
High Effort, Low Gain: Fundamental Limits of Active Learning for Linear Dynamical Systems

Nicolas Chatzikiriakos

Institute for Systems Theory and Automatic Control
University of Stuttgart
nicolas.chatzikiriakos@ist.uni-stuttgart.de

Kevin Jamieson

Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, WA
jamieson@cs.washington.edu

Andrea Iannelli

Institute for Systems Theory and Automatic Control
University of Stuttgart
andrea.iannelli@ist.uni-stuttgart.de

Abstract

In this work, we consider the problem of identifying an unknown linear dynamical system given a finite hypothesis class. In particular, we analyze the effect of the excitation input on the sample complexity of identifying the true system with high probability. To this end, we present sample complexity lower bounds that capture the choice of the selected excitation input. The sample complexity lower bound gives rise to a system theoretic condition to determine the potential benefit of experiment design. Informed by the analysis of the sample complexity lower bound, we propose a persistent excitation (PE) condition tailored to the considered setting, which we then use to establish sample complexity upper bounds. Notably, the PE condition is weaker than in the case of an infinite hypothesis class and allows analyzing different excitation inputs modularly. Crucially the lower and upper bounds share the same dependency on key problem parameters. Finally, we leverage these insights to propose an active learning algorithm that sequentially excites the system optimally with respect to the current estimate, and provide sample complexity guarantees for the presented algorithm. Concluding simulations showcase the effectiveness of the proposed algorithm.

1 Introduction

The problem of learning a model of an unknown dynamical system from data is important across domains such as reinforcement learning (RL), data-driven control and robotics. Hereby, it is of particular interest to obtain an accurate model with high confidence from as few samples as possible, since data collection is often expensive. Specifically, in online settings such as RL exploration comes at the cost of loss in performance. Hence, it is of particular interest to understand the mechanisms and benefits of (active) exploration since this enables finding suitable tradeoffs between exploration and exploitation. While there exists a large literature on experiment design both in dynamical systems and learning theory, there are still fundamental challenges in the field that have not been addressed even in simple settings. Specifically a study of the potential benefit of experiment design algorithms

is lacking in the literature of dynamical systems, raising the question of whether experiment design algorithms should be applied universally, or whether a more nuanced answer is appropriate. Motivated by these challenges, we consider the setting where the learner has prior knowledge of the true system through a finite hypothesis class. This reflects cases where some prior knowledge is available, e.g., based on first principles, yet certain parameters are hard to model or vary across different instances. When it comes to dynamical systems, the data is usually collected from a single trajectory and hence is highly correlated. This poses a key challenge when analyzing learning in the finite sample regime. In particular, existing works considering the identification of dynamical systems¹ mostly consider the case of an infinite hypothesis class and rely on the least squares estimator. Early works derived sample complexity upper bounds for Gaussian inputs for linear [21, 24] and certain classes of non-linear systems [7, 23, 22]. Further, sample complexity lower bounds for Gaussian excitations and linear systems have been considered in [12, 25]. For the case of a finite hypothesis class, [3, 18] provide sample complexity upper and lower bounds when the excitation is Gaussian. The problem of experiment design has a long history in system identification (see [2, 10]). While these classical works consider the asymptotic case, recent works provided a finite sample perspective on the topic for linear [28, 4] and certain classes of non-linear systems [16, 17]. While the underlying principles in the finite-sample analysis are always similar, there exists no modular framework for this analysis. Further, the benefit of experiment design over random excitations has not been explicitly considered in dynamical systems. This work addresses these gaps by answering the following questions to improve our understanding of experiment design and active learning:

- Q1) How large is the problem-specific benefit of oracle experiment design over isotropic Gaussian excitations in terms of the sample complexity?
- Q2) How can sample complexity upper bounds be established for general excitation inputs?
- Q3) How can active learning algorithms generate excitation inputs to provably outperform random isotropic excitations?
- Q4) How does the finite hypothesis class simplify experiment design compared to an infinite hypothesis class?

Even though we consider these questions in a linear setup, the improved understanding of key mechanisms in exploration and experiment design can also be beneficial in more complex settings such as RL. Our primary contributions can be summarized as follows:

1. We provide fundamental instance-dependent lower bounds on the sample complexity of identifying the true system with high probability. To this end, we leverage tools from information theory and derive the optimal oracle excitation input during data collection. We further provide a detailed analysis of the system-theoretic quantities dictating how the sample complexity lower bounds differ for different excitation inputs.
2. Building on the notion of persistency of excitation (PE), we propose a modular framework to derive sample complexity upper bounds for different excitation inputs. While the probability of miss-specification decays exponentially for all excitations satisfying PE, we establish that the decay rate depends on system theoretic quantities that match the lower bounds.
3. Using the notion of PE we establish an interpretable condition pointing out when experiment design using certainty equivalence is more efficient than random excitations. Interestingly, the derived condition can be satisfied even before data-collection, indicating that the additional knowledge of the finite hypothesis class can be leveraged for experiment design.
4. Notably, our analysis also uncovers cases where the benefit of experiment design is small even in an oracle setup, indicating that the usefulness of experiment design needs to be evaluated on a problem-specific level.

Notation: The n -dimensional simplex is denoted by Δ_n . We denote the set of symmetric positive (semi-)definite matrices of dimension $n \times n$ by \mathbb{S}_{++}^n (\mathbb{S}_+^n). Given a vector $v \in \mathbb{R}^n$ and a matrix $M \in \mathbb{S}_{++}^n$ we define $\|v\|_M := \sqrt{v^\top M v}$. Given a symmetric matrix $M \in \mathbb{R}^{n \times n}$ we denote its largest eigenvalue by $\lambda_{\max}(M)$ and the mean of all eigenvalues by $\lambda_{\text{mean}}(M) = \frac{1}{n} \text{tr}(M)$, where $\text{tr}(M)$ indicates the trace of the matrix M . Given a sequence $\{u(t)\}_{t=t_0}^{t_1}$ we denote the stacked collection as $U = [u(t_0)^\top \dots u(t_1)^\top]^\top$, where the boundaries of the interval will be clear from the context. We use $\text{diag}_t(M)$ to denote a block-diagonal matrix which repeats the matrix M on its diagonal t -times, i.e., $\text{diag}_t(M) := I_t \otimes M$, where \otimes denotes the Kronecker product.

¹We provide an extended overview of the related literature in Appendix A.

2 Preliminaries

We consider the unknown discrete-time linear time-invariant (LTI) dynamical system

$$x(t+1) = A_*x(t) + B_*u(t) + w(t), \quad w(t) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_w), \quad (1)$$

where $x(t) \in \mathbb{R}^{n_x}$, $u(t) \in \mathbb{R}^{n_u}$ are the state and the input of the system at timestep t , and $w(t)$ is process noise. While we assume knowledge of the covariance matrix $\Sigma_w \in \mathbb{S}_{++}^{n_x}$, all the presented results and the proposed algorithm can be adapted to the case where Σ_w is unknown and only an upper bound $\sigma_w^2 I \succeq \Sigma_w$ or an informative prior is available. For notational simplicity, we assume $x(0) = 0$ unless stated otherwise. Further, we assume the learner has knowledge of a finite set of systems, which contains the true system, i.e.,

$$\theta_* = (A_*, B_*) \in \mathcal{S} := \{(A_0, B_0), \dots, (A_N, B_N)\}, \quad (2)$$

where \mathcal{S} is known to the learner, and we assume $(A_*, B_*) = (A_0, B_0)$ for notational simplicity. The goal of the learner is to identify the true system matrices $\theta_* \in \mathcal{S}$ from a single data trajectory $\mathcal{D}_T = (\{x(t)\}_{t=0}^T, \{u(t)\}_{t=0}^{T-1})$, where the excitation input u can be selected by the learner. Thus, in the context of this work an identification problem is uniquely defined by the tuple (θ_*, \mathcal{S}) . For our analysis it is important to understand how predictions using θ_* compare to predictions using θ_i , for some $i \in [1, N]$. To this end, we denote the sum of one-step prediction errors between θ_* and θ_i by

$$\tilde{\varepsilon}_{\theta_i}(0, \tau) := \sum_{t=0}^{\tau-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2, \quad (3)$$

where we defined $\Delta A_i := A_* - A_i$ and $\Delta B_i := B_* - B_i$. Recall that the data \mathcal{D}_τ is collected along a trajectory of the dynamical system (1). Recursively plugging in the dynamics (1) for $x(t)$ yields

$$\tilde{\varepsilon}_{\theta_i}(0, \tau) = \sum_{t=0}^{\tau-1} \left\| \Delta A_i \left(\sum_{s=0}^{t-1} A_*^{t-1-s} B_* u(s) + A_*^{t-1-s} w(s) \right) + \Delta B_i u(t) \right\|_{\Sigma_w^{-1}}^2, \quad (4)$$

which reveals that $\tilde{\varepsilon}_{\theta_i}(0, \tau)$ is a random quantity even if the excitation input is deterministic². This is due to the process noise affecting the system (1) and hence the data. Taking the expectation and using $w(t) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_w)$ it can be shown that

$$\mathbb{E}[\tilde{\varepsilon}_{\theta_i}(0, \tau)] = \mathbb{E} \left[\sum_{t=0}^{\tau-1} \left\| \Delta A_i \sum_{s=0}^{t-1} A_*^{t-1-s} B_* u(s) + \Delta B_i u(t) \right\|_{\Sigma_w^{-1}}^2 \right] + \mathbb{E} \left[\left\| \Delta A_i \sum_{s=0}^{t-1} A_*^{t-1-s} w(s) \right\|_{\Sigma_w^{-1}}^2 \right]. \quad (5)$$

To simplify notation we define the Toeplitz matrices

$$S_u(\tau) := \begin{bmatrix} B_* & 0 & \dots & 0 \\ A_* B_* & B_* & 0 & \vdots \\ \vdots & & \ddots & 0 \\ A_*^{\tau-1} B_* & A_*^{\tau-2} B_* & \dots & B_* \end{bmatrix}, \quad S_w(\tau) := \begin{bmatrix} \Sigma_w^{1/2} & 0 & \dots & 0 \\ A_* \Sigma_w^{1/2} & \Sigma_w^{1/2} & 0 & \vdots \\ \vdots & & \ddots & 0 \\ A_*^{\tau-1} \Sigma_w^{1/2} & A_*^{\tau-2} \Sigma_w^{1/2} & \dots & \Sigma_w^{1/2} \end{bmatrix}, \quad (6)$$

that can be used to map from the t -step noise and input trajectories to corresponding the state trajectory. Using (6), computing the expectation in (5) and writing it in matrix form simplifies to

$$\begin{aligned} \mathbb{E}[\tilde{\varepsilon}_{\theta_i}(0, \tau)] &= \mathbb{E} \left[U^\top (R_{\Sigma_w^{-1}}^i(\tau) + S_u(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_u(\tau) \right. \\ &\quad \left. + N_{\Sigma_w^{-1}}^i(\tau) S_u(\tau) + (N_{\Sigma_w^{-1}}^i(\tau) S_u(\tau))^\top) U \right] + \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right), \end{aligned} \quad (7)$$

where $U \in \mathbb{R}^{n_u \tau}$ and we introduced the block-diagonal matrices

$$Q_{\Sigma_w^{-1}}^i(\tau) := \text{diag}_\tau(\Delta A_i^\top \Sigma_w^{-1} \Delta A_i), \quad R_{\Sigma_w^{-1}}^i(\tau) := \text{diag}_\tau(\Delta B_i^\top \Sigma_w^{-1} \Delta B_i) \quad (8a)$$

$$N_{\Sigma_w^{-1}}^i(\tau) := \text{diag}_\tau(\Delta B_i^\top \Sigma_w^{-1} \Delta A_i) \quad (8b)$$

²A detailed derivation, for the general case $x(0) \in \mathbb{R}^{n_x}$ that includes the case where $u(t)$ consists of both a deterministic and a random part is presented in Appendix G.

Note that $U^\top R_{\Sigma_w^{-1}}^i(\tau)U = \sum_{t=0}^\tau \|\Delta B_i u(t)\|_{\Sigma_w^{-1}}^2$ measures to what extent the difference between B_* and B_i can be seen in the weighted prediction error when applying the input sequence U . Similarly, $U^\top S_u(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau)S_u(\tau)U$ measures to what extent the difference between A_* and A_i can be seen in the weighted prediction error when applying U and $\text{tr}(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau)S_w(\tau))$ measures the influence of the process noise on this quantity. Thus, accounting for the cross term, the matrices

$$W_i(\tau) := R_{\Sigma_w^{-1}}^i(\tau) + S_u(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau)S_u(\tau) + N_{\Sigma_w^{-1}}^i(\tau)S_u(\tau) + (N_{\Sigma_w^{-1}}^i(\tau)S_u(\tau))^\top, \quad \forall i \in [1, N], \quad (9)$$

measure the difficulty of distinguishing between θ_* and θ_i and the sensitivity of (3) to changes in the excitation input.

