subject to the constraint on bounded distance to the origin (specified by q). Later on in our theoretical result, we will demonstrate how q will affect the finite sample estimation error bound for learning Θ .

We establish some definitions now. Define the batch matrices

$$\begin{aligned} X &= [x_1^1 \ x_1^2 \ \dots \ x_1^N] \in \mathbb{R}^{n \times N} \\ W &= [w_0^1 \ w_0^2 \ \dots \ w_0^N] \in \mathbb{R}^{n \times N} \\ R &= [r_0^1 \ r_0^2 \ \dots \ r_0^N] \in \mathbb{R}^{n \times N}. \end{aligned} \quad (3)$$

Let $\hat{z}_0^i = [z_0^{i'} \ 1]^\top \in \mathbb{R}^{n+p+1}$. Define the regressor matrix

$$Z = [\hat{z}_0^1 \ \hat{z}_0^2 \ \dots \ \hat{z}_0^N] \in \mathbb{R}^{(n+p+1) \times N}. \quad (4)$$

We have the following relationship

$$X = \Theta Z + W + R. \quad (5)$$

To learn the linear model Θ , we would like to solve the following regularized least squares problem

$$\min_{\tilde{\Theta} \in \mathbb{R}^{n \times (n+p+1)}} \{ \|X - \tilde{\Theta}Z\|_F^2 + \lambda \|\tilde{\Theta}\|_F^2 \},$$

where $\lambda \geq 0$ is a regularization parameter. The closed-form solution of the above problem is given by

$$\hat{\Theta} = XZ'(ZZ' + \lambda I_{n+p+1})^{-1}, \quad (6)$$

under the invertibility assumption [27]. The estimation error is then given by

$$\begin{aligned} \|\hat{\Theta} - \Theta\| &= \| -\lambda\Theta(ZZ' + \lambda I_{n+p+1})^{-1} \\ &\quad + WZ'(ZZ' + \lambda I_{n+p+1})^{-1} \\ &\quad + RZ'(ZZ' + \lambda I_{n+p+1})^{-1} \|. \end{aligned} \quad (7)$$

For the ease of reference, the above steps are encapsulated in Algorithm 2.

Algorithm 2 System Identification Using Multiple Length 1 Trajectories

Input Dataset $\{(x_1^i, x_0^i, u_0^i) : 1 \leq i \leq N\}$, regularization parameter $\lambda \geq 0$

```
1: Construct the matrices  $X, Z$ . Compute  $\hat{\Theta} = XZ'(ZZ' + \lambda I_{n+p+1})^{-1}$ .
2: Extract the estimated system matrices  $A, B, o$  from the estimate  $\hat{\Theta} = [\hat{A} \ \hat{B} \ \hat{o}]$ .
```

In the next section, we will provide a finite sample bound of the system identification error (7) using Algorithm 1 and Algorithm 2. The bound explicitly characterizes how the error depends on N, q, σ_w, λ , and other system parameters, and will provide guidance on selecting q, λ .

IV. THEORETICAL ANALYSIS

To upper bound the system identification error in (7) with high probability, we bound the terms $\| -\lambda\Theta(ZZ' + \lambda I_{n+p+1})^{-1} \|, \| WZ'(ZZ' + \lambda I_{n+p+1})^{-1/2} \|, \| (ZZ' + \lambda I_{n+p+1})^{-1/2} \|, \| RZ'(ZZ' + \lambda I_{n+p+1})^{-1} \|$ separately. We provide some intermediate results first in Section IV-A. Our main result is presented in Section IV-B. Due to space constraints, some of the proofs are provided in [28].

A. Intermediate results

The following result shows the persistent excitation property of Algorithm 1. Note that the requirement on $N \geq 4(n+p)$ below is mainly used for numerical simplification.

Lemma 1: Suppose that Algorithm 1 is used to generate data. Let $N \geq 4(n+p)$. Then we have the following inequalities

$$\lambda_{\min}(ZZ') \geq N \min\left\{ \frac{q^2}{2(n+p)}, \frac{1}{2} \right\},$$

$$\lambda_{\max}(ZZ') \leq N \max\left\{ \frac{2q^2}{n+p}, 2 \right\}.$$

□

We will use the following lemma, which generalizes the upper bound for self-normalized martingales in [29, Theorem 1] to the multi-dimensional case. The proof can be found in [30, Lemma 5].

