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

Nicolas Chatzikiriakos*
University of Stuttgart

Kevin Jamieson†
University of Washington

Andrea Iannelli*
University of Stuttgart

Abstract

We consider the problem of identifying an unknown linear dynamical system from 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 persistency of 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.

*Institute for Systems Theory and Automatic Control, nchatzi@ist.uni-stuttgart.de, iannelli@ist.uni-stuttgart.de

†Allen School of Computer Science & Engineering, jamieson@cs.washington.edu

Proceedings of the 29th International Conference on Artificial Intelligence and Statistics (AISTATS) 2026, Tangier, Morocco. PMLR: Volume 300. Copyright 2026 by the author(s).

1 INTRODUCTION

The problem of learning a model of an unknown dynamical system from data is important across domains such as reinforcement learning, 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. To tackle this problem, there exists a large literature on experiment design both in dynamical systems and learning theory. However, experiment design algorithms can be computationally demanding and a study of the potential benefit of experiment design algorithms is lacking in the literature of dynamical systems. This raises 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 systems[1] 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 (Sarkar and Rakhlin, 2019; Simchowitz et al., 2018) and certain classes of non-linear systems (Foster et al., 2020; Sattar et al., 2022; Sattar and Oymak, 2022). These results have been complemented by the sample complexity lower bounds for Gaussian excitations and linear systems presented by Jedra and Proutiere (2022); Tsiamis and Pappas (2021). For the case of a finite hypothesis class, Chatzikiriakos and Iannelli (2024); Muehlebach

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

et al. (2025) 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, e.g., Bombois et al. (2011); Goodwin and Payne (1977)). While these classical works consider the asymptotic case, recent works provided a finite sample perspective on the topic for linear (Wagenmaker and Jamieson, 2020; Wagenmaker et al., 2021; Chatzikiriakos et al., 2025) and certain classes of non-linear systems (Lee et al., 2024; Mania et al., 2022). While the underlying principles in the finite-sample analysis are always similar, there exists no modular framework for the analysis. Further, the benefit of experiment design over randomly exciting the systems has not been explicitly investigated, even for linear 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 a finite hypothesis class simplify experiment design compared to an infinite hypothesis class?

With respect to these key questions, our primary contributions can be summarized as follows:

1. We provide 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 analyze 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.
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 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 $U = \begin{bmatrix} u(t_0)^\top & \dots & u(t_1)^\top \end{bmatrix}^\top$, 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.

2 PRELIMINARIES

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

$$x(t+1) = A_* x(t) + B_* u(t) + w(t), \tag{1}$$

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. We assume $w(t) \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, \Sigma_w)$ with a known covariance matrix $\Sigma_w \in \mathbb{S}^{n_x}_{++}$. However, all the presented results and the proposed algorithm can be adapted to the case where Σ_w is unknown and only an upper bound $\sigma_w^2 I_{n_x} \succeq \Sigma_w$ or an informative prior is available. For notational simplicity, we assume $x(0) = 0$ unless stated otherwise. Further, we assume the learner has access to a finite set of systems, containing the true system, i.e.,

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

where $\mathcal{S}$ is known to the learner. 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 defined by the tuple $(\theta_*, \mathcal{S})$. The problem of identifying the true system from a finite hypothesis class naturally emerges in many settings, such as fault detection (Miljković, 2011) or validation of ecological and evolutionary models (Johnson and Omland, 2004). More generally, the set $\mathcal{S}$ can be viewed as a principled, model-driven way

to introduce knowledge of the real process into the identification problem.

The performance of the learner is evaluated in terms of the samples required to identify the true system with confidence δ. 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

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

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

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

which reveals that $\tilde{\varepsilon}_{\theta_i}(0,\tau)$ is a random quantity even if the excitation input is deterministic[2]. This is due to the process noise affecting the evolution of the system (1). Taking the expectation and using $w(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_w)$ it can be shown that

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

To simplify notation, we define the Toeplitz matrices

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

These matrices are tightly connected to the controllability of (A_*, B_*), and are commonly used in systems theory to map from the τ-step noise and input trajectories to the corresponding state trajectory. Using (6), computing the expectation in (5) and writing the statement in matrix form simplifies to

$$\begin{aligned}\mathbb{E}\left[\tilde{\varepsilon}_{\theta_i}(0,\tau)\right] = \mathbb{E}\Big[U^\top \big(& R^i_{\Sigma_w^{-1}}(\tau) + S_u(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_u(\tau) \\ &+ N^i_{\Sigma_w^{-1}}(\tau) S_u(\tau) + (N^i_{\Sigma_w^{-1}}(\tau) S_u(\tau))^\top\big) U\Big] \\ &+ \operatorname{tr}\Big(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\Big), \end{aligned} \tag{7}$$

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

$$Q^i_{\Sigma_w^{-1}}(\tau) \coloneqq \operatorname{diag}_\tau(\Delta A_i^\top \Sigma_w^{-1} \Delta A_i), \tag{8a}$$

$$R^i_{\Sigma_w^{-1}}(\tau) \coloneqq \operatorname{diag}_\tau(\Delta B_i^\top \Sigma_w^{-1} \Delta B_i) \tag{8b}$$

$$N^i_{\Sigma_w^{-1}}(\tau) \coloneqq \operatorname{diag}_\tau(\Delta B_i^\top \Sigma_w^{-1} \Delta A_i) \tag{8c}$$

Note that, leveraging the system-theoretic interpretation of the Toeplitz matrices (6), $U^\top R^i_{\Sigma_w^{-1}}(\tau) U = \sum_{t=0}^{\tau} \|\Delta B_i u(t)\|^2_{\Sigma_w^{-1}}$ 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, $U^\top S_u(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_u(\tau) U$ measures to what extent the difference between A_* and A_i can be seen in the weighted prediction error when applying U and $\operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right)$ measures the influence of the process noise on this quantity. Accounting for the cross term, for any $i \in [1, N]$, the matrices

$$\begin{aligned} W_i(\tau) \coloneqq& R^i_{\Sigma_w^{-1}}(\tau) + S_u(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_u(\tau) \\ &+ N^i_{\Sigma_w^{-1}}(\tau) S_u(\tau) + (N^i_{\Sigma_w^{-1}}(\tau) S_u(\tau))^\top, \end{aligned} \tag{9}$$

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 ENHANCE 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 make the following assumption on the excitation input, which is standard in the relevant literature (see, e.g., Wagenmaker and Jamieson (2020))

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., $\mathbb{E}\left[\sum_{t=0}^{T-1} \|u(t)\|^2\right] \le \gamma_u^2 T$.*

[2] A 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.