3 Can active learning algorithms improve identification?

To analyze the influence of the data collection scheme on the sample complexity, we first derive a sample complexity lower bound that holds for any *reasonable* algorithm. This instance-specific lower bound quantifies the hardness of learning in our setup and provides a nuanced answer on whether active learning algorithms can significantly reduce the sample complexity. Further, in Section 3.2 we provide a modular framework for deriving complementary sample complexity upper bounds. In this work, we assume the excitation input satisfies the following standard assumption (see, e.g., [27])

Assumption 1 (Bounded input power). *The expected average power of the (potentially random) input sequence $\{u(t)\}_{t=0}^{T-1}$ is bounded by γ_u^2 , i.e., $\mathbb{E} \left[\sum_{t=0}^{T-1} \|u(t)\|^2 \right] \leq \gamma_u^2 T$.*

3.1 Sample complexity lower bounds

To formalize the algorithms that we consider in this section, we define the class of δ -correct algorithms.

Definition 1 (δ -correct algorithms). *Consider the setup described in Section 2. An algorithm is called δ -correct, if for all $\delta \in (0, 1)$, and any θ_* and \mathcal{S} there exists a finite time \bar{T} such that for all $t \geq \bar{T}$ the algorithm returns an estimate $\hat{\theta}_t$ that satisfies $\mathbb{P}[\hat{\theta}_t = \theta_*] \geq 1 - \delta$.*

Note that by restricting to δ -correct algorithms, we can obtain instance-specific sample complexity lower bounds that hold for any reasonable algorithm and any input sequence.

Theorem 3.1. *Consider the unknown dynamical system (1) with $x(0) = 0$ and the set \mathcal{S} defined in (2). Then for any (potentially random) excitation input sequence³ $U \in \mathbb{R}^{n_u \bar{T}}$ and any δ -correct algorithm it holds that*

$$\min_{i \in [1, N]} \mathbb{E} [U^\top W_i(\bar{T})U] + \text{tr}(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T})S_w(\bar{T})) \geq 2 \log \left(\frac{1}{2.4\delta} \right). \quad (10)$$

Furthermore, under Assumption 1 the lower bound is minimized by the excitation input

$$U^* \in \arg \max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1, N]} U^\top W_i(\bar{T})U + \text{tr}(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T})S_w(\bar{T})), \quad (11)$$

and when applying U^* any δ -correct algorithm satisfies

$$\min_{p \in \Delta_N} \gamma_u^2 \bar{T} \lambda_{\max} \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right) + \max_{i \in [1, N]} \text{tr}(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T})S_w(\bar{T})) \geq 2 \log \left(\frac{1}{2.4\delta} \right). \quad (12)$$

While using the maximum of the trace term in (12) introduces some conservatism, in practical applications the noise is often significantly smaller than the input. In these cases (12) is dominated by the first term and the added conservatism is small. The proof of Theorem 3.1 is presented in Appendix B.1. Note that while the solution to (11) might not be unique, any optimizer suffices to achieve the optimal lower bound. As shown in Appendix C, the optimal solution exists.

³Random and deterministic input sequences are denoted identically. The interpretation will be clear from the context.

Remark 3.1. Recall from (7)-(9) that, given $x(0) = 0$, (11) is equivalent to

$$U^* \in \arg \max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1, N]} \mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right]. \quad (13)$$

Thus, Theorem 3.1 implies that the optimal excitation U^* maximizes $\mathbb{E} [\tilde{\varepsilon}_{\theta_i}(0, \bar{T})]$ uniformly over all $\theta_i \in \mathcal{S} \setminus \{\theta_*\}$. Thus, U^* maximizes the distance between $\mathbb{E} [\tilde{\varepsilon}_{\theta^*}(0, \bar{T})] = 0$ and $\mathbb{E} [\tilde{\varepsilon}_{\theta_i}(0, \bar{T})]$, to separate the true system θ_* from all other systems $\theta_i \in \mathcal{S} \setminus \{\theta_*\}$ as clearly as possible.

Theorem 3.1 can be used to obtain a sample complexity lower bound for isotropic Gaussian inputs.

Corollary 3.2. Consider the unknown dynamical system (1) with $x(0) = 0$ and the set \mathcal{S} defined in (2). Suppose that $u(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$. Then for any δ -correct algorithm it holds that

$$\min_{i \in [1, N]} \gamma_u^2 \bar{T} \lambda_{\text{mean}}(W_i(\bar{T})) + \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_w(\bar{T}) \right) \geq 2 \log \left(\frac{1}{2.4\delta} \right). \quad (14)$$

The proof is provided in Appendix B.2. To gain an intuition for the results in this chapter consider the following example, where \mathcal{S} consists of only two systems which eliminates the need for minimization.

Example 3.1. First, consider the identification problem defined by

$$A_* = \begin{bmatrix} 0 & 0.1 \\ 0 & 0 \end{bmatrix} \quad B_* = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad A_1 = \begin{bmatrix} 0 & 0.2 \\ 0 & 0 \end{bmatrix} \quad B_1 = B_* \quad (15)$$

and $\mathcal{S} = \{(A_*, B_*), (A_1, B_*)\}$. A straightforward calculation yields $W_1(t) = \text{diag}_t([0.01])$ and $\lambda_j(W_1(t)) = 0.01, \forall j \in [1, t], \forall t \geq 1$. Thus, in this case isotropic Gaussian excitations are optimal in the sense of the sample complexity lower bound. Next, consider $\tilde{\mathcal{S}} = \{(\tilde{A}_*, \tilde{B}_*), (\tilde{A}_1, \tilde{B}_*)\}$ with

$$\tilde{A}_* = \begin{bmatrix} A_* & 0_{2 \times d} \\ 0_{d \times 2} & I_d \end{bmatrix} \quad \tilde{B}_* = \begin{bmatrix} B_* & 0_{2 \times d} \\ 0_{d \times 1} & I_d \end{bmatrix} \quad \tilde{A}_1 = \begin{bmatrix} A_1 & 0_{2 \times d} \\ 0_{d \times 2} & I_d \end{bmatrix} \quad \tilde{B}_1 = \tilde{B}_*, \quad (16)$$

for some fixed $d > 0$. Through direct calculations, we obtain $\tilde{W}_1(t) = \text{diag}_t \left(\begin{bmatrix} 0.01 & 0_{1 \times d} \\ 0_{d \times 1} & 0_{d \times d} \end{bmatrix} \right)$,

$\forall t \geq 1$ and hence $\lambda_{\text{mean}}(\tilde{W}_1(t)) = \frac{1}{d+1} \lambda_{\text{max}}(\tilde{W}_1(t))$. Hence, when d is large, the sample complexity lower bound for isotropic Gaussian excitations is significantly larger than for the optimal oracle excitation. This indicates that using isotropic Gaussian excitations requires more samples to achieve the same confidence δ if an algorithm that matches the lower bound is used. This can be observed in the sample complexity upper bounds we present in the next section. Further, under the same process noise conditions the sample complexity lower bound of the identification problems (θ_*, \mathcal{S}) and $(\tilde{\theta}_*, \tilde{\mathcal{S}})$ is identical given the respective optimal oracle excitation is used for both instances, as $\lambda_{\text{max}}(W_1(t)) \equiv \lambda_{\text{max}}(\tilde{W}_1(t))$. Hence, the difficulty of problems (θ_*, \mathcal{S}) and $(\tilde{\theta}_*, \tilde{\mathcal{S}})$ is identical. This is due to the fact that the last d modes of the system (16) are not relevant for the identification problem since they are decoupled from the unknown part. Confirming this intuition, \tilde{U}^* only excites the system with the first input. Crucially, our mathematical formulation enables the analysis of complex problem setups, where the optimal excitation and hardness cannot be determined intuitively.

3.2 Sample complexity upper bounds

In this section, we establish a modular framework to derive sample complexity upper bounds for identifying θ_* with high probability. Our proposed framework builds on the notion of persistency of excitation which is closely connected to the observations made in the previous section.

3.2.1 Persistency of Excitation

Recall that in Example 3.1, the goal of identifying the true system can be accomplished by exciting only the first mode of the system using only the first control input. Clearly, this input does not satisfy

$$\mathbb{E} \left[\sum_{t=0}^{\tau-1} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \begin{bmatrix} x(t)^\top & u(t)^\top \end{bmatrix} \right] \succeq c I_{n_x + n_u}, \quad c > 0 \quad (17)$$

which is the PE condition required in the case of an infinite hypothesis class [26]. In fact, it is only necessary to excite the parts of the system which carry uncertainty. Thus, we introduce weaker PE condition, which has been used as an assumption in a similar form in [18].

Definition 2 (Persistency of excitation (PE)). *Consider an input sequence $\{u(t)\}_{t=0}^{\tau-1}$, satisfying $\mathbb{E}[\sum_{t=0}^{\tau-1} \|u(t)\|^2] = \gamma_u^2 \tau$ for some $\gamma_u > 0$. We say $\{u(t)\}_{t=0}^{\tau-1}$ is persistently exciting for (θ_*, \mathcal{S}) if there exist constants $c_u(\tau) > 0$, $c_w(\tau) > 0$ such that for any $x(0) \in \mathbb{R}^{n_x}$ and $\theta_i \in \mathcal{S} \setminus \{\theta_*\}$*

$$\frac{1}{\tau} \sum_{t=0}^{\tau-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq c_u(\tau) \gamma_u^2 + c_w(\tau). \quad (18)$$

Note that, the constants c_u and c_w depend on the block length τ . Further they are instance specific, i.e., they depend on the set \mathcal{S} . By the sample complexity lower bound the excitation input needs to satisfy Definition 2 for some $\tau > 0$ to guarantee identification of θ_* with high probability (see Remark 3.1 and Theorem B.1). If not, there exists some $\theta_i \neq \theta_*$ which yields the same state trajectory and hence is indistinguishable from θ_* given the data. Given the PE condition (18), it follows that isotropic Gaussian excitations and the optimal oracle excitation U^* defined in (11) satisfy Definition 2.

Lemma 3.3 (PE for isotropic Gaussian inputs). *Consider the system (1) and the set \mathcal{S} defined in (2). Then $u(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma^2}{n_u} I_{n_u})$ is PE for any block length $\tau > 0$ with*

$$c_u^{\text{rand}}(\tau) = \min_{p \in \Delta_N} \lambda_{\text{mean}} \left(\sum_{i=1}^N p_i W_i(\tau) \right), \quad c_w(\tau) = \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_w(\tau) \right). \quad (19)$$

Lemma 3.4 (PE for optimal oracle excitation). *Consider the system (1) and let the set \mathcal{S} defined in (2). Then*

$$U^*(x(0)) \in \arg \max_{U^\top U \leq \gamma_u^2 \tau} \min_{i \in [1, N]} \sum_{t=0}^{\tau} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \quad (20)$$

is PE for any block length τ with

$$c_u^{\text{opt}}(\tau) = \min_{p \in \Delta_N} \lambda_{\text{max}} \left(\sum_{i=1}^N p_i W_i(\tau) \right), \quad c_w(\tau) = \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_w(\tau) \right). \quad (21)$$

The proof of both results is given in Appendices D.1 and D.2, respectively. Importantly, the PE coefficients $c_u^{\text{rand}}(\tau)$ and $c_w(\tau)$ coincide with the key problem parameters influencing the sample complexity lower bounds presented in Theorem 3.1 and Corollary 3.2. Further, the optimal excitation U^* derived in Theorem 3.1 maximizes the PE coefficients $c_u(\tau)$, $c_w(\tau)$ in (18).

Remark 3.2. *Lemmas 3.3 and 3.4 also shed light on a tradeoff when selecting the block size τ . The properties of the dynamical system (1) dictate how fast $c_u(\tau)$ grows, so that a smaller (larger) block length τ might provide better guarantees.*

Remark 3.3. *If Σ_w is unknown and instead an estimate $\hat{\Sigma}_w$ is used to solve (20), the corresponding input sequence is PE, although with potentially suboptimal coefficients. The degree of sub-optimality is instance-dependent and can be analyzed by computing and comparing the coefficients.*

3.2.2 High probability identification with finite samples

Having established PE, we now show that PE guarantees fast identification of the true system θ_* with high probability. To do so, we introduce a sequential estimation algorithm for general input sequences. In particular, we derive a sample complexity upper bound for Algorithm 1 that holds for any PE input sequence. For estimation, the algorithm evaluates the sum of weighted one-step prediction errors of θ_i up to time t defined as

$$\varepsilon_{\theta_i}(t) := \sum_{s=0}^{t-1} \|x(s+1) - A_i x(s) - B_i u(s)\|_{\Sigma_w^{-1}}^2. \quad (22)$$

Clearly, $\varepsilon_{\theta_i}(t)$ can be interpreted as the negative log-likelihood of system i given the data collected from time 0 to t . Thus, the termination criterion of Algorithm 1 is equivalent to a log-likelihood test.

Algorithm 1 Sequential identification algorithm

Require: \mathcal{S} , episode length τ , desired confidence δ

- 1: **for** $k = 1, 2, \dots$ **do**
- 2: Collect data using PE excitation input $\{u(t)\}_{t=(k-1)\tau}^{k\tau-1}$ with coefficients $c_{u_k}(\tau), c_w(\tau)$
- 3: Compute $\varepsilon_{\theta_i}(k\tau)$ for all $i \in [0, N]$
- 4: **if** $\exists \hat{\theta} \in \mathcal{S} : \varepsilon_{\theta_i}(k\tau) - \varepsilon_{\hat{\theta}}(k\tau) > 2 \log\left(\frac{N}{\delta}\right)$ for all $\theta_i \in \mathcal{S} \setminus \hat{\theta}$ **then**
- 5: Stop and **return** estimate $\hat{\theta}$

While related works considering this setup [3, 18] only analyze the case where the data is generated by Gaussian excitations, the results in this work hold for *any* input satisfying PE. In particular, this allows us to reason about how learning can be accelerated through particular choices of the excitation input. To formalize this, we propose the following result.