Lemma 2: Let $\{\mathcal{F}_t\}_{t \geq 0}$ be a filtration. Let $\{w_t\}_{t \geq 1}$ be a \mathbb{R}^n -valued stochastic process such that w_t is \mathcal{F}_t -measurable, and w_t is conditionally sub-Gaussian on \mathcal{F}_{t-1} with parameter R^2 . Let $\{z_t\}_{t \geq 1}$ be an \mathbb{R}^m -valued stochastic process such that z_t is \mathcal{F}_{t-1} -measurable. Assume that V is a $m \times m$ dimensional positive definite matrix. For all $t \geq 0$, define

$$\bar{V}_t = V + \sum_{s=1}^t z_s z_s', S_t = \sum_{s=1}^t z_s w_s'.$$

Then, for any $\delta \in (0, 1)$, and for all $t \geq 0$,

$$\begin{aligned} P(\|\bar{V}_t^{-\frac{1}{2}} S_t\| \leq \sqrt{\frac{32}{9} R^2 (\log \frac{9^n}{\delta} + \frac{1}{2} \log \det(\bar{V}_t V^{-1})})} \\ \geq 1 - \delta. \end{aligned}$$

□

We have the following result that upper bounds the contribution from noise.

Lemma 3: Suppose that Algorithm 1 is used to generate data. Let $N \geq 4(n+p)$ and $q \leq \sqrt{n+p}$. Then for any fixed $\delta \in (0, 1)$, we have with probability at least $1 - \delta$

$$\begin{aligned} & \|WZ'(ZZ' + \lambda I_{n+p+1})^{-1/2}\| \\ & \leq 3\sigma_w \sqrt{\log \frac{9^n}{\delta} + (n+p+1) \log(1 + \frac{4(n+p)}{q^2 + \zeta})}, \end{aligned}$$

where $\zeta = \frac{4\lambda(n+p)}{N}$. □

Next, we bound the contribution from the higher order terms.

Lemma 4: Suppose that Algorithm 1 is used to generate data. Let $N \geq 4(n+p)$ and $q \leq \sqrt{n+p}$. Fix constants c and β that satisfy Assumption 1, and denote $\gamma = \frac{\lambda(n+p)}{Nq^2}$. Then if $q < c$, we have

$$\begin{aligned} & \|RZ'(ZZ' + \lambda I_{n+p+1})^{-1}\| \\ & \leq \sqrt{\frac{2\beta^2(n^2 + np)}{1 + \gamma}} q + \frac{2(n+p)\sqrt{\lambda N n \beta^2 q^4}}{Nq^2 + 2\lambda(n+p)}. \end{aligned} \quad (8)$$

□

B. Main Result

Now we present our main result, a finite sample upper bound of the system identification error (7).

Theorem 1: Suppose that Algorithm 1 is used to generate data. Let $N \geq 4(n+p)$ and $q \leq \sqrt{n+p}$. Fix constants c, β that satisfy Assumption 1, and a confidence parameter $\delta \in (0, 1)$. Then if $q < c$, with probability at least $1 - \delta$, the estimation error of Algorithm 2 satisfies

$$\begin{aligned} \|\hat{\Theta} - \Theta\| & \leq \underbrace{\frac{5\sigma_w \sqrt{\log \frac{9^n}{\delta} + (n+p+1) \log(1 + \frac{4(n+p)}{q^2})}}{\sqrt{Nq^2/(n+p) + \lambda}}}_{\text{Error due to noise}} \\ & + \underbrace{\sqrt{\frac{2(n^2 + np)}{1 + \gamma}} \beta q}_{\text{Error due to nonlinearity}} \\ & + \underbrace{\frac{2(n+p)(\lambda\|\Theta\| + \sqrt{\lambda N n \beta^2 q^4})}{2\lambda(n+p) + Nq^2}}_{\text{Error due to regularization}}, \end{aligned} \quad (9)$$

where $\gamma = \frac{\lambda(n+p)}{Nq^2}$. □

Proof: Recall the estimation error in (7). We have

$$\begin{aligned} \|\hat{\Theta} - \Theta\| & \leq \lambda\|\Theta\| \|(ZZ' + \lambda I_{n+p+1})^{-1}\| \\ & + \|RZ'(ZZ' + \lambda I_{n+p+1})^{-1}\| \\ & + \|WZ'(ZZ' + \lambda I_{n+p+1})^{-1/2}\| \times \\ & \|(ZZ' + \lambda I_{n+p+1})^{-1/2}\|. \end{aligned} \quad (10)$$