3.1 Sample complexity lower bounds

To formalize the class of algorithms considered 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 $(\theta_*, \mathcal{S})$ 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 sequence[3] $U \in \mathbb{R}^{n_u \bar{T}}$ and any δ-correct algorithm it holds that*

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

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

$$U^* \in \operatorname*{arg\,max}_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i\in[1,N]} U^\top W_i(\bar{T}) U + \operatorname{tr}\big(S_w(\bar{T})^\top Q^i_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T})\big), \tag{11}$$

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

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

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, an optimal solution exists.

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

$$U^* \in \operatorname*{arg\,max}_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i\in[1,N]} \mathbb{E}\left[\tilde{\varepsilon}_{\theta_i}(0,\bar{T})\right]. \tag{13}$$

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

Thus, Theorem 3.1 implies that the optimal excitation U^* maximizes $\mathbb{E}\left[\tilde{\varepsilon}_{\theta_i}(0,\bar{T})\right]$ uniformly over $\theta_i \in \mathcal{S}\setminus\{\theta_*\}$. Thus, U^* maximizes the distance between $\mathbb{E}\left[\tilde{\varepsilon}_{\theta_*}(0,\bar{T})\right] = 0$ and $\mathbb{E}\left[\tilde{\varepsilon}_{\theta_i}(0,\bar{T})\right]$, 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 $u(t) \overset{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$. Then for any δ-correct algorithm it holds that*

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

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*

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

and $\mathcal{S} = \{(A_*, B_*), (A_1, B_*)\}$. A straightforward calculation yields $W_1(t) = \operatorname{diag}_t\left([0.01]\right)$ 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

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

for some fixed $d > 0$. Through direct calculations, we obtain $\tilde{W}_1(t) = \operatorname{diag}_t\left(\begin{bmatrix} 0.01 & 0_{1\times d} \\ 0_{d\times 1} & 0_{d\times d} \end{bmatrix}\right)$, $\forall t \geq 1$ and hence $\lambda_{\text{mean}}(\tilde{W}_1(t)) = \frac{1}{d+1}\lambda_{\max}(\tilde{W}_1(t))$. 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 algorithms matching the lower bound are used. This can also 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 $(\theta_*, \mathcal{S})$ and $(\tilde{\theta}_*, \tilde{\mathcal{S}})$ is identical given the respective optimal oracle excitation is

used for both instances, as $\lambda_{\max}(W_1(t)) \equiv \lambda_{\max}(\tilde{W}_1(t))$. *Hence, in this sense the difficulty of the problems* $(\theta_*, \mathcal{S})$ *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, there exists no $c > 0$ for which this input satisfies

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

which is the PE condition required in the case of an infinite hypothesis class (cf. Tsiamis et al. (2023)). In fact, it is only necessary to excite the parts of the system that carry uncertainty. Thus, we introduce a weaker PE condition, which has been used as an assumption in a similar form by Muehlebach et al. (2025).

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* $(\theta_*, \mathcal{S})$ *if there exist coefficients* $c_u(\tau) > 0$, $c_w(\tau) > 0$ *such that for any* $x(0) \in \mathbb{R}^{n_x}$ *and* $\theta_i \in \mathcal{S} \setminus \{\theta_*\}$

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

Note that the PE coefficients c_u and c_w are problem-specific and depend on the block length τ. 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 as θ_* and hence is indistinguishable from θ_* given the data. In the following, we show that isotropic Gaussian excitations and the optimal oracle excitation U^* defined in (11) satisfy Definition 2, i.e., satisfy PE.

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

$$c_u^{\text{rand}}(\tau) = \min_{p \in \Delta_N} \lambda_{\text{mean}}\Big(\sum_{i=1}^N p_i W_i(\tau)\Big), \tag{19}$$

$$c_w(\tau) = \min_{i \in [1,N]} \frac{1}{\tau}\text{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_w(\tau)\right). \tag{20}$$

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

$$U^* \in \underset{U^\top U \leq \gamma_u^2 \tau}{\arg\max} \min_{i \in [1,N]} \sum_{t=0}^{\tau} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \tag{21}$$

is PE for block length τ *with*

$$c_u^{\text{opt}}(\tau) = \min_{p \in \Delta_N} \lambda_{\max}\Big(\sum_{i=1}^N p_i W_i(\tau)\Big), \tag{22}$$

$$c_w(\tau) = \min_{i \in [1,N]} \frac{1}{\tau}\text{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_w(\tau)\right). \tag{23}$$

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 $c_u^{\text{opt}}(\tau)$ coincide with the key problem parameters influencing the sample complexity lower bounds 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. *If* Σ_w *is unknown and instead an estimate* $\hat{\Sigma}_w$ *or upper bound* $\sigma_w^2 I_{n_x} \succeq \Sigma_w$ *is used to solve* (21), *the corresponding input sequence is PE, although with potentially suboptimal coefficients. This sub-optimality is due to the lack of absence of information which could be used when designing excitation inputs. 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

After establishing 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 Alg. 1 that holds for any PE input sequence. For estimation,

Algorithm 1 Sequential identification algorithm

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

1: **for** $k = 1, 2, \ldots$ **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 $\hat{\theta}$

the algorithm evaluates the sum of weighted empirical one-step prediction errors of θ_i up to time t defined as

$$\varepsilon_{\theta_i}(t) := \sum_{s=0}^{t-1} \|x(s+1) - A_i x(s) - B_i u(s)\|^2_{\Sigma_w^{-1}}. \quad (24)$$

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 Alg. 1 is equivalent to a log-likelihood test. While related works considering this setup (Chatzikiriakos and Iannelli, 2024; Muehlebach et al., 2025) 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 systematically 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 Alg. 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \leq \delta$ and with probability at least $1-\delta$ terminates no later than when k satisfies*

$$k c_w(\tau) + \sum_{j=1}^{k} c_{u_j}(\tau)\gamma_u^2 \geq c' \log\left(\frac{N}{\delta}\right), \quad (25)$$

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 input sequence in each epoch is PE, 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 Alg. 1 it is only necessary to show the excitation satisfies Definition 2. Given the PE coefficients $c_{u_j}(\tau)$, $c_w(\tau)$ the sample complexity upper bound follows immediately. Clearly, the level of PE of the excitation dictates the speed of identification, where larger PE coefficients result in faster identification. For the input sequences analyzed previously, the following result holds.