Theorem 3.5. *Consider the unknown system (1), set \mathcal{S} as defined in (2). Then Algorithm 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates no later than when k satisfies*

$$\tau \sum_{j=1}^k c_{u_j}(\tau) \gamma_u^2 + c_w(\tau) \geq c' \log\left(\frac{N}{\delta}\right), \quad (23)$$

where c' is a constant influenced by the variance in $\tilde{\varepsilon}_{\theta_i}$.

The proof of Theorem 3.5 is presented in Appendix E.1, where we also present the full version of the result. If the PE coefficients are similar throughout the episodes Theorem 3.5 guarantees that the risk of miss-specification decays exponentially in k . Further, Theorem 3.5 establishes a modular framework to derive sample complexity upper bounds for different excitations. That is, to provide a sample complexity upper bound for Algorithm 1 it is only necessary to show the excitation satisfies Definition 2. Given the PE coefficients $c_u(\tau), c_w(\tau)$ the sample complexity upper bound then follows immediately. Hence, for the input sequences analyzed previously, the following result holds.

Corollary 3.6. *Consider the same setup as in Theorem 3.5. If $u(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$ then Algorithm 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates at the latest when $T = k\tau$ first satisfies*

$$T \left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\text{mean}} \left(\sum_{i=1}^N p_i W_i(\tau) \right) + \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{I_{n_x}}^i S_w(\tau) \right) \right) \geq c' \log\left(\frac{N}{\delta}\right). \quad (24)$$

Further, if the optimal oracle excitation input U^* defined in (20) is applied the estimate $\hat{\theta}$ satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$, and Algorithm 1 terminates at the latest when $T = k\tau$ first satisfies

$$T \left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\text{max}} \left(\sum_{i=1}^N p_i W_i(\tau) \right) + \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{I_{n_x}}^i S_w(\tau) \right) \right) \geq c'' \log\left(\frac{N}{\delta}\right). \quad (25)$$

where c' and c'' are constants influenced by the variance in $\tilde{\varepsilon}_{\theta_i}$.

After leveraging PE, Corollary 3.6 follows directly from Theorem 3.5. We provide the proof and the full version in Appendix E.3. Clearly, the level of PE of the excitation dictates the speed of identification of the true system. Since the PE coefficients coincide with the key parameters entering the sample complexity lower bounds we observe the same dependency in the sample complexity upper bounds, as well. The block parameter τ can be used as a degree of freedom to optimize the bounds. As discussed in Remark 3.2 the optimal choice of τ depends on the problem setup.

4 Sequential input design algorithm

In the previous section, we established that the matrices $W_i(\tau)$ are the key quantities that determine the degree of sub-optimality of random excitations compared to the optimal oracle excitation. However, since the optimal oracle excitation input depends on θ_* it cannot be computed in practice. This

Algorithm 2 Input design subroutine

Require: \mathcal{S} , episode length τ , prediction errors $\varepsilon_{\theta_i}((k-1)\tau)$, state x , scaling $\rho_k \in [0, 1]$

- 1: Compute weights using exponential weighting $w_{k+1}(i) = \exp(-\eta \varepsilon_{\theta_i}((k-1)\tau))$
- 2: Sample $\hat{i}_k \sim p_k(i)$, where $p_k(i) = w_k(i) / \sum_{j=0}^N w_k(j)$
- 3: Set $\hat{\theta}_k = \theta_{\hat{i}_k}$ and compute optimal input sequence $U_{\hat{\theta}_k}^*(x)$ by solving (26)
- 4: Define excitation according to ρ_k : $u_k^*(t) = \sqrt{1 - \rho_k} u_{\hat{\theta}_k}^*(t) + \sqrt{\rho_k} u_\eta(t)$, $u_\eta(t) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$

return excitation sequence u_k^*

phenomenon is well known in the literature of experiment design and is usually tackled by using sequential algorithms to generate excitations that are close to optimal. In line with these approaches, our proposed input design subroutine, which is presented as Algorithm 2, can be called by Algorithm 1 in each episode and returns an excitation input sequence. To derive the excitation, the subroutine draws an estimate $\hat{\theta}_k$ using exponential weights and solves the optimization (20) using certainty equivalence (CE), i.e., using $\hat{\theta}_k$ as if it were the true system. Thus, after sampling $\hat{\theta}_k = (\hat{A}_k, \hat{B}_k)$ and defining $\Delta\hat{A}_i := \hat{A}_k - A_i$ and $\Delta\hat{B}_i := \hat{B}_k - B_i$ the excitation input is computed by solving

$$U_{\hat{\theta}_k}^*(x(k\tau)) \in \arg \max_{U^\top U \leq \gamma_u^2 \tau} \min_{\theta_i \neq \hat{\theta}_k} \mathbb{E} \left[\sum_{t=k\tau}^{\tau(k+1)-1} \|\Delta\hat{A}_i x(t) + \Delta\hat{B}_i u(t)\|_{\Sigma_w^{-1}}^2 \right]. \quad (26)$$

Due to the uncertainty in the estimate, we add random excitations to the optimal excitation input according to a scaling parameter ρ_k , which can be adapted over time. We will analyze the choice of ρ_k and its consequences on the speed of identification throughout this section. Note that any $\rho_k \in [0, 1]$ yields excitation inputs satisfying Assumption 1. As shown above, establishing sample complexity upper bounds reduces to showing that Algorithm 2 produces PE inputs.

Lemma 4.1. *Consider the system (1), set \mathcal{S} as defined in (2). Consider the excitation input sequence u_k^* generated by Algorithm 2 at some fixed iterate k and let $\mathbb{P}[\hat{\theta}_k \neq \theta_*] \leq p_k \in [0, 1]$, where $\hat{\theta}_k$ is the estimate drawn using exponential weights in round k . Then u_k^* satisfies PE with*

$$c_{u_k}^{\text{Alg}}(\tau) = (1 - p_k)(1 - \rho_k)c_u^{\text{opt}}(\tau) + \rho_k c_u^{\text{rand}}(\tau), \quad (27a)$$

$$c_w(\tau) = \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_w(\tau) \right), \quad (27b)$$

where $c_u^{\text{opt}}(\tau)$ is the PE coefficient for optimal excitation as defined in (21) and $c_u^{\text{rand}}(\tau)$ is the PE coefficient for isotropic Gaussian excitations as defined in (19).

The proof of Lemma 4.1 can be found in Appendix D.3. Note that while it might seem intuitive that any convex combination of two PE inputs is also PE this is not necessarily true. Given the setup in Example 3.1 with $d = 1$ it is easy to see that $u_1^\top(t) = [1 \ 0]$, $\forall t \geq 0$ and $u_2^\top(t) = [-1 \ 0]$, $\forall t \geq 0$ are both PE. However, $u(t) = 0.5(u_1(t) + u_2(t)) = 0$ is clearly not PE. With Lemma 4.1 we can directly obtain the following sample complexity upper bound.

Theorem 4.2. *Let Algorithm 1 be used with the excitation input derived by Algorithm 2. Then, the estimate $\hat{\theta}$ of Algorithm 1 satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates no later than when k satisfies*

$$\tau \left(\sum_{j=1}^k c_{u_j}^{\text{Alg}}(\tau) + c_w(\tau) \right) \geq c' \log \left(\frac{N}{\delta} \right), \quad (28)$$

where $c_{u_j}^{\text{Alg}}(\tau)$, $c_w(\tau)$ are defined in (27) and c' is a constant influenced by the variance in ε_{θ_i} .

The full version of Theorem 4.2 and its proof are presented in Appendix E.4. In essence, Lemma 4.1 specifies that whenever the uncertainty is small enough compared to the expected benefit in PE from the optimal excitation CE will provably improve the PE coefficients compared to random Gaussian excitations. Theorem 4.2 then shows that experiment design using CE provably improves the sample complexity of identification in these cases. Ideally, ρ_k is selected to maximize (27a) which results in

$$\rho_k = \begin{cases} 0, & \text{if } (1 - p_k)c_u^{\text{opt}} \geq c_u^{\text{rand}} \\ 1, & \text{else.} \end{cases} \quad (29)$$

By (29) the optimal choice for ρ_k switches from one to zero once the confidence in the estimates is large enough relative to the expected benefit in PE. When $\rho_k = 1$, the learner relies only on isotropic Gaussian excitations to gain some initial knowledge from PE data. Since this might not be true for $U_{\theta_k}^*$, depending on the particular problem setup, the learner only switches to the CE-based excitation once the confidence is large enough. Since p_k decays to zero exponentially as shown in Theorem F.1, thus, the excitation selected by Algorithm 2 approaches the optimal oracle excitation as k grows.

Example 3.1 (continued). Recall that in Example 3.1 knowledge of the true system is not necessary to select the optimal excitation input sequence. This is due to the structure in \mathcal{S} , which yields greedy CE-based experiment design be optimal (see numerical evaluation in Appendix H.1). Deriving properties of \mathcal{S} resulting in this characteristic is an interesting topic for future research.

While these findings improve the understanding of the learning problem and the underlying mechanisms, the selection rule (29) cannot be implemented without oracle knowledge. In our numerical studies, selecting $\rho_k = 0$, i.e., greedily using CE, yielded the best results. In practice, ρ_k can be thought of as a safety feature to ensure PE. Thus, setting $\rho_0 = c > 0$ and decreasing it as k grows is a practical option to ensure PE (see Appendix H.2 for an evaluation of different selections of ρ_k).

5 Numerical experiments

To gain additional insights, we consider a problem which is only slightly more complicated than the setup in Example 3.1 but provides and reinforces some fundamental insights⁴. Specifically, we consider the set $\mathcal{S} = \{(A_*, B_*), (A_1, B_*), (A_2, B_*), (A_3, B_*)\}$ with $B_*^\top = [0_{2 \times 1} \quad I_2]$

$$A_* = \begin{bmatrix} 0 & 0.1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \quad A_1 = \begin{bmatrix} 0 & 0 & 0.1 \\ 0 & 0 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 & 0 & 0.1 \\ 0 & 0 & 0 \\ 0 & 0 & 0.8 \end{bmatrix} \quad A_3 = \begin{bmatrix} 0 & 0.1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0.8 \end{bmatrix}.$$

The noise variance and input power are given by $\Sigma_w = I_{n_x}$ and $\gamma_u = 1$, respectively. For the numerical evaluation, we consider Algorithm 1 where the excitation input is computed using Algorithm 2 with $\rho_k \equiv 0$ and $\eta = 0.01$, and compare it to the following data collection strategies used with Algorithm 1: 1) Algorithm 2 with oracle knowledge (in each episode (26) is solved using θ_*), 2) the optimization criterion in [27] with oracle knowledge, 3) Oracle optimal excitation for the considered data length (fully offline), and 4) $u(t) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{2} I_{n_u})$. We conduct 100 Monte Carlo simulations with 5 episodes and an episode length of 15 time steps. In Figure 1 we display the mean and $\frac{\sigma}{2}$ -band of the likelihood of θ_* given the data over the episodes for the different data collection procedures. It is easy to see that the optimal oracle excitation performs best. Further, since the setup is constructed in a way that input design yields significant benefits, isotropic Gaussian excitations perform significantly worse compared to all other input sequences. Interestingly, Algorithm 2 with $\rho_k \equiv 0$ performs almost as good as the offline oracle excitation and on par with the online oracle approach, even though it lacks the knowledge of the true system and relies on CE. This means that, even though Algorithm 2 does not necessarily plan with the correct system initially, the obtained data is still informative, and in particular significantly more informative than data that is collected using isotropic Gaussian excitations. Capturing this effect mathematically and understanding how it depends on the structure of the set \mathcal{S} is an interesting direction for future

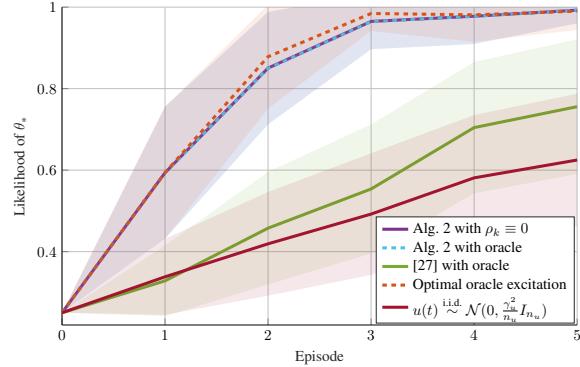


Figure 1: Mean and $\frac{\sigma}{2}$ -band of the likelihood of θ_* given the data for several data collection strategies.

⁴All experiments are carried out in Python using a standard notebook. Additional numerical evaluations, including a larger problem, are provided in Appendix H. The code for the numerical example can be accessed at: <https://github.com/col-tasas/2025-high-effort-low-gain>

research. While the active learning criterion in [27], which does not account for the finite hypothesis class, outperforms isotropic Gaussian excitations, it falls short of the criterion presented in this work. This highlights the importance of using the knowledge of the finite hypothesis class in input design.