Noting that

$$\begin{aligned} \|(ZZ' + \lambda I_{n+p+1})^{-1/2}\| & = \frac{1}{\sqrt{\lambda_{\min}(ZZ' + \lambda I_{n+p+1})}} \\ & \leq \frac{1}{\sqrt{\lambda_{\min}(ZZ') + \lambda}}, \end{aligned} \quad (11)$$

from Weyl's inequality [31], the result directly follows from applying Lemma 1, Lemma 3, and Lemma 4 after some algebraic manipulations. ■

Remark 1: Interpretation of Theorem 1. Note that Theorem 1 holds irrespective of the spectral radius of the system matrix A , which captures a well known advantage of the multiple trajectories setup. Below we discuss other key insights provided by Theorem 1.

Trade-off between error due to noise and error due to nonlinearity: Suppose that $\lambda = 0$ for now. When the system is perfectly linear, one has $\beta = 0$. Consequently, the upper bound in Theorem 1 only contains the error due to noise, which goes to zero with a rate of $\mathcal{O}(\frac{1}{\sqrt{N}})$. This implies a consistent estimator of which the convergence rate matches the results in the existing literature for learning perfectly linear system using random inputs [7], [16]. When there does exist nonlinearity, i.e., $\beta > 0$, one can observe that the error due to nonlinearity can be made arbitrarily small by choosing a smaller q used in Algorithm 1 (recall that q captures the magnitude of the initial conditions). On the other hand, a smaller q would also make the denominator of the term capturing error due to noise small, thus leading to a larger error due to noise. In other words, if one starts close enough to the reference point (by setting q to be small), one would have less bias due to nonlinearity, at the cost of having a smaller signal to noise ratio (thus a larger error due to noise). However, the error due to noise can always be made almost zero by increasing the number of experiments N . Consequently, if one can afford to generate a large amount of data, it is preferable to use a small q due to the low bias introduced by the nonlinear terms, and the small error introduced by the noise (which is due to the large amount of samples). These insights are different from system identification for truly linear systems, where it is commonly believed that a larger signal to noise ratio is always better. We will also illustrate these ideas empirically in Section V.

Role of regularization: Suppose that N, q are fixed. As λ becomes larger, we can observe that both the error due to noise and the error due to nonlinearity goes to zero, and the error due to regularization converges to $\|\Theta\|$. This result implies that setting λ to be relatively large can be helpful if σ_w is large (system is very noisy) or β is large (system has strong nonlinearity), while $\|\Theta\|$ is small. However, the optimal λ can be hard to obtain if (some upper bounds of) the parameters in (9) are unknown in advance. In practice, cross validation techniques [32] are commonly used to select a good value of λ . □

V. NUMERICAL EXAMPLES

In this section, we provide simulated numerical examples to validate the insights for system identification using AI-

gorithm 1 and Algorithm 2. We also compare the results against the single trajectory setup, where the input is set to be independent zero mean Gaussian, with slight adjustments to deal with the offset o in our setup (2), i.e., by appending ones in the regressor matrix. More specifically, we still use Algorithm 2 in the single trajectory setup, but the dataset is generated without restarting the system, see [16], [33] for examples. Such comparisons are made since Gaussian inputs are commonly used in the literature on linear system identification [7], [12]. For simplicity, we set $\lambda = 0$ for all experiments. All results are averaged over 10 independent experiments.

A. System with mild nonlinearity

In the first example, we investigate the performance of the system identification algorithms under mild nonlinearity. The model we use here captures the dynamics of a nonlinear pendulum.¹ The system states are the pendulum angle and its velocity, and the input is the torque applied. We set the mass and length of the pendulum to be 1 kg and 1 meter, respectively. After discretization using Euler's method by setting the sampling time to be 0.05 seconds, the dynamics is given by

$$\begin{bmatrix} x_{1,k+1} \\ x_{2,k+1} \end{bmatrix} = \begin{bmatrix} x_{1,k} + 0.05x_{2,k} \\ -0.49 \sin(x_{1,k}) + x_{2,k} + 0.05u_k \end{bmatrix} + w_k, \quad (12)$$

where we set w_k to be independent Gaussian random vectors with zero mean and covariance matrix given by $0.25I_2$. The linearized system matrices around the origin are given by

$$A = \begin{bmatrix} 1 & 0.05 \\ -0.49 & 1 \end{bmatrix}, B = \begin{bmatrix} 0 \\ 0.05 \end{bmatrix}, o = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (13)$$

We plot the system identification error using Algorithm 1 and Algorithm 2 versus the number of experiments N for $q = 1.2, 0.9, 0.6$ in Fig. 1. As can be observed, a smaller q could lead to a larger error when N is small, due to a smaller signal to noise ratio. However, a smaller q can eventually result in a smaller error when N is large enough due to less bias, which confirms our observations in Theorem 1.