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

$$k\Big(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\text{mean}}\Big(\sum_{i=1}^{N} p_i W_i(\tau)\Big) + \min_{i \in [1,N]} \frac{1}{\tau}\text{tr}\left(S_w(\tau)^\top Q^i_{I_{n_x}} S_w(\tau)\right)\Big) \geq c' \log\left(\frac{N}{\delta}\right). \quad (26)$$

Further, if the optimal oracle excitation input U^* defined in (21) is applied the estimate $\hat{\theta}$ satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$, and with probability at least $1-\delta$ Alg. 1 terminates at the latest when k first satisfies

$$k\Big(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\max}\Big(\sum_{i=1}^{N} p_i W_i(\tau)\Big) + \min_{i \in [1,N]} \frac{1}{\tau}\text{tr}\left(S_w(\tau)^\top Q^i_{I_{n_x}} S_w(\tau)\right)\Big) \geq c'' \log\left(\frac{N}{\delta}\right), \quad (27)$$

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. Note that since PE coefficients (19),(22) coincide with the key parameters entering the sample complexity lower bounds, the sample complexity upper bounds depend on the same parameters as the lower bounds. Thus, the lower and upper bounds qualitatively match, suggesting that they capture the key quantities contributing to the sample complexity. The epoch length τ can be used as a degree of freedom to optimize the sample complexity bound. Comparing Corollaries 3.2 and 3.6, we see that, loosely speaking, whenever the matrices W_i are ill-conditioned, input design can provide significant benefits in terms of the sample complexity. On the other hand, when all eigenvalues are identical, isotropic Gaussian excitation is already optimal and cannot be improved upon (cf. Example 3.1). Notably, the matrices W_i are often ill-conditioned in settings where prior knowledge is available. Hence, in settings where the uncertainty only lies in certain directions, input design can yield large benefits, which confirms intuition. If, however, the uncertainty is unstructured, the benefits of experiment design are significantly smaller since this yields matrices W_i which are well-conditioned.

4 SEQUENTIAL INPUT DESIGN ALGORITHM

In the previous section, we established that the matrices $W_i(\tau)$ are the key quantities that determine the

Algorithm 2 Input design subroutine

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

1: Compute weights using exponential weighting $w_{k+1}(i) = \exp\left(-\eta\varepsilon_{\theta_i}((k-1)\tau)\right)$

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 (28)

4: Define excitation according to ρ_k: $u^*_k(t) = \sqrt{1-\rho_k}u^*_{\hat{\theta}_k}(t) + \sqrt{\rho_k}u_\eta(t)$, $u_\eta(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u}I_{n_u})$

return excitation sequence u^*_k

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 phenomenon is well known in the literature of experiment design and is usually tackled by using sequential algorithms to generate approximately optimal excitations. In line with these approaches, our proposed input design subroutine, which is presented as Alg. 2, can be called by Alg. 1 in each epoch and returns an excitation input sequence. To derive the excitation, Alg. 2 draws an estimate $\hat{\theta}_k$ using exponential weights and solves the optimization (21) 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 optimization problem

$$U^*_{\hat{\theta}_k}(x(k\tau)) \in \underset{U^\top U \le \gamma_u^2\tau}{\arg\max} \min_{\theta_i \neq \hat{\theta}_k} \mathbb{E}\Big[\sum_{t=k\tau}^{\tau(k+1)-1} \|\Delta\hat{A}_i x(t) + \Delta\hat{B}_i u(t)\|^2_{\Sigma_w^{-1}}\Big] \tag{28}$$

is solved to compute the excitation input. While this optimization problem uses knowledge of Σ_w, this is not required by our algorithm, and Alg. 2 still provides PE inputs even when only an upper bound or no knowledge of Σ_w is available (cf. Remark 3.2). We note that the optimization problem (28) is generally non-convex and, especially for $N \gg 0$ might be challenging to solve. Since the focus of this work is statistical, we defer finding scalable approximations, as done by Chatzikiriakos et al. (2025) for a similar setting, to future research. Due to the uncertainty in the estimate, we add random excitations to the optimal excitation input (28) 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 Alg. 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 Alg. 2 at some fixed iterate k and let $\mathbb{P}[\hat{\theta}_k \neq \theta_*] \le p_k \in [0,1]$, where $\hat{\theta}_k$ is the estimate drawn using exponential weights in round k. Then u^*_k satisfies PE with*

$$c^{\text{Alg}}_{u_k}(\tau) = (1-p_k)(1-\rho_k)c^{\text{opt}}_u(\tau) + \rho_k c^{\text{rand}}_u(\tau), \tag{29a}$$

$$c_w(\tau) = \min_{i\in[1,N]} \frac{1}{\tau}\text{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau)S_w(\tau)\right), \tag{29b}$$

where $c^{\text{opt}}_u(\tau)$ is the PE coefficient for optimal excitation as defined in (22) and $c^{\text{rand}}_u(\tau)$ 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 $u_1^\top(t) = \begin{bmatrix}1 & 0\end{bmatrix}$, $\forall t \ge 0$ and $u_2^\top(t) = \begin{bmatrix}-1 & 0\end{bmatrix}$, $\forall t \ge 0$ 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. *Consider Alg. 1 where the excitation input is derived by Alg. 2. Then, the estimate $\hat{\theta}$ of Alg. 1 satisfies $\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \le \delta$ and terminates no later than when k satisfies*

$$kc_w(\tau) + \sum_{j=1}^{k} c^{\text{Alg}}_{u_j}(\tau) \ge c' \log\left(\frac{N}{\delta}\right), \tag{30}$$

where $c^{\text{Alg}}_{u_j}(\tau)$, $c_w(\tau)$ are defined in (29) and c' is a constant influenced by the variance in $\tilde{\varepsilon}_{\theta_i}$.

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 shows that experiment design using CE provably improves the sample complexity of identification in these cases. Building on these arguments, the following result shows that as $\delta \to 0$ the proposed Algorithm exhibits a sample complexity upper bound that matches the lower bound up to constants.

Theorem 4.3. *Consider Alg. 1 where the excitation input is derived by Alg. 2, with η small enough. Let further $\rho_0 = c_0 \in (0,1)$, $\lim_{k\to\infty}\rho_k = c_\infty \in [0, c_0]$, $\rho_{k-1} \ge \rho_k \ge (1-\varepsilon)\rho_{k-1}$ where ε is a small enough non-negative constant. Then for any $\delta \in (0,1)$ there exists a k_δ, such that with probability $1-\delta$ Alg. 1*

terminates after at most k_δ epochs with an estimate $\hat{\theta}$ that satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and

$$\lim_{\delta \to 0} \frac{c' \log\left(\frac{N}{\delta}\right)}{k_\delta} \leq \gamma_u^2 (1 - c_\infty) \min_{p \in \Delta_N} \lambda_{\max}\Big(\sum_{i=1}^{N} p_i W_i(\tau)\Big) + \min_{i \in [1,N]} \frac{1}{\tau} \mathrm{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau) S_w(\tau)\right) \tag{31}$$

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

The full version of Theorem 4.3 and its proof are presented in Appendix F. This result shows that, Alg. 2 recovers the sample complexity lower bound up to constants as $\delta \to 0$ (cf. Theorem 3.1).