6 Conclusion

In this work, we analyze the problem of identifying an unknown linear dynamical system from a finite hypothesis class. We present sample complexity lower bounds for isotropic Gaussian and arbitrary excitation, which can be used to determine the instance specific benefit of experiment design. We introduce the notion of PE which gives rise to a modular framework to establish sample complexity upper bounds for any excitation input that satisfies PE. Based on the findings we propose a CE-based algorithm for experiment design and derive sample complexity upper bounds. Numerical studies showcase that the proposed algorithm is highly competitive and achieves a performance close to the optimal oracle excitation. Numerically, it was observed that for hypothesis classes with a certain structure the algorithm was able to match the optimal oracle performance. It is an interesting direction for future research to investigate this further.

Acknowledgments

Nicolas Chatzikiriakos and Andrea Iannelli acknowledge the support by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 – 390740016 and by the Stuttgart Center for Simulation Science (SimTech).

References

- [1] Y. Abbasi-Yadkori, D. P., and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24, 2011.
- [2] X. Bombois, M. Gevers, R. Hildebrand, and G. Solari. Optimal experiment design for open and closed-loop system identification. *Commun. in Information and Systems*, 11(3):197–224, 2011.
- [3] N. Chatzikiriakos and A. Iannelli. Sample complexity bounds for linear system identification from a finite set. *IEEE Control Systems Letters*, 8:2751–2756, 2024.
- [4] N. Chatzikiriakos, B. Song, P. Rank, and A. Iannelli. Hidden convexity in active learning: A convexified online input design for ARX systems. *arXiv preprint arXiv:2509.03257*, 2025.
- [5] H. Chernoff. Sequential Design of Experiments. *The Annals of Mathematical Statistics*, 30(3):755 – 770, 1959.
- [6] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020.
- [7] D. Foster, T. Sarkar, and A. Rakhlin. Learning nonlinear dynamical systems from a single trajectory. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120, pages 851–861. PMLR, 2020.
- [8] A. Garivier, P. Ménard, and G. Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399, 2019.
- [9] L. Gerencsér and H. Hjalmarsson. Adaptive input design in system identification. In *44th IEEE Conf. on Decision and Control*, pages 4988–4993, 2005.
- [10] G. Goodwin and R. Payne. *Dynamic System Identification: Experiment Design and Data Analysis*. Mathematics in Science and Engineering. Academic Press, 1977.
- [11] Y. Jedra and A. Proutiere. Sample complexity lower bounds for linear system identification. In *IEEE 58th Conf. on Decision and Control (CDC)*, 2019.
- [12] Y. Jedra and A. Proutiere. Finite-time identification of linear systems: Fundamental limits and optimal algorithms. *IEEE Transactions on Automatic Control*, 2022.
- [13] J. Kiefer. General equivalence theory for optimum designs (approximate theory). *The Annals of Statistics*, 2(5):849–879, 1974.
- [14] J. Kiefer and J. Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960.
- [15] T. Lattimore and C. Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [16] B. D. Lee, I. Ziemann, G. J. Pappas, and N. Matni. Active learning for control-oriented identification of nonlinear systems. In *IEEE 63rd Conf. on Decision and Control (CDC)*, 2024.
- [17] H. Mania, M. I. Jordan, and B. Recht. Active learning for nonlinear system identification with guarantees. *J. of Machine Learning Research*, 23(32):1–30, 2022.
- [18] M. Muehlebach, Z. He, and M. I. Jordan. The sample complexity of online reinforcement learning: A multi-model perspective. *arXiv:2501.15910*, 2025.
- [19] P. Rigollet and J.-C. Hüttner. High-dimensional statistics. *arXiv:2310.19244*, 2023.

- [20] C. R. Rojas, J. S. Welsh, G. C. Goodwin, and A. Feuer. Robust optimal experiment design for system identification. *Automatica*, 43(6):993–1008, 2007.
- [21] T. Sarkar and A. Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conf. on Machine Learning*, pages 5610–5618. PMLR, 2019.
- [22] Y. Sattar and S. Oymak. Non-asymptotic and accurate learning of nonlinear dynamical systems. *Journal of Machine Learning Research*, 23(140):1–49, 2022.
- [23] Y. Sattar, S. Oymak, and N. Ozay. Finite sample identification of bilinear dynamical systems. In *IEEE 61st Conference on Decision and Control (CDC)*, pages 6705–6711, 2022.
- [24] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conf. On Learning Theory*, pages 439–473. PMLR, 2018.
- [25] A. Tsiamis and G. J. Pappas. Linear systems can be hard to learn. In *60th IEEE Conf. on Decision and Control (CDC)*. IEEE, 2021.
- [26] A. Tsiamis, I. Ziemann, N. Matni, and G. J. Pappas. Statistical learning theory for control: A finite-sample perspective. *IEEE Control Systems*, 43(6):67–97, 2023.
- [27] A. Wagenmaker and K. Jamieson. Active learning for identification of linear dynamical systems. In *Conf. on Learning Theory*, pages 3487–3582. PMLR, 2020.
- [28] A. J. Wagenmaker, M. Simchowitz, and K. Jamieson. Task-optimal exploration in linear dynamical systems. In *38th International Conf. on Machine Learning*, pages 10641–10652. PMLR, 2021.
- [29] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.

Contents

1	Introduction	1
2	Preliminaries	3
3	Can active learning algorithms improve identification?	4
3.1	Sample complexity lower bounds	4
3.2	Sample complexity upper bounds	5
3.2.1	Persistency of Excitation	5
3.2.2	High probability identification with finite samples	6
4	Sequential input design algorithm	7
5	Numerical experiments	9
6	Conclusion	10
A	Related Works	14
B	Proofs for sample complexity lower bounds	14
B.1	Proof of Theorem 3.1	15
B.2	Proof of Corollary 3.2	16
C	Existence of the optimal solution to (11)	16
D	Proving PE for different excitations	17
D.1	Proof of Lemma 3.3	17
D.2	Proof of Lemma 3.4	17
D.3	Proof of Lemma 4.1	17
E	Proofs for sample complexity upper bounds	18
E.1	Proof of Theorem 3.5	18
E.2	Intermediate results used in the proof of Theorem 3.5	20
E.3	Sample complexity Upper bounds for isotropic Gaussian and oracle excitations	23
E.4	Sample complexity upper bound for Algorithm 2	23
F	Convergence of Algorithm 2 to the optimal excitation	24
G	Auxiliary results	24
H	Additional numerical examples	26
H.1	Numerical evaluation of Example 3.1	26
H.2	Identification for a random candidate set	27

A Related Works

Finite Sample System Identification Novel results in high-dimensional statistics [1, 29] have sparked an increased interest in the finite sample analysis of system identification. Hereby the sample complexity of identification is of particular importance and has been analyzed mostly for random excitations and for an infinite hypothesis class using the closed-form solution of the ordinary least squares estimator (OLS) estimator. Since the data collected from dynamical systems is highly correlated, first works [6] only provided sample complexity upper bounds for the case where the data was collected from multiple trajectories. Subsequent works were able to overcome this restrictive assumption and provided sample complexity upper bounds for marginally stable [24] and also unstable linear dynamical systems [21] for trajectory data. Hereby it has been shown that the OLS is statistically inconsistent when the data is collected from certain classes of unstable systems. Additional works extended these works to certain classes of nonlinear dynamical systems such as bilinear systems [23] generalized linear systems [7, 22]. While sample complexity upper bounds are valuable since they establish identification error guarantees, sample complexity lower bounds are important to understand the hardness of the identification problem and to judge the tightness of the upper bounds. Building on information theoretic tools [12] provide sample complexity lower bounds that hold independently of the used algorithm. Similar tools have been used in [25] to understand when learning is hard with isotropic Gaussian excitations. All previous works consider the case where the hypothesis class is infinite. However, in many applications there might exist some prior knowledge restricting the hypothesis class to be finite, i.e., there exists a finite set of systems the true system belongs to. This setup has been previously considered in [3, 18] where it has been observed that stability of the true system does not seem to influence learning negatively as is the case for the infinite hypothesis class. While, the previously mentioned works consider several system classes and setups, they all assume the data is collected using (sub-)Gaussian inputs, which might be suboptimal in many cases.

Experiment Design: The design of optimal experiments has a long history in learning theory [5, 13, 14] where a large body of works exists concerning the question of how data should be sampled to obtain the most accurate estimate of an unknown quantity. One particular field of interest for experiment design is the identification of unknown dynamical systems, where compared to the classical setup in learning theory the data cannot be sampled arbitrarily but rather the unknown dynamical system governs the way the data can be collected (see [2, 10] for surveys on the topic). One key difficulty introduced by the dynamics of the unknown system is that the Fisher-information matrix naturally depends on the unknown system parameters. To tackle this problem, several algorithms have been proposed which can be broadly categorized into robust approaches that optimize for the worst-case using minmax objectives (see [20] and the references therein) and sequential approaches that rely on a running estimate [9]. While classical works, only consider the asymptotic case, recently several works considered the problem of experiment design for the identification of an unknown dynamical system under a finite sample perspective. In particular [27, 28] consider LTI systems and provide finite sample guarantees by leveraging sequential algorithms. Further, [4] provides an exact convex reformulation of the design criterion in [27], enabling the efficient application of the algorithm. The proposed algorithm optimizes over periodic excitation signals using the CE principle and augments the input with exploratory noise. Extending previous works [16, 17] consider certain classes of non-linear systems and provide finite sample guarantees for them. While optimality of the proposed experiment design algorithms is considered in the respective works, they do not analyze how large the benefit of experiment design is compared isotropic Gaussian excitations. Furthermore, the input space is restricted to periodic inputs, which although shown to be sufficient [27] is restrictive.

B Proofs for sample complexity lower bounds

Before we proceed with presenting the proofs of the sample complexity lower bound results in Section 3.1 we first provide the following input-dependent sample complexity result, which serves as the starting point for the subsequent results and is akin to [11, Theorem 2] which considers the case of an infinite hypothesis class instead of the finite hypothesis class considered in this work.

Theorem B.1 (Input-dependent sample complexity lower bound). *Consider the unknown dynamical system (1) and the hypothesis class \mathcal{S} as defined in (2). Then any δ -correct algorithm satisfies*

$$\mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq 2 \log \left(\frac{1}{2.4\delta} \right) \quad (30)$$

for any excitation input sequence $\{u(t)\}_{t=0}^{\bar{T}-1}$.

Proof. Define the data as $\mathcal{D}_{\bar{T}} := \{x(0), u(0), \dots, u(\bar{T}-1), x(\bar{T})\}$ and the probability of the observing $\mathcal{D}_{\bar{T}}$ under system θ_i as $\mathbb{P}_{\theta_i}(\mathcal{D}_{\bar{T}})$. Then, we define the log-likelihood ratio of the first \bar{T} observations under θ_* and some $\theta_i \in \mathcal{S} \setminus \{\theta_*\}$ as $L_{\bar{T}} = \log \left(\frac{\mathbb{P}_{\theta_*}(\mathcal{D}_{\bar{T}})}{\mathbb{P}_{\theta_i}(\mathcal{D}_{\bar{T}})} \right)$. Following the change of measurement argument in [11], we use the generalized data processing inequality [8, Lemma 1] to obtain the lower bound

$$\begin{aligned} \mathbb{E}[L_{\bar{T}}] &= \text{KL}(\mathbb{P}_{\theta_*}(\mathcal{D}_{\bar{T}}) \parallel \mathbb{P}_{\theta_i}(\mathcal{D}_{\bar{T}})) \\ &\geq \sup_{\mathcal{E} \in \mathcal{F}_{\bar{T}}} \text{kl}(\mathbb{P}_{\theta_*}(\mathcal{E}) \parallel \mathbb{P}_{\theta_i}(\mathcal{E})), \end{aligned}$$

where $\text{kl}(x\|y)$ is the KL-divergence of two Bernoulli distributions of means x and y , respectively. Since we analyze δ -correct algorithms we define the event $\mathcal{E} := \{\hat{\theta}_{\bar{T}} = \theta_*\}$ which yields $\mathbb{P}_{\theta_*}(\mathcal{E}) \geq 1 - \delta$ and $\mathbb{P}_{\theta_i}(\mathcal{E}) \leq \delta$ and hence

$$\text{kl}(\mathbb{P}_{\theta_*}(\mathcal{E}) \parallel \mathbb{P}_{\theta_i}(\mathcal{E})) \geq (2\delta - 1) \log \left(\frac{1 - \delta}{\delta} \right) \geq \log \left(\frac{1}{2.4\delta} \right). \quad (31)$$

Further, we follow [11, Section IV.A] to obtain

$$\begin{aligned} \mathbb{E}[L_{\bar{T}}] &= \frac{1}{2} \mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} [x(t)^\top \ u(t)^\top] \begin{bmatrix} \Delta A_i^\top \\ \Delta B_i^\top \end{bmatrix} \Sigma_w^{-1} [\Delta A_i \ \Delta B_i] \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \right] \\ &= \frac{1}{2} \mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right], \end{aligned}$$

which concludes the proof. \square

B.1 Proof of Theorem 3.1

The first statement of the theorem follows directly from the derivations in Section 2. To derive the optimal excitation input we leverage that by Theorem B.1 for any δ -correct algorithm it holds that

$$\mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq 2 \log \left(\frac{1}{2.4\delta} \right) \quad \forall i \in [1, N], \quad (32)$$

where $x(t)$ is generated by (1). Since (32) needs to hold for all $i \in [1, N]$ it follows immediately that it suffices to consider

$$\min_{i \in [1, N]} \mathbb{E} \left[\sum_{s=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq 2 \log \left(\frac{1}{2.4\delta} \right). \quad (33)$$