In the single trajectory setup, we plot the error using i.i.d zero mean Gaussian inputs with different variance σ_u^2 , where N here represents the number of samples used in the single trajectory. A common heuristic is that one should apply small inputs to learn a good linear approximation around a given reference point, i.e., the variance σ_u^2 should be small. However, as shown in Fig. 2, the error plateaus at around 0.6, even for small variance inputs. The key reason is that the random input and process noise can always drive the system states to undesired regions and excite the higher order terms, unless the input is carefully designed. In fact, the paper [17] shows that random inputs in the single trajectory setup could result in inconsistent estimation under certain conditions even for Lipschitz nonlinearity.

¹<https://courses.engr.illinois.edu/ece486/fa2019/handbook/lec02.html>




Fig. 1: System identification error using Algorithms 1-2 with different q , mild nonlinearity




Fig. 2: System identification error using a single trajectory with different σ_u , mild nonlinearity

B. System with strong nonlinearity

In the second example, we investigate the performance of the system identification algorithms under strong nonlinearity (where the assumption of Lipschitzness used in [17] no longer holds). The virtual model we use here is given by

$$\begin{bmatrix} x_{1,k+1} \\ x_{2,k+1} \end{bmatrix} = \begin{bmatrix} 0.9 & 0.5 \\ 0 & 0.8 \end{bmatrix} \begin{bmatrix} x_{1,k} \\ x_{2,k} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} u_k + \begin{bmatrix} x_{1,k}^3 + x_{2,k}^5 \\ x_{1,k}x_{2,k} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} + w_k, \quad (14)$$

where we again set w_k to be independent Gaussian random vectors with zero mean and covariance matrix given by $0.25I_2$.

Again, we plot the system identification error using Algorithm 1 and Algorithm 2 versus the number of experiments N for $q = 0.6, 0.4, 0.2$. As can be observed, similar trends still hold, i.e., a smaller q results in a larger error when N is small, but is beneficial in the long run, even for system with relatively strong nonlinearity.

In contrast, in the single trajectory setup, we applied i.i.d zero mean Gaussian inputs with variance $\sigma_u^2 =$

$0.1^2, 0.01^2, 0.001^2$. However, all of them fail to converge and result in numerical issues since the noise and non-zero offset drive the system states to regions where nonlinearity dominates.

VI. CONCLUSION AND FUTURE WORK

In this paper, we proposed system identification algorithms to learn the linearized model of a system. Unlike existing works, we assume that the underlying dynamics could be nonlinear. We presented a finite sample error bound of the algorithms, which shows that one can learn the linearized dynamics with arbitrarily small error given sufficiently many samples, and demonstrates a trade-off between the error due to noise and the error due to nonlinearity. Our bound further characterizes the benefits of using regularization. As shown in [34], initializing states at different locations might come at different costs. Consequently, future work would focus on studying how to optimize the data collection procedure under constraints on initial state/input.




Fig. 3: System identification error using Algorithms 1-2 with different q , strong nonlinearity

REFERENCES

- [1] S. Athey, "The impact of machine learning on economics," in *The economics of artificial intelligence: An agenda*. University of Chicago Press, 2018, pp. 507–547.
- [2] T. M. Mitchell et al., *Machine learning*. McGraw-hill New York, 2007, vol. 1.
- [3] L. Ljung, "System identification," *Wiley encyclopedia of electrical and electronics engineering*, pp. 1–19, 1999.
- [4] W. J. Rugh, *Linear system theory*. Prentice-Hall, Inc., 1996.
- [5] D. Bauer, M. Deistler, and W. Scherrer, "Consistency and asymptotic normality of some subspace algorithms for systems without observed inputs," *Automatica*, vol. 35, no. 7, pp. 1243–1254, 1999.
- [6] M. Jansson and B. Wahlberg, "On consistency of subspace methods for system identification," *Automatica*, vol. 34, no. 12, pp. 1507–1519, 1998.
- [7] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, pp. 1–47, 2019.
- [8] S. Fattahi and S. Sojoudi, "Data-driven sparse system identification," in *Proc. Allerton Conference on Communication, Control, and Computing*, 2018, pp. 462–469.
- [9] Y. Zheng and N. Li, "Non-asymptotic identification of linear dynamical systems using multiple trajectories," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1693–1698, 2020.