Remark 4.1. *Note that we need to resort to a blocking technique to show the result since the data is correlated. This has two consequences. First, compared to Theorem 3.1 the upper bound depends on $W_i(\tau)$ instead of $W_i(\bar{T})$, where $\tau \ll \bar{T}$. For stable systems $W_i(\tau)$ converges quickly as τ increases, i.e., we have $W_i(\tau) \approx W_i(\bar{T})$ even for small τ. However, if the system (1) is unstable, $W_i(\tau)$ can grow unbounded. Hence, while the upper bound remains valid for unstable systems, the gap to the lower bound is larger. Second, the sample complexity upper bound depends on the number of epochs k instead of the total number of samples $T = k\tau$. Since the epoch length τ is usually a small integer, this only introduces a small constant factor.*

Further, as shown by Proposition F.4, if ρ_k decreases slow enough the PE-coefficient of Alg. 2, $c^{\mathrm{Alg}}_{u_k}(\tau)$, increases as k increases. Hence, Alg. 2 outperforms random excitations after a finite number of iterations and, according to Proposition F.3, up to constants converges to the optimal PE coefficients. Note that while ε in Theorem 4.3 depends on oracle knowledge, any decrease that is slower than exponential can be used in practice, so that Theorem 4.3 guides the practical selection of ρ_k. In particular, by Theorem 4.3, even selecting a constant $\rho_k \in (0,1)$ (e.g., $\rho_k \equiv \frac{1}{2}$) will only impact the optimality by a constant as $\delta \to 0$. We numerically evaluate how different selections of ρ_k influence identification in Section 5. Interestingly, in our numerical studies, selecting $\rho_k = 0$, i.e., greedily using CE, yielded the best results.

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 optimal (see numerical evaluation in Appendix H). Deriving properties of $\mathcal{S}$ causing this is an interesting topic for future research.*

5 NUMERICAL EXPERIMENTS

In this section, we provide two numerical evaluations that showcase when active learning is beneficial and when its benefits are significantly smaller.[4] First, we consider an example where $|\mathcal{S}|$ is small and exhibits structure. In a second example, we consider the case where $\mathcal{S}$ is larger and generated randomly. While in the first case, active learning provides significant benefits in terms of identification speed, these benefits are significantly smaller in the second case.

Active learning can be high gain: Consider the set $\mathcal{S} = \{(A_*, B_*), (A_1, B_*), (A_2, B_*), (A_3, B_*)\}$ with $B_*^\top = \begin{bmatrix} 0_{2\times 1} & I_2 \end{bmatrix}$ and

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

The noise covariance and input power are given by $\Sigma_w = I_{n_x}$ and $\gamma_u = 1$, respectively. For the numerical evaluation, we consider Alg. 1 where the excitation input is computed using Alg. 2 with $\rho_k \equiv 0$ and $\eta = 0.01$, and compare it to the following data collection strategies used with Alg. 1: 1) Alg. 2 with oracle knowledge (in each epoch (28) is solved using θ_*), 2) the optimization criterion proposed by Wagenmaker and Jamieson (2020) with oracle knowledge, 3) $u(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{2} I_{n_u})$, and 4) optimal oracle excitation for the considered data length (fully offline). We conduct 100 Monte Carlo simulations with 5 epochs and an epoch length of 15 time steps. In Figure 1a we display the mean and $\frac{\sigma}{2}$-band of the likelihood of θ_* given the data over the epochs for the different data collection procedures. We investigate the likelihood since it is directly linked to the termination criterion of Alg. 1 and provides more insights than the termination times. It is easy to see that the optimal oracle excitation performs best. To see why this is the case, note that the uncertainty in $\mathcal{S}$ is structured and only affects particular directions in the state space. Hence, according to the theoretical insights presented in the previous sections, this setting admits large benefits of active learning. Consequently, isotropic Gaussian excitations perform significantly worse compared to all other input sequences. While the active learning criterion by Wagenmaker and Jamieson (2020), which does not account for the finite hypothesis class, outperforms isotropic Gaussian excitations, it falls short of the criterion presented in this work even when using oracle knowledge. This highlights the importance of leveraging knowledge of the finite hypothesis class in input design, when it is available.

Experiment design can be low gain: Now we consider

[4] All experiments are carried out in Python using a standard notebook. The code for the numerical example can be accessed at: https://github.com/col-tasas/2025-high-effort-low-gain

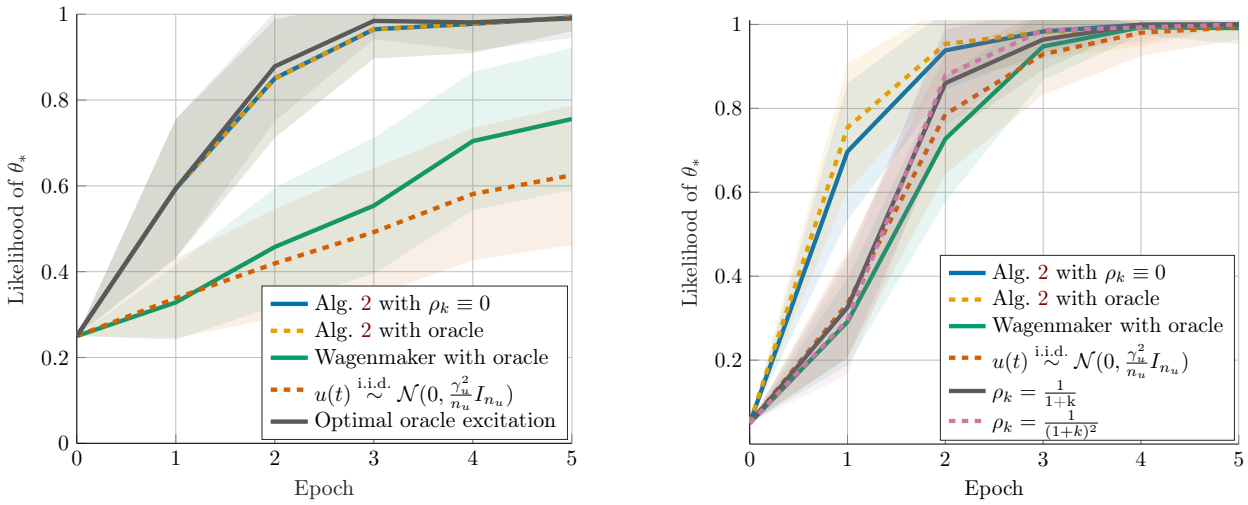

(a) Structured set $\mathcal{S}$.