Since we seek a result that holds for all possible input sequences that satisfy $\frac{1}{\bar{T}} \sum_{t=0}^{\bar{T}-1} u(t)^\top u(t) \leq \gamma_u^2$ we maximize the l.h.s. of (40) over all admissible input sequences. This yields

$$\max_{U^\top U \leq \gamma_u^2} \min_{i \in [1, N]} \mathbb{E} \left[\sum_{s=0}^{t-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq 2 \log \left(\frac{1}{2.4\delta} \right). \quad (34)$$

Now, using Lemma G.1 with $\rho = 0$ and $M = \Sigma_w^{-1}$ we can equivalently rewrite the l.h.s. of (34) as

$$\max_{U^\top U \leq \gamma_u^2} \min_{i \in [1, N]} U^\top W_i(\bar{T}) U + \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_w(\bar{T}) \right) \quad (35)$$

and U^* is the corresponding optimizer. For the last result, consider that

$$\max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1, N]} U^\top W_i(\bar{T}) U + \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_w(\bar{T}) \right) \quad (36)$$

$$\leq \max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1, N]} U^\top W_i(\bar{T}) U + \max_{j \in [1, N]} \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^j(\bar{T}) S_w(\bar{T}) \right) \quad (37)$$

$$= \max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{p \in \Delta_N} U^\top \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right) U + \max_{j \in [1, N]} \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^j(\bar{T}) S_w(\bar{T}) \right) \quad (38)$$

$$= \min_{p \in \Delta_N} \max_{U^\top U \leq \gamma_u^2 \bar{T}} U^\top \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right) U + \max_{j \in [1, N]} \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^j(\bar{T}) S_w(\bar{T}) \right). \quad (39)$$

Selecting U^* in the direction of $v_{\max} \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right)$ yields the result.

B.2 Proof of Corollary 3.2

By Theorem B.1 for any δ -correct algorithm it holds that

$$\mathbb{E} \left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq 2 \log \left(\frac{1}{2.4\delta} \right) \quad \forall i \in [1, N], \quad (40)$$

where $x(t)$ is generated by (1). Applying Lemma G.1 with $\rho = 1$, $M = \Sigma_w^{-1}$ and $\sigma_u^2 = \frac{\gamma_u^2}{n_u}$ we obtain that for any δ -correct algorithm

$$\frac{\gamma_u^2}{n_u} \text{tr} \left(W_i(\bar{T}) \right) + \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_w(\bar{T}) \right) \geq 2 \log \left(\frac{1}{2.4\delta} \right) \quad \forall i \in [1, N]. \quad (41)$$

After noting that $\frac{1}{n_u \bar{T}} \text{tr} \left(W_i(\bar{T}) \right) = \lambda_{\text{mean}}(W_i(\bar{T}))$ the result follows immediately after realizing that it suffices to consider the minimum of the l.h.s. of (41).

C Existence of the optimal solution to (11)

For clarity of exposition for each $i \in [1, N]$ we define

$$W_i := R_{\Sigma_w^{-1}}^i(\bar{T}) + S_u(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_u(\bar{T}) + 2N_{\Sigma_w^{-1}}^i(\bar{T}) S_u(\bar{T}), \quad (42)$$

$$c_i := \text{tr} \left(S_w(\bar{T})^\top Q_{\Sigma_w^{-1}}^i(\bar{T}) S_w(\bar{T}) \right), \quad (43)$$

where $W_i \succeq 0$ and $c_i \geq 0$, for all $i \in [1, N]$. To show that the optimal solution to (11) exists we consider the equivalent optimization problem

$$\min_{U, \xi} \quad -\xi \quad (44a)$$

$$\text{s.t.} \quad U^\top U - \gamma_u^2 \bar{T} \leq 0 \quad (44b)$$

$$\xi - U^\top W_i U - c_i \leq 0 \quad \forall i \in [1, N], \quad (44c)$$

$$-\xi \leq 0 \quad (44d)$$

where ξ is a slack variable. Now we can leverage the following theorem due to Weierstrass.

Theorem C.1. *Consider the constrained optimization problem $\min_{x \in \mathcal{X}} f(x)$. If the objective function f is continuous and the feasible region \mathcal{X} is closed and bounded, then there exists a global optimum.*

Clearly $f(\xi, U) = -\xi$ is continuous. Further, the set of feasible input sequences U is closed and bounded by (44b). The feasible region of ξ is defined by the constraints (44c). Clearly, for a given \bar{T} under (44b) ξ is upper bounded by a finite, non-negative value. Combining this with (44d) the feasible region of ξ is bounded and closed, from which we can conclude that the optimal solution to (44) and hence also of (11) exists.

D Proving PE for different excitations

In this section we present all proofs related to PE in the same order the results are presented in the main part of the paper.

D.1 Proof of Lemma 3.3

Proof. Applying Lemma G.1 with $\rho = 1$ and $M = \Sigma_x^{-1}$ for any $i \in [1, N]$ yields

$$\frac{1}{\tau} \sum_{t=0}^{\tau} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \geq \frac{\gamma_u^2}{n_u \tau} \text{tr}(W_i(\tau)) + \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_w(\tau) \right).$$

Since we require a lower bound that hold uniformly over $i \in [1, N]$ we consider

$$\min_{i \in [1, N]} \frac{\gamma_u^2}{n_u \tau} \text{tr}(W_i(\tau)) + \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i(\tau) S_w(\tau) \right) \quad (45)$$

$$\geq \min_{i \in [1, N]} \frac{\gamma_u^2}{n_u \tau} \text{tr}(W_i(\tau)) + \min_{j \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^j(\tau) S_w(\tau) \right) \quad (46)$$

$$= \min_{p \in \Delta_N} \frac{\gamma_u^2}{n_u \tau} \text{tr} \left(\sum_{i=1}^N p_i W_i(\tau) \right) + \min_{j \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^j(\tau) S_w(\tau) \right) \quad (47)$$

$$= \min_{p \in \Delta_N} \lambda_{\text{mean}} \left(\sum_{i=1}^N p_i W_i(\tau) \right) \gamma_u^2 + \min_{j \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^j(\tau) S_w(\tau) \right). \quad (48)$$

Comparing terms with (18) yields the result. \square

D.2 Proof of Lemma 3.4

Proof. Applying Lemma G.1 with $\rho = 0$ and $M = \Sigma_w^{-1}$ for any $i \in [1, N]$ yields

$$\sum_{t=0}^{\tau-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2 \right] = U^\top W_i(\tau) U + 2U^\top m_i(x(0)) + c_i(x(0)) + \text{tr} \left(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau) \right) \quad (49)$$

for any $U \in \mathbb{R}^{n_u \tau}$. Thus, defining $\bar{c}_i = \text{tr} \left(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau) \right)$ by our choice of the excitation input we obtain

$$\max_{U^\top U \leq \gamma_u^2 \tau} \min_{i \in [1, N]} U^\top W_i(\tau) U + 2U^\top m_i(x_0) + c_i(x(0)) + \bar{c}_i \quad (50)$$

$$\geq \max_{U^\top U \leq \gamma_u^2 \tau} \min_{i \in [1, N]} U^\top W_i(\tau) U + \bar{c}_i \quad (51)$$

$$\geq \max_{U^\top U \leq \gamma_u^2 \tau} \min_{i \in [1, N]} U^\top W_i(\tau) U + \min_{j \in [1, N]} \bar{c}_j, \quad (52)$$

where the first inequality holds since $c_i(x(0))$ is non-negative by definition and $U^\top m_i(x_0)$ is non-negative when U is selected to maximize the expression. Similar to the proof of Theorem 3.1, solving the optimization yields the result. \square

D.3 Proof of Lemma 4.1

This proof is carried out for the case $k = 0$ without loss of generality. To obtain the result we analyze the expected excitation of the optimal control input based on true system and the expected excitation of a suboptimal control input based on a faulty estimate separately. We start with

$$\begin{aligned} \sum_{t=0}^{\tau-1} \mathbb{E}[\|\Delta A_i x(t) + \Delta B_i u_k^*(t)\|_{\Sigma_w^{-1}}^2] &= \mathbb{P}[\hat{\theta}_0 = \theta_*] \sum_{t=0}^{\tau-1} \mathbb{E}[\|\Delta A_i x(t) + \Delta B_i u_{\theta_*}^*(t)\|_{\Sigma_w^{-1}}^2] \\ &\quad + \mathbb{P}[\hat{\theta}_0 \neq \theta_*] \sum_{t=0}^{\tau-1} \mathbb{E}[\|\Delta A_i x(t) + \Delta B_i u_{\hat{\theta}}^*(t)\|_{\Sigma_w^{-1}}^2]. \end{aligned} \quad (53)$$

The first terms in (53) represents the excitation due to the optimal input sequence and hence after applying Lemma G.1 with $\nu = \rho_0$ and $M = \Sigma_w^{-1}$ we obtain

$$\begin{aligned} \sum_{t=0}^{\tau-1} \mathbb{E} \left[\left\| \Delta A_i x(t) + \Delta B_i u_{\theta_*}^*(t) \right\|_{\Sigma_w^{-1}}^2 \right] &\geq \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) \\ &\quad + (1 - \rho_0) (U_{\theta_*}^*)^\top W_i(\tau) U_{\theta_*}^* + \frac{\rho_0 \gamma_u^2}{n_u} \text{tr} (W_i(\tau)). \end{aligned} \quad (54)$$

To tackle the second term in (53) we follow the same steps to obtain

$$\begin{aligned} \sum_{t=0}^{\tau-1} \mathbb{E} \left[\left\| \Delta A_i x(t) + \Delta B_i u_{\bar{\theta}}^*(t) \right\|_{\Sigma_w^{-1}}^2 \right] &\geq \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) \\ &\quad + (1 - \rho_0) (U_{\bar{\theta}}^*)^\top W_i(\tau) U_{\bar{\theta}}^* + \frac{\rho_0 \gamma_u^2}{n_u} \text{tr} (W_i(\tau)). \end{aligned} \quad (55)$$

Plugging (54) and (55) into (53) and collecting terms we obtain

$$\begin{aligned} \sum_{t=1}^{\tau-1} \mathbb{E} [\| \Delta A_i x(t) + \Delta B_i u_k^*(t) \|_{\Sigma_w^{-1}}^2] &\geq \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) + \frac{\rho_0 \gamma_u^2}{n_u} \text{tr} (W_i(\tau)) \\ &\quad + \mathbb{P}[\hat{\theta}_0 = \theta_*] (1 - \rho_0) (U_{\theta_*}^*)^\top W_i(\tau) U_{\theta_*}^* + \mathbb{P}[\hat{\theta}_0 \neq \theta_*] (1 - \rho_0) (U_{\bar{\theta}}^*)^\top W_i(\tau) U_{\bar{\theta}}^*. \end{aligned} \quad (56)$$

The last term in (56) is non-negative. Hence, we obtain that

$$\begin{aligned} \sum_{t=1}^{\tau-1} \mathbb{E} [\| \Delta A_i x(t) + \Delta B_i u_k^*(t) \|_{\Sigma_w^{-1}}^2] &\geq \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) + \frac{\rho_0 \gamma_u^2}{n_u} \text{tr} (W_i(\tau)) + (1 - \rho_0) (1 - \rho_0) (U_{\theta_*}^*)^\top W_i(\tau) U_{\theta_*}^*, \end{aligned} \quad (57)$$

where we used that $\mathbb{P}[\hat{\theta}_0 = \theta_*] \geq 1 - p_0$. Dividing by $\tau > 0$ yields

$$\begin{aligned} \frac{1}{\tau} \sum_{t=1}^{\tau-1} \mathbb{E} [\| \Delta A_i x(t) + \Delta B_i u_k^*(t) \|_{\Sigma_w^{-1}}^2] &\geq \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) + \frac{\rho_0 \gamma_u^2}{\tau n_u} \text{tr} (W_i(\tau)) + (1 - p_0) (1 - \rho_0) \frac{1}{\tau} (U_{\theta_*}^*)^\top W_i(\tau) U_{\theta_*}^* \\ &\geq \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{\Sigma_w^{-1}}^i S_w(\tau) \right) + \frac{\rho_0 \gamma_u^2}{\tau n_u} \text{tr} (W_i(\tau)) + (1 - p_0) (1 - \rho_0) \frac{1}{\tau} (U_{\theta_*}^*)^\top W_i(\tau) U_{\theta_*}^*. \end{aligned} \quad (58)$$

Following similar steps as in the previous proofs yields the result.

E Proofs for sample complexity upper bounds

This section contains the proofs related to the sample complexity upper bounds and related intermediate results.

E.1 Proof of Theorem 3.5

We provide the full version of Theorem 3.5 and its proof.

Theorem E.1. *Consider the unknown system (1), set \mathcal{S} as defined in (2). Then Algorithm 1 yields an estimate $\hat{\theta}_T$ satisfying $\mathbb{P} [\hat{\theta}_T \neq \theta_*] \leq \delta$ and terminates at the latest for the first k satisfying*

$$\tau \sum_{j=1}^k c_{u_j}(\tau) \gamma_u^2 + c_w(\tau) \geq 8 \left(2a + \frac{1}{2\eta} \right) \log \left(\frac{N}{\delta} \right), \quad (59)$$

where

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)} \quad (60)$$

and $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (79).