- [10] L. Xin, G. Chiu, and S. Sundaram, "Learning the dynamics of autonomous linear systems from multiple trajectories," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 3955–3960.
- [11] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *Proc. Conference On Learning Theory*, 2018, pp. 439–473.
- [12] S. Oymak and N. Ozay, "Non-asymptotic identification of LTI systems from a single trajectory," in *American control conference*. IEEE, 2019, pp. 5655–5661.
- [13] M. Simchowitz, R. Boczar, and B. Recht, "Learning linear dynamical systems with semi-parametric least squares," in *Proc. Conference on Learning Theory*, 2019, pp. 2714–2802.
- [14] T. Sarkar, A. Rakhlin, and M. A. Dahleh, "Nonparametric finite time LTI system identification," *arXiv preprint arXiv:1902.01848*, 2019.
- [15] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, "Finite time identification in unstable linear systems," *Automatica*, vol. 96, pp. 342–353, 2018.
- [16] T. Sarkar and A. Rakhlin, "Near optimal finite time identification of arbitrary linear dynamical systems," in *Proc. International Conference on Machine Learning*, 2019, pp. 5610–5618.
- [17] A. Sarker, P. Fisher, J. E. Gaudio, and A. M. Annaswamy, "Accurate parameter estimation for safety-critical systems with unmodeled dynamics," *Artificial Intelligence*, p. 103857, 2023.
- [18] Y. Sattar and S. Oymak, "Non-asymptotic and accurate learning of nonlinear dynamical systems," *Journal of Machine Learning Research*, vol. 23, no. 140, pp. 1–49, 2022.
- [19] H. Mania, M. I. Jordan, and B. Recht, "Active learning for nonlinear system identification with guarantees," *arXiv preprint arXiv:2006.10277*, 2020.
- [20] D. Foster, T. Sarkar, and A. Rakhlin, "Learning nonlinear dynamical systems from a single trajectory," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 851–861.
- [21] A. Mauroy and J. Goncalves, "Linear identification of nonlinear systems: A lifting technique based on the koopman operator," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 6500–6505.
- [22] W. Hao, B. Huang, W. Pan, D. Wu, and S. Mou, "Deep koopman representation of nonlinear time varying systems," *arXiv preprint arXiv:2210.06272*, 2022.
- [23] Ş. Cobzaş, R. Miculescu, A. Nicolae et al., *Lipschitz functions*. Springer, 2019.
- [24] O. Rivasplata, "Subgaussian random variables: An expository note," *Internet publication, PDF*, vol. 5, 2012.
- [25] R. Courant, F. John, A. A. Blank, and A. Solomon, *Introduction to calculus and analysis*. Springer, 1965, vol. 1.
- [26] G. B. Folland, "Higher-order derivatives and Taylor's formula in several variables," *Preprint*, pp. 1–4, 2005.
- [27] A. E. Hoerl and R. W. Kennard, "Ridge regression: Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.
- [28] L. Xin, G. Chiu, and S. Sundaram, "Learning linearized models from nonlinear systems with finite data," *arXiv preprint*, 2023.
- [29] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, "Improved algorithms for linear stochastic bandits," *Advances in neural information processing systems*, vol. 24, 2011.
- [30] L. Xin, L. Ye, G. Chiu, and S. Sundaram, "Learning dynamical systems by leveraging data from similar systems," *arXiv preprint arXiv:2302.04344*, 2023.
- [31] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.
- [32] P. Refaeilzadeh, L. Tang, and H. Liu, "Cross-validation," *Encyclopedia of database systems*, vol. 5, pp. 532–538, 2009.
- [33] L. Ye, H. Zhu, and V. Gupta, "On the sample complexity of decentralized linear quadratic regulator with partially nested information structure," *IEEE Transactions on Automatic Control*, 2022.
- [34] A. A. Ahmadi, A. Chaudhry, V. Sindhvani, and S. Tu, "Safely learning dynamical systems from short trajectories," in *Learning for Dynamics and Control*. PMLR, 2021, pp. 498–509.