(b) Randomly generated set $\mathcal{S}$.

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

a 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_*, B_*). The true system matrices (A_*, B_*) and all other problem parameters are the same as in the previous section, to single out the effect structure in $\mathcal{S}$ has on the results. We again compare Alg. 2 with $\rho_k \equiv 0$ to 1) Alg. 2 with oracle knowledge (in each epoch (28) is solved using θ_*), 2) the optimization criterion proposed by Wagenmaker and Jamieson (2020) with oracle knowledge, and 3) $u(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{1}{2} I_{n_u})$. Further, we consider different time-varying choices of the weighting parameter ρ_k, namely $\rho_k = \frac{1}{1+k}$ and $\rho_k = \frac{1}{(1+k)^2}$. The resulting mean and $\frac{\sigma}{2}$-band of the likelihood of θ_* given the data over 100 Monte Carlo simulations are displayed in Figure 1b. Again, isotropic Gaussian excitations perform worse than all excitation sequences generated by Alg. 2. Notice, however, that compared to the previous setting, 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. This is in accordance with our theoretical results and indicates that there exist non-trivial cases where experiment design provides only a small benefit over isotropic random excitations. The criterion presented in Wagenmaker and Jamieson (2020) performs similarly to random excitations and worse than our approach for all considered choices of ρ_k. Interestingly, in all our numerical evaluations, Alg. 2 with $\rho_k \equiv 0$ performs on par with the online oracle approach, even though it lacks the knowledge of the true system and relies on CE. In particular, this choice outperformed all other selections of ρ_k. This means that, even though Alg. 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 research.

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 a tailored 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 our findings, we propose a CE-based algorithm for experiment design and derive matching sample complexity upper bounds. Our analysis enables a quantitative assessment of the benefits provided by input design algorithms and identifies settings in which these benefits are limited, especially compared to the often considerable implementation complexity and computational effort of the underlying optimization problems. Numerical studies showcase that the proposed algorithm is highly competitive and achieves a performance close to the optimal oracle excitation. We see extensions to systems with partially observed states, where the decay rate of the Markov parameters impacts the optimal excitation, and investigations of the case where the true system does not lie in $\mathcal{S}$ as interesting extensions for future research.

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

Abbasi-Yadkori, Y., P., D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24.

Bombois, X., Gevers, M., Hildebrand, R., and Solari, G. (2011). Optimal experiment design for open and closed-loop system identification. *Commun. in Information and Systems*, 11(3):197–224.

Chatzikiriakos, N. and Iannelli, A. (2024). Sample complexity bounds for linear system identification from a finite set. *IEEE Control Systems Letters*, 8:2751–2756.

Chatzikiriakos, N., Song, B., Rank, P., and Iannelli, A. (2025). Hidden convexity in active learning: A convexified online input design for arx systems. In *IEEE 64th Conf. on Decision and Control (CDC)*, pages 63–68.

Chernoff, H. (1959). Sequential Design of Experiments. *The Annals of Mathematical Statistics*, 30(3):755 – 770.

Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. (2020). On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679.

Foster, D., Sarkar, T., and Rakhlin, A. (2020). Learning nonlinear dynamical systems from a single trajectory. In *Proceedings of the 2nd Conf. on Learning for Dynamics and Control*, volume 120, pages 851–861. PMLR.

Garivier, A., Ménard, P., and Stoltz, G. (2019). Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399.

Gerencsér, L. and Hjalmarsson, H. (2005). Adaptive input design in system identification. In *44th IEEE Conf. on Decision and Control*, pages 4988–4993.

Goodwin, G. and Payne, R. (1977). *Dynamic System Identification: Experiment Design and Data Analysis*. Mathematics in Science and Engineering. Academic Press.

Jedra, Y. and Proutiere, A. (2019). Sample complexity lower bounds for linear system identification. In *IEEE 58th Conf. on Decision and Control (CDC)*.

Jedra, Y. and Proutiere, A. (2022). Finite-time identification of linear systems: Fundamental limits and optimal algorithms. *IEEE Transactions on Automatic Control*.

Johnson, J. and Omland, K. (2004). Model selection in ecology and evolution. *Trends in Ecology & Evolution*, 19(2):101–108.

Kiefer, J. (1974). General equivalence theory for optimum designs (approximate theory). *The Annals of Statistics*, 2(5):849–879.

Kiefer, J. and Wolfowitz, J. (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366.

Lattimore, T. and Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.

Lee, B. D., Ziemann, I., Pappas, G. J., and Matni, N. (2024). Active learning for control-oriented identification of nonlinear systems. In *IEEE 63rd Conf. on Decision and Control (CDC)*.

Mania, H., Jordan, M. I., and Recht, B. (2022). Active learning for nonlinear system identification with guarantees. *J. of Machine Learning Research*, 23(32):1–30.

Miljković, D. (2011). Fault detection methods: A literature survey. In *2011 Proceedings of the 34th international convention MIPRO*, pages 750–755. IEEE.

Muehlebach, M., He, Z., and Jordan, M. I. (2025). The sample complexity of online reinforcement learning: A multi-model perspective. *arXiv:2501.15910*.

Rigollet, P. and Hütter, J.-C. (2023). High-dimensional statistics. *arXiv:2310.19244*.

Rojas, C. R., Welsh, J. S., Goodwin, G. C., and Feuer, A. (2007). Robust optimal experiment design for system identification. *Automatica*, 43(6):993–1008.

Sarkar, T. and Rakhlin, A. (2019). Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conf. on Machine Learning*, pages 5610–5618. PMLR.

Sattar, Y. and Oymak, S. (2022). Non-asymptotic and accurate learning of nonlinear dynamical systems. *Journal of Machine Learning Research*, 23(140):1–49.

Sattar, Y., Oymak, S., and Ozay, N. (2022). Finite sample identification of bilinear dynamical systems. In *IEEE 61st Conf. on Decision and Control (CDC)*, pages 6705–6711.

Simchowitz, M., Mania, H., Tu, S., Jordan, M. I., and Recht, B. (2018). Learning without mixing: Towards a sharp analysis of linear system identification. In *Conf. On Learning Theory*, pages 439–473. PMLR.

Tsiamis, A. and Pappas, G. J. (2021). Linear systems can be hard to learn. In *60th IEEE Conf. on Decision and Control (CDC)*. IEEE.