Proof. This proof consists of two parts. First we show that if the algorithm terminates it holds that $\mathbb{P}[\hat{\theta}_T \neq \theta_*] \leq \delta$, i.e., we show the algorithm is δ -correct. Then derive the upper bound on the stopping time and hence on the sample complexity.

Correctness: For the remainder of this proof we use \mathbb{P}_{θ_i} and \mathbb{E}_{θ_i} to denote probability and expectation of an event under the hypothesis that $\theta_* = \theta_i$. To analyze correctness of the algorithm we analyze the likelihood-ratio

$$L_{\theta_i, \theta_j}(t) = \frac{\mathbb{P}_{\theta_i}[\mathcal{D}_t]}{\mathbb{P}_{\theta_j}[\mathcal{D}_t]} = \frac{\exp\left(-\frac{1}{2}\varepsilon_{\theta_i}(t)\right)}{\exp\left(-\frac{1}{2}\varepsilon_{\theta_j}(t)\right)}. \quad (61)$$

It can easily be shown that the likelihood-ratio is a martingale sequence (see, e.g., [29, Example 2.18]). By the termination rule the algorithm terminates under the event

$$\mathcal{E} = \left\{ \exists \hat{\theta} : \log(L_{\hat{\theta}, \theta_i}(t)) > \log\left(\frac{N}{\delta}\right), \quad \forall \theta_i \in \mathcal{S} \setminus \{\hat{\theta}\} \right\} \quad (62)$$

$$= \left\{ \exists \hat{\theta} : L_{\hat{\theta}, \theta_i}(t) > \frac{N}{\delta}, \quad \forall \theta_i \in \mathcal{S} \setminus \{\hat{\theta}\} \right\}. \quad (63)$$

Thus, $\hat{\theta} \neq \theta_*$ requires $L_{\theta_i, \theta_0}(t) \geq \frac{N}{\delta}$ for at least one $\theta_i \neq \theta_*$ and some $t \in \mathbb{Z}_+$. Thus

$$\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \mathbb{P}_{\theta_0} \left[\bigcup_{i=1}^N \exists t \in \mathbb{Z}_+ : L_{\theta_i, \theta_0}(t) \geq \frac{N}{\delta} \right]. \quad (64)$$

Thus,

$$\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \mathbb{P}_{\theta_0} \left[\bigcup_{i=1}^N \exists t \in \mathbb{Z}_+ : L_{\theta_i, \theta_0}(t) \geq \frac{N}{\delta} \right] \quad (65)$$

$$\leq \sum_{i=1}^N \mathbb{P}_{\theta_0} \left[\exists t \in \mathbb{Z}_+ : L_{\theta_i, \theta_0}(t) \geq \frac{N}{\delta} \right] \leq \sum_{i=1}^N \frac{\delta}{N} = \delta, \quad (66)$$

where the second inequality uses a union bound, and the last inequality follows from [15, Theorem 3.9].

Stopping Time: Since in each epoch, the input is PE with coefficients $c_{u_j}(\tau)$ and $c_w(\tau)$ we can use Proposition E.2 to show that under the choice of η the sequence

$$S_i(k) = \exp \left(-\eta \left(\varepsilon_{\theta_i}(t_k) - \varepsilon_{\theta_*}(t_k) - \frac{1}{4} \sum_{j=0}^{k-1} \sum_{t=t_j}^{t_{j+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \mid \mathcal{F}_j \right] \right) \right) \quad (67)$$

with $S_i(0) = 1$ is a super-martingale. Hence, again using the maximal inequality [15, Theorem 3.9] we obtain

$$\mathbb{P} \left[\exists k : S_i(k) \geq \frac{N}{\delta} \right] \leq \frac{\delta}{N} \quad (68)$$

Thus using union bound arguments we obtain

$$\mathbb{P} \left[\bigcup_{i=1}^N \exists k : S_i(k) \geq \frac{N}{\delta} \right] \leq \sum_{i=1}^N \mathbb{P} \left[\exists k : S_i(k) \geq \frac{N}{\delta} \right] \leq \delta \quad (69)$$

Note that, given the definition of PE, $S_i(k) \geq \frac{N}{\delta}$ is implied by

$$\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k) \geq \frac{1}{\eta} \log \left(\frac{N}{\delta} \right) - \frac{\tau}{4} \sum_{j=0}^{k-1} c_{u_j}(\tau) \gamma_u^2 + c_w(\tau). \quad (70)$$

Thus by (69) it holds with probability at least $1 - \delta$ that

$$\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k) \leq \frac{1}{\eta} \log \left(\frac{N}{\delta} \right) - \frac{\tau}{4} \sum_{j=0}^{k-1} c_{u_j}(\tau) \gamma_u^2 + c_w(\tau). \quad (71)$$

Thus, as soon as

$$8 \left(\frac{1}{2\eta} + 2 \right) \log \left(\frac{N}{\delta} \right) \leq \tau \sum_{j=0}^{k-1} c_{u_j}(\tau) \gamma_u^2 + c_w(\tau) \quad (72)$$

we obtain

$$\log(L_{\theta_i, \theta_*}(k\tau)) = \frac{1}{2} (\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k)) \quad (73)$$

$$\leq \frac{1}{2\eta} \log \left(\frac{N}{\delta} \right) - \frac{\tau}{8} \sum_{j=0}^{k-1} c_{u_j}(\tau) \gamma_u^2 + c_w(\tau) \quad (74)$$

$$\leq \frac{1}{2\eta} \log \left(\frac{N}{\delta} \right) - \left(\frac{1}{2\eta} + 2 \right) \log \left(\frac{N}{\delta} \right) \quad (75)$$

$$= 2 \log \left(\frac{N}{\delta} \right) \quad (76)$$

Thus, when (72) is satisfied Algorithm 1 terminates. This happens at the latest when (59) is satisfied and yields $\hat{\theta} = \theta_*$ with probability at least $1 - \delta$. \square

E.2 Intermediate results used in the proof of Theorem 3.5

In the following, we present intermediate results which are used in the proof of Theorem 3.5.

Proposition E.2. *Consider the unknown system (1), and the set \mathcal{S} as defined in (2). Let the data be collected, by exciting the true system with $u(t) = u_d(t) + u_r(t)$, where $u_d(t)$ is a deterministic sequence and $u_r(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \Sigma_u)$. Divide the data into K blocks of length $\tau > 0$ and define $t_k = \tau k$, $k \in [0, K]$. Further, define the filtration \mathcal{F}_k consisting of all random variables observed at the start of block k , i.e., $\mathcal{F}_k = \{x(0), u(0), \dots, u(t_k - 1), x(t_k)\}$. Then the sequence*

$$S_i(k) = \exp \left(-\eta \left(\varepsilon_{\theta_i}(t_k) - \varepsilon_{\theta_*}(t_k) - \frac{1}{4} \sum_{j=0}^{k-1} \sum_{t=t_j}^{t_{j+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 | \mathcal{F}_j \right] \right) \right) \quad (77)$$

with $S_i(0) = 1$ is a supermartingale⁵ for any $\eta > 0$ satisfying

$$\eta\tau \leq \frac{1}{512\lambda_{\max} \left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau) \right)^{\top} \Sigma_w^{-1} \Sigma_{\Delta_i}^{\frac{1}{2}}(\tau) \right)} \quad (78)$$

with

$$\Sigma_{\Delta_i}(\tau) := \Delta A_i \left(\sum_{s=0}^{\tau-1} A_*^{\tau-s} B \Sigma_u B^{\top} (A^{\tau-s})^{\top} + A_*^{\tau-s} \Sigma_w (A_*^{\tau-s})^{\top} \right) \Delta A_i^{\top} + \Delta B_i \Sigma_u \Delta B_i^{\top}. \quad (79)$$

Proof. For the proof of this result, we overload notation and define

$$\varepsilon_{\theta_i}(t_0, t_1) := \sum_{t=t_0}^{t_1-1} \|x(t+1) - A_i x(t) - B_i u(t)\|_{\Sigma_w^{-1}}^2. \quad (80)$$

⁵A supermartingale is a sequence $X(0), X(1), \dots$, of integrable random variables satisfying $\mathbb{E}[X(k+1)|\mathcal{F}_k] \leq X(k)$.

To show $S_i(k)$ is a supermartingale we consider

$$\begin{aligned} \mathbb{E}[S_i(k+1)|\mathcal{F}_k] &= S_i(k) \mathbb{E} \left[\exp \left(-\eta \left(\varepsilon_{\theta_i}(t_k, t_{k+1}) \right. \right. \right. \\ &\quad \left. \left. \left. - \varepsilon_{\theta_*}(t_k, t_{k+1}) - \frac{1}{4} \left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \right) \right) \right) \Big| \mathcal{F}_k \right]. \end{aligned} \quad (81)$$

Thus, to show the result requires showing that

$$\mathbb{E} \left[e^{-\eta(\varepsilon_{\theta_i}(t_k, t_{k+1}) - \varepsilon_{\theta_*}(t_k, t_{k+1}) - \frac{1}{4} \left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \right))} \Big| \mathcal{F}_k \right] \leq 1. \quad (82)$$

To this end, we define

$$\xi_{\theta_i}(t) := \|x(t+1) - A_i x(t) - B_i u(t)\|_{\Sigma_w^{-1}}^2. \quad (83)$$

Note that

$$\mathbb{E} \left[e^{-\eta(\varepsilon_{\theta_i}(t_k, t_{k+1}) - \varepsilon_{\theta_*}(t_k, t_{k+1}))} \Big| \mathcal{F}_k \right] = \mathbb{E} \left[e^{-\eta \sum_{t=t_k}^{t_{k+1}-1} \xi_{\theta_i}(t) - \xi_{\theta_*}(t)} \Big| \mathcal{F}_k \right] \quad (84)$$

$$\leq \prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E} \left[e^{-\eta \tau (\xi_{\theta_i}(t) - \xi_{\theta_*}(t))} \Big| \mathcal{F}_k \right] \right)^{\frac{1}{\tau}} \quad (85)$$

$$\leq \prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E} \left[e^{-\frac{\eta}{2} \tau \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2} \Big| \mathcal{F}_k \right] \right)^{\frac{1}{\tau}}, \quad (86)$$

where the first inequality follows from Hölder's inequality and the last inequality follows from Lemma E.3 given that $\eta \leq \frac{1}{4}$. Note now that conditioned on \mathcal{F}_k we have $\Delta A_i x(t_k + t) + \Delta B_i u(t_k + t) \sim \mathcal{N}(\mu_{\Delta_i}(t), \Sigma_{\Delta_i}(t))$, where $\Sigma_{\Delta_i}(t)$ is defined in (79) and $\mu_{\Delta_i}(t)$ depends on $x(t_k)$, the input sequence $u_d(t)$, and the system matrices. Hence, conditioned on \mathcal{F}_k , by Proposition E.4 we have

$$\begin{aligned} &\frac{1}{2} \left(\|\Delta A_i x(t_k + t) + \Delta B_i u(t_k + t)\|_{\Sigma_w^{-1}}^2 \right. \\ &\quad \left. - \mathbb{E} \left[\|\Delta A_i x(t_k + t) + \Delta B_i u(t_k + t)\|_{\Sigma_w^{-1}}^2 \right] \right) \sim \text{subExp}(\nu(t)), \forall t \in [0, \tau - 1] \end{aligned} \quad (87)$$

with $\nu(t) = 16\lambda_{\max} \left((\Sigma_{\Delta_i}(t)^{\frac{1}{2}})^{\top} \Sigma_w^{-1} \Sigma_{\Delta_i}(t)^{\frac{1}{2}} \right)$. Taking $\eta\tau \leq \frac{1}{\nu(\tau)}$ we can use the definition of subexponential random variables to obtain

$$\prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E} \left[e^{-\frac{\eta}{2} \tau \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2} \Big| \mathcal{F}_k \right] \right)^{\frac{1}{\tau}} \quad (88)$$

$$\leq \left(\prod_{t=t_k}^{t_{k+1}-1} e^{\frac{\eta^2}{2} \tau^2 \nu(t)^2} e^{-\frac{\eta}{2} \tau \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right]} \right)^{\frac{1}{\tau}} \quad (89)$$

$$= \left(e^{-\frac{\eta}{2} \tau \left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] - \tau \eta \nu^2(t) \right)} \right)^{\frac{1}{\tau}} \quad (90)$$

$$= e^{-\frac{\eta}{2} \left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] - \tau \eta \nu^2(t) \right)}. \quad (91)$$

Imposing

$$\tau \eta \nu^2(t) \leq \frac{1}{2} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right] \quad (92)$$

for $t \in [t_k, t_{k+1} - 1]$ and combining the above arguments results in

$$\mathbb{E} \left[e^{-\eta(\varepsilon_{\theta_i}(t_k, t_{k+1}) - \varepsilon_{\theta_*}(t_k, t_{k+1}))} \Big| \mathcal{F}_k \right] \leq \exp \left(-\frac{\eta}{4} \sum_{t=t_k}^{t_{k+1}-1} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \Big| \mathcal{F}_k \right] \right) \leq 1, \quad (93)$$

where the last inequality follows since $\mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 | \mathcal{F}_k \right] > 0, \forall t \in [t_k, t_{k+1} - 1]$. Note finally that

$$\mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 | \mathcal{F}_k \right] \quad (94)$$

$$= \mathbb{E} \left[(\Delta A_i x(t) + \Delta B_i u(t))^{\top} \Sigma_w^{-1} (\Delta A_i x(t) + \Delta B_i u(t)) | \mathcal{F}_k \right] \quad (95)$$

$$= \text{tr} \left(\Sigma_w^{-1} \mathbb{E} \left[(\Delta A_i x(t) + \Delta B_i u(t)) (\Delta A_i x(t) + \Delta B_i u(t))^{\top} | \mathcal{F}_k \right] \right) \quad (96)$$

$$= \text{tr} \left(\Sigma_w^{-1} (\Sigma_{\Delta_i}(t) + \mu_{\Delta_i}(t) \mu_{\Delta_i}(t)^{\top}) \right) \quad (97)$$

$$= \text{tr} \left((\Sigma_{\Delta_i}(t)^{\frac{1}{2}})^{\top} \Sigma_w^{-1} \Sigma_{\Delta_i}(t)^{\frac{1}{2}} + \mu_{\Delta_i}(t)^{\top} \Sigma_w^{-1} \mu_{\Delta_i}(t) \right) \quad (98)$$

$$\geq \lambda_{\max} \left((\Sigma_{\Delta_i}(t)^{\frac{1}{2}})^{\top} \Sigma_w^{-1} \Sigma_{\Delta_i}(t)^{\frac{1}{2}} \right) = \frac{\nu(t)}{16}, \quad (99)$$

hence condition (92) is implied by

$$\tau \eta \nu^2(t) \leq \frac{\nu(t)}{32} \stackrel{\nu(t) > 0}{\iff} \tau \eta \nu(t) \leq \frac{1}{32} \forall t \in [0, \tau] \quad (100)$$

or since the $\nu(t)$ only grows as t increases by

$$\tau \eta \leq \frac{1}{32 \nu(\tau)} = \frac{1}{512 \lambda_{\max} \left((\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau))^{\top} \Sigma_w^{-1} \Sigma_{\Delta_i}^{\frac{1}{2}}(\tau) \right)}. \quad (101)$$

□

Lemma E.3. Define

$$\xi_{\theta_i}(t) := \|x(t+1) - A_i x(t) - B_i u(t)\|_{\Sigma_w^{-1}}^2 \quad (102)$$

and let \mathcal{F}_k be a filtration containing all random variables up to time k . Then, if $\eta \leq \frac{1}{4}$ we have

$$\mathbb{E} \left[e^{-\eta(\xi_{\theta_i}(t) - \xi_{\theta_*}(t))} | \mathcal{F}_k \right] \leq \mathbb{E} \left[\exp \left(-\frac{\eta}{2} \|\Delta A_i x(t) - \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right) | \mathcal{F}_k \right] \quad \forall k = 1, 2, \dots, t. \quad (103)$$

Proof. By the definition of $\xi_{\theta_i}(t)$ we have

$$\xi_{\theta_i}(t) = \|\Delta A_i x(t) - \Delta B_i u(t) + w(t)\|_{\Sigma_w^{-1}}^2 \quad (104)$$

$$= \|\Delta A_i x(t) - \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 + 2w(t)^{\top} \Sigma_w^{-1} (A_i x(t) - \Delta B_i u(t)) + w(t)^{\top} \Sigma_w^{-1} w(t) \quad (105)$$

and hence

$$\xi_{\theta_i}(t) - \xi_{\theta_*}(t) = \|\Delta A_i x(t) - \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 + 2w(t)^{\top} \Sigma_w^{-1} (A_i x(t) - \Delta B_i u(t)). \quad (106)$$

Conditioning on $x(t)$, $u(t)$ the randomness in $\xi_{\theta_i}(t) - \xi_{\theta_*}(t)$ is due to $2w(t)^{\top} \Sigma_w^{-1} (A_i x(t) - \Delta B_i u(t))$. Realize that $w(t)^{\top} \Sigma_w^{-\frac{1}{2}} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{n_x})$. Recall that the definition of the moment generating function for a random variable $w \sim \mathcal{N}(0, I_{n_x})$ is given by

$$M_w(\lambda) = \mathbb{E} \left[e^{w^{\top} \lambda} \right] = \exp \left(\frac{1}{2} \lambda^{\top} \lambda \right). \quad (107)$$

Using this definition with $\lambda = 2\eta \Sigma_w^{-\frac{1}{2}} (A_i x(t) - \Delta B_i u(t))$ we obtain

$$\begin{aligned} & \mathbb{E} \left[\exp(-\eta(\xi_{\theta_i}(t) - \xi_{\theta_*}(t))) | x(t), u(t) \right] \\ &= \exp \left(-\eta \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right) \mathbb{E} \left[\exp \left(-2\eta w(t)^{\top} \Sigma_w^{-1} (A_i x(t) + \Delta B_i u(t)) \right) | x(t), u(t) \right] \\ &= \exp \left(-\eta \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 + 2\eta^2 \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 \right) \\ &= \exp \left(-\eta \|\Delta A_i x(t) + \Delta B_i u(t)\|_{\Sigma_w^{-1}}^2 (1 - 2\eta) \right). \end{aligned}$$

Taking $\eta \leq \frac{1}{4}$ the desired bound holds. □

Proposition E.4. Consider the random Gaussian vector $z \sim \mathcal{N}(\mu_z, \Sigma_z)$ of dimension n_z and a symmetric matrix $M \in \mathbb{S}_{++}^{n_z}$. Then $\frac{1}{2} (\|z\|_M^2 - \mathbb{E} [\|z\|_M^2])$ is sub-exponential⁶ with parameter ν ,

⁶A random variable X is said to be sub-exponential with parameter ν (denoted by $X \sim \text{subG}(\nu)$) if $\mathbb{E}[X] = 0$ and its moment generating function satisfies $\mathbb{E}[e^{sX}] \leq e^{\frac{s^2 \nu^2}{2}}$, $\forall |s| \leq \frac{1}{\nu}$.

where

$$\nu = 16\lambda_{\max} \left(\Sigma_z^{\frac{1}{2}} M \Sigma_z^{\frac{1}{2}} \right). \quad (108a)$$

Proof. Define $\zeta = \Sigma_z^{-\frac{1}{2}} (z - \mu_z)$, yielding $\|z\|_M = \|\zeta + \Sigma_z^{-\frac{1}{2}} \mu_z\|_{\Sigma_z^{\frac{1}{2}} M \Sigma_z^{\frac{1}{2}}}$ with $\zeta \sim \mathcal{N}(0, I_{n_z})$. Note that the $f(\zeta) = \frac{1}{\sqrt{2}} \|\zeta + \mu_z\|_M$ is Lipschitz continuous with Lipschitz constant $L = \sqrt{\lambda_{\max}(M)}$. Thus, by [29, Theorem 2.26]

$$\|\zeta + \mu_z\|_M \sim \text{subG} \left(\lambda_{\max} \left(\Sigma_z^{\frac{1}{2}} M \Sigma_z^{\frac{1}{2}} \right) \right). \quad (109)$$

Applying [19, Lemma 1.12] we obtain the result. \square

E.3 Sample complexity Upper bounds for isotropic Gaussian and oracle excitations

In the following, we present the full version of Corollary 3.6 along with its proof.

Corollary E.5. *Consider the same setup as in Theorem 3.5. If $u(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$ then Algorithm 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates no later than when $T = k\tau$ satisfies*

$$T \left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\text{mean}} \left(\sum_{i=1}^N p_i W_i(\tau) \right) + \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{I_{n_x}}^i S_w(\tau) \right) \right) \quad (110)$$

$$\geq 8 \left(2 + \frac{1}{2\eta} \right) \log \left(\frac{N}{\delta} \right). \quad (111)$$

Further, if the optimal oracle excitation input U^* defined in (20) is applied the estimate $\hat{\theta}$ satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$, and Algorithm 1 terminates no later than when $T = k\tau$ satisfies

$$T \left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\text{max}} \left(\sum_{i=1}^N p_i W_i(\tau) \right) + \min_{i \in [1, N]} \frac{1}{\tau} \text{tr} \left(S_w(\tau)^\top Q_{I_{n_x}}^i S_w(\tau) \right) \right) \quad (112)$$

$$\geq 8 \left(2 + \frac{1}{2\eta} \right) \log \left(\frac{N}{\delta} \right). \quad (113)$$

For both results we have

$$\eta\tau \leq \frac{1}{512\lambda_{\max} \left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau) \right)^\top \Sigma_w^{-1} \Sigma_{\Delta_i}^{\frac{1}{2}}(\tau) \right)}, \quad (114)$$

where $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (79).

Proof. The result follows directly by using Theorem E.1 with the PE coefficients derived in Lemmas 3.3 and 3.4. \square

E.4 Sample complexity upper bound for Algorithm 2

In the following, we present the full version of Theorem 4.2 along with its proof. Note that, the Theorem is presented in full generality, to provide theoretical insights. Depending on the update rule for ρ_k corollaries can be derived directly.

Theorem E.6. *Let Algorithm 1 be used with the excitation input derived by Algorithm 2. Then, the estimate $\hat{\theta}$ of Algorithm 1 satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates no later than when k satisfies*

$$\tau \left(\sum_{j=1}^k c_{u_j}^{\text{Alg}}(\tau) + c_w(\tau) \right) \geq 8 \left(2 + \frac{1}{2\eta} \right) \log \left(\frac{N}{\delta} \right), \quad (115)$$

where

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)}, \quad (116)$$

with $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ as defined in (79) and $c_{u_j}^{\text{Alg}}(\tau)$ and $c_w(\tau)$ as defined in (27).

Proof. The result follows directly by using Theorem E.1 and the PE coefficients derived in Lemma 4.1. \square

F Convergence of Algorithm 2 to the optimal excitation

Theorem F.1. Consider the unknown system (1), set \mathcal{S} as defined in (2) and Algorithm 2, where the input applied in round k is PE with $c_{u_j}(\tau)$ and $c_w(\tau)$. Let η chosen to satisfy

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)} \quad (117)$$

where $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (79). Then the estimate $\hat{\theta}_k$ drawn in Line 2 of Algorithm 2 satisfies

$$\mathbb{P}\left[\hat{\theta}_k = \theta_i\right] \leq N \exp\left(-\frac{\eta\tau}{4} \sum_{j=1}^k c_{u_j}(\tau)\gamma_u^2 + c_w(\tau)\right). \quad (118)$$

Proof. For the proof of this result, we overload notation and define

$$\varepsilon_{\theta_i}(t_0, t_1) := \sum_{t=t_0}^{t_1-1} \|x(t+1) - A_i x(t) - B_i u(t)\|_{\Sigma_w^{-1}}^2. \quad (119)$$

First note, that by similar arguments as in [18, Lemma B.1] we have

$$\mathbb{P}\left[\hat{\theta}_k = \theta_i\right] \leq \mathbb{E}\left[e^{-\eta(\varepsilon_{\theta_i}(0, k\tau) - \varepsilon_{\theta_*}(0, k\tau))}\right]. \quad (120)$$

By using Proposition E.2 we obtain

$$\mathbb{P}\left[\hat{\theta}_k = \theta_i\right] \leq \mathbb{E}\left[e^{-\eta(\varepsilon_{\theta_i}(0, k\tau) - \varepsilon_{\theta_*}(0, k\tau))}\right] \leq \exp\left(-\frac{\eta\tau}{4} \sum_{j=1}^k c_{u_j}(\tau)\gamma_u^2 + c_w(\tau)\right). \quad (121)$$

Taking a union bound for all $i \in [1, N]$ and solving for δ yields the result. \square

Note that the result can also be used to derive a similar result for the estimation of the true system using the softmax maximum likelihood estimator (MLE), similar to the results in [18].

G Auxiliary results

The following result is a consequence of the superposition principle in LTI systems and leverages that all random quantities are independent. It is a core ingredient in the proofs of the sample complexity lower and upper bounds.