Tsiamis, A., Ziemann, I., Matni, N., and Pappas, G. J. (2023). Statistical learning theory for control: A finite-sample perspective. *IEEE Control Systems*, 43(6):67–97.

Wagenmaker, A. and Jamieson, K. (2020). Active learning for identification of linear dynamical systems. In *Conf. on Learning Theory*, pages 3487–3582. PMLR.

Wagenmaker, A. J., Simchowitz, M., and Jamieson, K. (2021). Task-optimal exploration in linear dynamical systems. In *38th International Conf. on Machine Learning*, pages 10641–10652. PMLR.

Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No, will be included in final submission]

2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [Yes]
   (b) Complete proofs of all theoretical results. [Yes]
   (c) Clear explanations of any assumptions. [Yes]

3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [No, will be included in final submission]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [Yes]
   (b) The license information of the assets, if applicable. [Not Applicable]
   (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
   (d) Information about consent from data providers/curators. Not Applicable]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [Not Applicable]
   (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
   (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Contents

1 INTRODUCTION **1**

2 PRELIMINARIES **2**

3 CAN ACTIVE LEARNING ENHANCE IDENTIFICATION? **3**

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 5

4 SEQUENTIAL INPUT DESIGN ALGORITHM **6**

5 NUMERICAL EXPERIMENTS **8**

6 CONCLUSION **9**

A Related Works **13**

B Proofs for sample complexity lower bounds **13**

B.1 Proof of Theorem 3.1 14

B.2 Proof of Corollary 3.2 15

C Existence of the optimal solution to (11) **15**

D Proving PE for different excitations **16**

D.1 Proof of Lemma 3.3 16

D.2 Proof of Lemma 3.4 16

D.3 Proof of Lemma 4.1 16

E Proofs for sample complexity upper bounds **17**

E.1 Proof of Theorem 3.5 18

E.2 Intermediate results used in the proof of Theorem 3.5 19

E.3 Sample complexity Upper bounds for isotropic Gaussian and oracle excitations 22

E.4 Sample complexity upper bound for Algorithm 2 23

F Asymptotic optimality of Algorithm 2 **23**

F.1 Intermediate Results for the Proof of Theorem F.1 24

G Auxiliary results **26**

H Numerical evaluation of Example 3.1 **28**

A Related Works

Finite Sample System Identification Novel results in high-dimensional statistics (Abbasi-Yadkori et al., 2011; Wainwright, 2019) 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, the first work by Dean et al. (2020) 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 (Simchowitz et al., 2018) and also unstable linear dynamical systems (Sarkar and Rakhlin, 2019) 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 (Sattar et al., 2022) generalized linear systems (Foster et al., 2020; Sattar and Oymak, 2022). 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, Jedra and Proutiere (2022) provide sample complexity lower bounds that hold independently of the used algorithm. Tsiamis and Pappas (2021) used similar tools 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 by Chatzikiriakos and Iannelli (2024); Muehlebach et al. (2025), who observed that the lack of stability of the true system does not 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 (Chernoff, 1959; Kiefer, 1974; Kiefer and Wolfowitz, 1960) 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 (cf. Bombois et al. (2011); Goodwin and Payne (1977) 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 (cf. Rojas et al. (2007) and the references therein) and sequential approaches that rely on a running estimate Gerencsér and Hjalmarsson (2005). While classical works only consider the asymptotic case, recently several works have considered the problem of experiment design for the identification of an unknown dynamical system under a finite sample perspective. In particular, Wagenmaker and Jamieson (2020); Wagenmaker et al. (2021) consider LTI systems and provide finite sample guarantees by leveraging sequential algorithms. Further, Chatzikiriakos et al. (2025) provide an exact convex reformulation of the design criterion presented by Wagenmaker and Jamieson (2020), 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 Lee et al. (2024); Mania et al. (2022) 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 by Wagenmaker and Jamieson (2020), 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 (Jedra and Proutiere, 2019, 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*

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

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

Proof. Define the data as $\mathcal{D}_{\bar{T}} := \{x(0), u(0), \ldots, u(\bar{T}-1), x(\bar{T})\}$ 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 presented by Jedra and Proutiere (2019), we use the generalized data processing inequality (Garivier et al., 2019, Lemma 1) to obtain the lower bound

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

where $\mathrm{kl}(x\|y)$ 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

$$\mathrm{kl}(\mathbb{P}_{\theta_*}(\mathcal{E}) \| \mathbb{P}_{\theta_i}(\mathcal{E})) \geq (2\delta - 1)\log\left(\frac{1-\delta}{\delta}\right) \geq \log\left(\frac{1}{2.4\delta}\right). \tag{33}$$

Further, we follow (Jedra and Proutiere, 2019, Section IV.A) to obtain

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

which concludes the proof. □

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

$$\mathbb{E}\left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \geq 2\log\left(\frac{1}{2.4\delta}\right) \quad \forall i \in [1, N], \tag{34}$$

where $x(t)$ is generated by (1). Since (34) needs to hold for all $i \in [1, N]$ is follows immediately that it suffices to consider

$$\min_{i\in[1,N]} \mathbb{E}\left[\sum_{s=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \geq 2\log\left(\frac{1}{2.4\delta}\right). \tag{35}$$

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 (42) over all admissible input sequences. This yields

$$\max_{U^\top U \leq \gamma_u^2 T} \min_{i\in[1,N]} \mathbb{E}\left[\sum_{s=0}^{t-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \geq 2\log\left(\frac{1}{2.4\delta}\right). \tag{36}$$

Now, using Lemma G.1 with $\rho = 0$ and $M = \Sigma_w^{-1}$ we can equivalently rewrite the l.h.s. of (36) as

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

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

$$\begin{aligned}
&\max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1,N]} U^\top W_i(\bar{T}) U + \operatorname{tr}\left(S_w(\bar{T})^\top Q^i_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right) && (38)\\
&\leq \max_{U^\top U \leq \gamma_u^2 \bar{T}} \min_{i \in [1,N]} U^\top W_i(\bar{T}) U + \max_{j \in [1,N]} \operatorname{tr}\left(S_w(\bar{T})^\top Q^j_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right) && (39)\\
&= \max_{U^\top U \leq \gamma_u^2 T} \min_{p \in \Delta_N} U^\top \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right) U + \max_{j \in [1,N]} \operatorname{tr}\left(S_w(\bar{T})^\top Q^j_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right) && (40)\\
&= \min_{p \in \Delta_N} \max_{U^\top U \leq \gamma_u^2 T} U^\top \left(\sum_{i=1}^N p_i W_i(\bar{T}) \right) U + \max_{j \in [1,N]} \operatorname{tr}\left(S_w(\bar{T})^\top Q^j_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right). && (41)
\end{aligned}$$

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

$$\mathbb{E}\left[\sum_{t=0}^{\bar{T}-1} \|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}} \right] \geq 2 \log\left(\frac{1}{2.4\delta} \right) \quad \forall i \in [1,N], \tag{42}$$

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

$$\frac{\gamma_u^2}{n_u} \operatorname{tr}\left(W_i(\bar{T}) \right) + \operatorname{tr}\left(S_w(\bar{T})^\top Q^i_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right) \geq 2 \log\left(\frac{1}{2.4\delta} \right) \quad \forall i \in [1,N]. \tag{43}$$

After noting that $\frac{1}{n_u \bar{T}} \operatorname{tr}\left(W_i(\bar{T}) \right) = \lambda_{\text{mean}}(W_i(\bar{T}))$ the result follows immediately after realizing that it suffices to consider the minimum of the l.h.s. of (43).

C Existence of the optimal solution to (11)

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

$$W_i \coloneqq R^i_{\Sigma_w^{-1}}(\bar{T}) + S_u(\bar{T})^\top Q^i_{\Sigma_w^{-1}}(\bar{T}) S_u(\bar{T}) + 2 N^i_{\Sigma_w^{-1}}(\bar{T}) S_u(\bar{T}), \tag{44}$$

$$c_i \coloneqq \operatorname{tr}\left(S_w(\bar{T})^\top Q^i_{\Sigma_w^{-1}}(\bar{T}) S_w(\bar{T}) \right), \tag{45}$$

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

$$\begin{aligned}
\min_{U,\xi} \quad & -\xi && (46a)\\
\text{s.t.} \quad & U^\top U - \gamma_u^2 \bar{T} \leq 0 && (46b)\\
& \xi - U^\top W_i U - c_i \leq 0 \qquad \forall i \in [1,N], && (46c)\\
& -\xi \leq 0 && (46d)
\end{aligned}$$

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 (46b). The feasible region of ξ is defined by the constraints (46c). Clearly, for a given $\bar{T}$ under (46b) ξ is upper bounded by a finite, non-negative value. Combining this with (46d) the feasible region of ξ is bounded and closed, from which we can conclude that the optimal solution to (46) 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

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

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

$$\min_{i\in[1,N]}\frac{\gamma_u^2}{n_u\tau}\mathrm{tr}\left(W_i(\tau)\right) + \frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}}(\tau)S_w(\tau)\right) \tag{47}$$

$$\geq \min_{i\in[1,N]}\frac{\gamma_u^2}{n_u\tau}\mathrm{tr}\left(W_i(\tau)\right) + \min_{j\in[1,N]}\frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^j_{\Sigma_w^{-1}}(\tau)S_w(\tau)\right) \tag{48}$$

$$= \min_{p\in\Delta_N}\frac{\gamma_u^2}{n_u\tau}\mathrm{tr}\left(\sum_{i=1}^{N}p_iW_i(\tau)\right) + \min_{j\in[1,N]}\frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^j_{\Sigma_w^{-1}}(\tau)S_w(\tau)\right) \tag{49}$$

$$= \min_{p\in\Delta_N}\lambda_{\text{mean}}\left(\sum_{i=1}^{N}p_iW_i(\tau)\right)\gamma_u^2 + \min_{j\in[1,N]}\frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^j_{\Sigma_w^{-1}}(\tau)S_w(\tau)\right). \tag{50}$$

Comparing terms with (18) yields the result. □

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

$$\begin{aligned}\sum_{t=0}^{\tau-1}\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_M\right] &= U^\top W_i(\tau)U + 2U^\top m_i\left(x(0)\right)\\ &\quad + c_i\left(x(0)\right) + \mathrm{tr}\left(S_w(\tau)^\top Q^i_M(\tau)S_w(\tau)\right)\end{aligned} \tag{51}$$

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

$$\max_{U^\top U\leq\gamma_u^2\tau}\min_{i\in[1,N]} U^\top W_i(\tau)U + 2U^\top m_i(x_0) + c_i\left(x(0)\right) + \bar{c}_i \tag{52}$$

$$\geq \max_{U^\top U\leq\gamma_u^2\tau}\min_{i\in[1,N]} U^\top W_i(\tau)U + \bar{c}_i \tag{53}$$

$$\geq \max_{U^\top U\leq\gamma_u^2\tau}\min_{i\in[1,N]} U^\top W_i(\tau)U + \min_{j\in[1,N]}\bar{c}_j, \tag{54}$$

where the first inequality holds since $c_i\left(x(0)\right)$ 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. □

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

$$\sum_{t=0}^{\tau-1}\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_k^*(t)\|^2_{\Sigma_w^{-1}}] = \mathbb{P}[\hat{\theta}_0=\theta_*]\sum_{t=0}^{\tau-1}\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_{\theta_*}^*(t)\|^2_{\Sigma_w^{-1}}] + \mathbb{P}[\hat{\theta}_0\neq\theta_*]\sum_{t=0}^{\tau-1}\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_{\bar{\theta}}^*(t)\|^2_{\Sigma_w^{-1}}]. \tag{55}$$

The first terms in (55) 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

$$\sum_{t=0}^{\tau-1}\mathbb{E}\left[\|\ \Delta A_i x(t)+\Delta B_i u_{\theta_*}^*(t)\|^2_{\Sigma_w^{-1}}\right] \geq \operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + (1-\rho_0)(U^*_{\theta_*})^\top W_i(\tau) U^*_{\theta_*} + \frac{\rho_0\gamma_u^2}{n_u}\operatorname{tr}\left(W_i(\tau)\right). \tag{56}$$

To tackle the second term in (55) we follow the same steps to obtain

$$\sum_{t=0}^{\tau-1}\mathbb{E}\left[\|\Delta A_i x(t)+\Delta B_i u_{\bar{\theta}}^*(t)\|^2_{\Sigma_w^{-1}}\right] \geq \operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + (1-\rho_0)(U^*_{\bar{\theta}})^\top W_i(\tau) U^*_{\bar{\theta}} + \frac{\rho_0\gamma_u^2}{n_u}\operatorname{tr}\left(W_i(\tau)\right). \tag{57}$$

Plugging (56) and (57) into (55) and collecting terms we obtain

$$\begin{aligned}\sum_{t=1}^{\tau-1}&\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_k^*(t)\|^2_{\Sigma_w^{-1}}]\\ &\geq \operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + \frac{\rho_0\gamma_u^2}{n_u}\operatorname{tr}\left(W_i(\tau)\right)\\ &\quad + \mathbb{P}[\hat{\theta}_0=\theta_*](1-\rho_0)(U^*_{\theta_*})^\top W_i(\tau)U^*_{\theta_*} + \mathbb{P}[\hat{\theta}_0\neq\theta_*](1-\rho_0)(U^*_{\bar{\theta}})^\top W_i(\tau)U^*_{\bar{\theta}}.\end{aligned} \tag{58}$$