Lemma G.1. Consider the system (1) with $x(0) \in \mathbb{R}^{n_x}$. Let $u(t) = (1 - \rho)u_d(t) + \rho u_p(t)$, with $\rho \in [0, 1]$, $u_p(t) \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma_u^2 I_{n_u})$ and a deterministic control input $u_d(t)$. Then for any fixed matrix $M \in \mathbb{S}_{++}^{n_x}$, any $i \in [1, N]$, and any $\tau \geq 0$ it holds that

$$\begin{aligned} \sum_{t=0}^{\tau} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2\right] &= (1 - \rho)^2 U_d^\top W_i(\tau) U_d + 2(1 - \rho) U_d^\top m_i(x_0) + c_i(x(0)) \\ &\quad + \sigma_u^2 \rho^2 \text{tr}(W_i(\tau)) + \text{tr}(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau)), \end{aligned} \quad (122)$$

where

$$\begin{aligned}
W_i(\tau) &:= R_M^i(\tau) + S_u(\tau)^\top Q_M^i(\tau) S_u(\tau) + N_M^i(\tau) S_u(\tau) + (N_M^i(\tau) S_u(\tau))^\top \\
m_i(x(0)) &:= (S_u(T)^\top Q_M^i(\tau) S_w(\tau) X(0) + N_M^i(\tau) S_w(\tau) X(0)) \\
c_i(x(0)) &:= \sum_{t=0}^{\tau} \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right)^\top M \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right) \\
X(0)^\top &:= \begin{bmatrix} (\Sigma_w^{-\frac{1}{2}} x)^\top & 0 & \dots & 0 \end{bmatrix}.
\end{aligned}$$

Furthermore,

- If $x(0) = 0$ then it holds for all $\rho \in [0, 1]$ that

$$\begin{aligned}
\sum_{t=0}^{\tau} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2 \right] &\geq (1 - \rho)^2 U_d^\top W_i(\tau) U_d + \sigma_u^2 \rho^2 \text{tr}(W_i(\tau)) \\
&\quad + \text{tr}(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau))
\end{aligned} \tag{123}$$

- If $\rho = 1$ then it holds for all $x(0) \in \mathbb{R}^{n_x}$ that

$$\sum_{t=0}^{\tau} \mathbb{E} \left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2 \right] \geq \sigma_u^2 \text{tr}(W_i(\tau)) + \text{tr}(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau)) \tag{124}$$

Proof. By recursively plugging in the $x(t) = A_* x(t-1) + B u(t-1) + w(t-1)$ we obtain

$$\Delta A_i x(t) + \Delta B_i u(t) = \Delta A_i x(t) + (1 - \rho) \Delta B_i u_d(t) + \rho \Delta B_i u_p(t) \tag{125}$$

$$\begin{aligned}
&= \Delta A_i \underbrace{\left(\sum_{s=0}^{t-1} (1 - \rho) A_*^{s-t-1} B_* u_d(s) + A_*^s x(0) \right)}_{\text{deterministic}} + (1 - \rho) \Delta B_i u_d(t) \\
&\quad + \rho \Delta A_i \underbrace{\sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_p(t)}_{\text{rand. vec., dep. on } u_p} + \rho \Delta B_i u_p(t) + \Delta A_i \underbrace{\sum_{s=0}^{t-1} A_*^{s-t-1} w(t)}_{\text{rand. vec., dep. on } w}.
\end{aligned} \tag{126}$$

Using that all, random vectors are independent and zero mean, and thus, cross terms are zero, it follows by plugging in (126) that

$$\begin{aligned}
&\sum_{t=0}^{\tau} \mathbb{E}[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2] \\
&= \sum_{t=0}^{\tau} \mathbb{E} \left[(\Delta A_i x(t) + \Delta B_i u(t))^\top M (\Delta A_i x(t) + \Delta B_i u(t)) \right] \\
&= \sum_{t=0}^{\tau} \left(\left\| \Delta A_i \left((1 - \rho) \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_d(s) + A_*^s x(0) \right) + (1 - \rho) \Delta B_i u_d(t) \right\|_M^2 \right. \\
&\quad \left. + \rho^2 \mathbb{E} \left[\left\| \Delta B_i u_p(t) + \Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_p(s) \right\|_M^2 \right] \right. \\
&\quad \left. + \mathbb{E} \left[\left\| \Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} w(s) \right\|_M^2 \right] \right).
\end{aligned} \tag{127}$$

We carry on by rewriting the first term in (127)

$$\begin{aligned}
& \sum_{t=0}^{\tau} \left\| \Delta A_i \left((1-\rho) \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_d(s) + A_*^s x(0) \right) + (1-\rho) \Delta B_i u_d(t) \right\|_M^2 \\
&= (1-\rho)^2 U_d^\top W_i(\tau) U_d + \underbrace{\sum_{t=0}^{\tau} \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right)^\top M \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right)}_{:=c_i(x(0)) \geq 0} \\
&\quad + 2(1-\rho) U_d^\top \underbrace{\left(S_u(\tau)^\top Q_M^i(\tau) S_w(\tau) X(0) + N_M^i(\tau) S_w(\tau) X(0) \right)}_{:=m_i(x_0)},
\end{aligned} \tag{128}$$

where $X(0)^\top = \begin{bmatrix} (\Sigma_w^{-\frac{1}{2}} x)^\top & 0 & \dots & 0 \end{bmatrix}$. Let us first consider $x(0) = 0$. Then we can use the matrices defined in (6) and (8) and plug them into (128) to obtain

$$\begin{aligned}
\sum_{t=0}^{\tau} \mathbb{E} \left[\left\| \Delta A_i x(t) + \Delta B_i u(t) \right\|_M^2 \right] &= (1-\rho)^2 U_d^\top W_i(\tau) U_d + 2(1-\rho) U_d^\top m_i(x_0) + c_i(x(0)) \\
&\quad + \rho^2 \mathbb{E} [U_p W_i(\tau) U_p] + \mathbb{E} [W^\top S_w(\tau)^\top Q_M^i(\tau) S_w(\tau) W]
\end{aligned} \tag{129}$$

Finally, we use $\mathbb{E}[x^\top M x] = \text{tr}(M \mathbb{E}[x x^\top])$ and $\text{tr}(A) = \sum_{j=1}^n \lambda_j(A)$, where $\lambda_j(A)$ is the j -th eigenvalue of the matrix $A \in \mathbb{R}^{n \times n}$ to obtain

$$\begin{aligned}
\sum_{t=0}^{\tau} \mathbb{E} \left[\left\| \Delta A_i x(t) + \Delta B_i u(t) \right\|_M^2 \right] &= (1-\rho)^2 U_d^\top W_i(\tau) U_d + 2(1-\rho) U_d^\top m_i(x_0) + c_i(x(0)) \\
&\quad + \sigma_u^2 \rho^2 \text{tr}(W_i(\tau)) + \text{tr}(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau))
\end{aligned} \tag{130}$$

which proves the first statement of the Lemma. The second statement follows by the fact that $m_i(0) = 0$ and $c_i(0) = 0$. The last statement follows, plugging in $\rho = 1$ and using $c_i(x(0)) \geq 0$, $\forall x(0) \in \mathbb{R}^{n_x}$. \square

H Additional numerical examples

In this section we provide some additional numerical examples. First we provide a numerical evaluation of Example 3.1. Then we numerically study how the choice of the tradeoff parameter ρ_k influences the performance of the algorithm.

H.1 Numerical evaluation of Example 3.1

In this section, we consider a variant of the setup considered in Example 3.1. To be precise, consider $\tilde{\mathcal{S}} = \{(\tilde{A}_*, \tilde{B}_*), (\tilde{A}_1, \tilde{B}_*)\}$ with

$$\begin{aligned}
\tilde{A}_* &= \begin{bmatrix} 0 & 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.9 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.9 \end{bmatrix}, \quad \tilde{B}_* = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \\
\tilde{A}_1 &= \begin{bmatrix} 0 & 0.2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.9 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.9 \end{bmatrix}, \quad \tilde{B}_1 = \tilde{B}_*.
\end{aligned}$$

Further, we fix $\gamma_u = 1$ and consider the process noise $w(t) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_w^2 I_6)$, with $\sigma_w = 0.1$. Following the steps outlined in Example 3.1 we obtain $\lambda_{\max}(\tilde{W}_1(t)) = 5\lambda_{\text{mean}}(W_1(t))$. Since we

initialize Algorithm 2 with equal weights for both systems we immediately obtain $\delta_0 = \frac{1}{2}$. Thus, by the results in Section 4 we expect CE to perform well. To highlight the effectiveness of Algorithm 2 we compare $\mathbb{P}[\hat{\theta}_k = \theta_*]$ for three different setups, where we select $\eta = 0.1$:

- Algorithm 2 with $\rho_k \equiv 0$, i.e., greedily using CE with the current estimate
- Algorithm 2 with $\rho_k \equiv 1$, i.e., isotropic Gaussian excitations
- Algorithm 2 with $\rho_k \equiv 0$ and oracle knowledge, i.e., computing the optimal excitation with θ_*

We select $\tau = 10$ and run 20 Monte Carlo simulations. The mean and σ -bands of the likelihood of θ_* given the data are displayed in Figure 2. It is immediately apparent that oracle knowledge does not provide any advantage in this setup, since the current estimate $\hat{\theta}_k$ does not influence the solution of (26) due to the structure of the problem. Further, as expected, the experiment design algorithm significantly outperforms the random excitations by allocating all input energy to the first input.

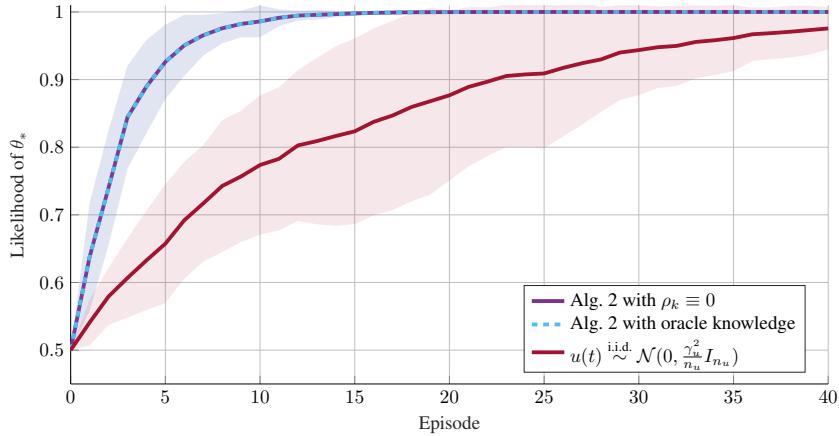


Figure 2: Mean and σ -band of the likelihood of θ_* given the data for different data collection strategies for the motivating example.

H.2 Identification for a random candidate set

In this section, we present the setup where the set \mathcal{S} is generated randomly. To be precise, we consider the case where \mathcal{S} consists of 20 additional systems, where the parameters are drawn from a Gaussian around the true parameters

$$A_* = \begin{bmatrix} 0 & 0.1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \quad B_* = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (131)$$

As before, we have $\Sigma_w = I_{n_x}$ and $\gamma_u = 1$. Furthermore, we consider different time-varying choices of the weighting parameter ρ_k and analyze how they influence the speed of identification. In particular we consider, $\rho_k = \frac{1}{1+\rho}$, $\rho_k = \frac{1}{(1+k)^2}$ and $\rho_k = e^{-k}$. Further, as for the numerical example in Section 5 we compare our algorithm to the active learning criterion proposed in [27] for the case of an infinite hypothesis class. Since the algorithm in [27] can not directly be compared to ours, since it has an exponentially increasing episode length, we use Algorithm 1 with the input defined by the objective in [27] with oracle knowledge. The resulting mean and σ -band of the likelihood of θ_* given the data over 100 simulations are displayed in Figure 3. It can be seen from the results that using CE in Algorithm 2 performs at least as good as the other time-varying selections for ρ_k . In particular, the performance is comparable to the performance of Algorithm 2 with oracle knowledge. Again, isotropic Gaussian excitations perform worse than all other excitation sequences. Notice that compared to the numerical example in Section 5 the gap is significantly smaller. This is due to the fact that the set \mathcal{S} consists of randomly generated systems. Hence, the set exhibits less structure, and uniform exploration is closer to optimal than in the previous example. This is in accordance with our

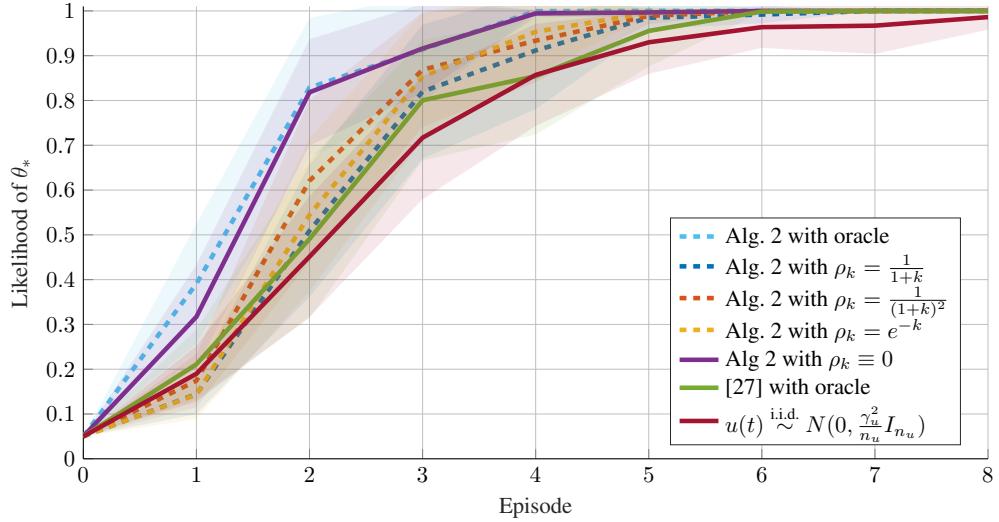


Figure 3: Mean and σ -band of the likelihood of θ_* given the data for different data collection strategies for the randomly generated set \mathcal{S} .

theoretical results and indicates that there exist non-trivial cases where experiment design provides only a small benefit over isotropic random excitations. Finally, the criterion presented in [27] again outperforms random excitations but does perform worse than our approach. This is the case because the criterion in [27] does not take into account the finiteness of the hypothesis class since it was not designed for this setup.