The last term in (58) is non-negative. Hence, we obtain that

$$\begin{aligned}\sum_{t=1}^{\tau-1}&\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_k^*(t)\|^2_{\Sigma_w^{-1}}]\\ &\geq \operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + \frac{\rho_0\gamma_u^2}{n_u}\operatorname{tr}\left(W_i(\tau)\right) + (1-p_0)(1-\rho_0)(U^*_{\theta_*})^\top W_i(\tau)U^*_{\theta_*},\end{aligned} \tag{59}$$

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

$$\begin{aligned}&\frac{1}{\tau}\sum_{t=1}^{\tau-1}\mathbb{E}[\|\Delta A_i x(t)+\Delta B_i u_k^*(t)\|^2_{\Sigma_w^{-1}}]\\ &\geq \frac{1}{\tau}\operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + \frac{\rho_0\gamma_u^2}{\tau n_u}\operatorname{tr}\left(W_i(\tau)\right) + (1-p_0)(1-\rho_0)\frac{1}{\tau}(U^*_{\theta_*})^\top W_i(\tau)U^*_{\theta_*}\\ &\geq \min_{i\in[1,N]}\frac{1}{\tau}\operatorname{tr}\left(S_w(\tau)^\top Q^i_{\Sigma_w^{-1}} S_w(\tau)\right) + \frac{\rho_0\gamma_u^2}{\tau n_u}\operatorname{tr}\left(W_i(\tau)\right) + (1-p_0)(1-\rho_0)\frac{1}{\tau}(U^*_{\theta_*})^\top W_i(\tau)U^*_{\theta_*}.\end{aligned} \tag{60}$$

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 $\mathbb{P}\left[\hat{\theta}_T \neq \theta_*\right] \leq \delta$ and with probability at least $1-\delta$ terminates at the latest for the first k satisfying*

$$\tau \sum_{j=1}^{k} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau) \geq 8\left(2a + \frac{1}{2\eta}\right)\log\left(\frac{N}{\delta}\right), \tag{61}$$

where

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^\top \Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)} \tag{62}$$

and $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (81).

Proof. This proof consists of two parts. First we show that if the algorithm terminates it holds that $\mathbb{P}\left[\hat{\theta}_T \neq \theta_*\right] \leq \delta$, 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 $\mathbb{P}_{\theta_i}$ and $\mathbb{E}_{\theta_i}$ 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

$$L_{\theta_i,\theta_j}(t) = \frac{\mathbb{P}_{\theta_i}[\mathcal{D}_t]}{\mathbb{P}_{\theta_j}[\mathcal{D}_t]} = \frac{\exp\left(-\frac{1}{2}\varepsilon_{\theta_i}(t)\right)}{\exp\left(-\frac{1}{2}\varepsilon_{\theta_j}(t)\right)}. \tag{63}$$

It can easily be shown that the likelihood-ratio is a martingale sequence (see, e.g., (Wainwright, 2019, Example 2.18)). By the termination rule the algorithm terminates under the event

$$\mathcal{E} = \left\{\exists\hat{\theta} : \log(L_{\hat{\theta},\theta_i}(t)) > \log\left(\frac{N}{\delta}\right), \quad \forall\theta_i \in \mathcal{S}\setminus\{\hat{\theta}\}\right\} \tag{64}$$

$$= \left\{\exists\hat{\theta} : L_{\hat{\theta},\theta_i}(t) > \frac{N}{\delta}, \quad \forall\theta_i \in \mathcal{S}\setminus\{\hat{\theta}\}\right\}. \tag{65}$$

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

$$\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \leq \mathbb{P}_{\theta_0}\left[\bigcup_{i=1}^{N} \exists t \in \mathbb{Z}_+ : L_{\theta_i,\theta_0}(t) \geq \frac{N}{\delta}\right]. \tag{66}$$

Thus,

$$\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \leq \mathbb{P}_{\theta_0}\left[\bigcup_{i=1}^{N} \exists t \in \mathbb{Z}_+ : L_{\theta_i,\theta_0}(t) \geq \frac{N}{\delta}\right] \tag{67}$$

$$\leq \sum_{i=1}^{N} \mathbb{P}_{\theta_0}\left[\exists t \in \mathbb{Z}_+ : L_{\theta_i,\theta_0}(t) \geq \frac{N}{\delta}\right] \leq \sum_{i=1}^{N}\frac{\delta}{N} = \delta, \tag{68}$$

where the second inequality uses a union bound, and the last inequality follows from (Lattimore and Szepesvári, 2020, 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

$$S_i(k) = \exp\left(-\eta\left(\varepsilon_{\theta_i}(t_k) - \varepsilon_{\theta_*}(t_k) - \frac{1}{4}\sum_{j=0}^{k-1}\sum_{t=t_j}^{t_{j+1}-1}\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}|\mathcal{F}_j\right]\right)\right) \tag{69}$$

with $S_i(0) = 1$ is a super-martingale. Hence, again using the maximal inequality (Lattimore and Szepesvári, 2020, Theorem 3.9) we obtain

$$\mathbb{P}\left[\exists k : S_i(k) \geq \frac{N}{\delta}\right] \leq \frac{\delta}{N} \tag{70}$$

Thus using union bound arguments we obtain

$$\mathbb{P}\left[\bigcup_{i=1}^{N} \exists k : S_i(k) \geq \frac{N}{\delta}\right] \leq \sum_{i=1}^{N} \mathbb{P}\left[\exists k : S_i(k) \geq \frac{N}{\delta}\right] \leq \delta \tag{71}$$

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

$$\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k) \geq \frac{1}{\eta}\log\left(\frac{N}{\delta}\right) - \frac{\tau}{4}\sum_{j=0}^{k-1} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau). \tag{72}$$

Thus by (71) it holds with probability at least $1-\delta$ that

$$\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k) \leq \frac{1}{\eta}\log\left(\frac{N}{\delta}\right) - \frac{\tau}{4}\sum_{j=0}^{k-1} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau). \tag{73}$$

Thus, as soon as

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

we obtain

$$\begin{aligned}
\log\left(L_{\theta_i,\theta_*}(k\tau)\right) &= \frac{1}{2}\left(\varepsilon_{\theta_*}(t_k) - \varepsilon_{\theta_i}(t_k)\right) && (75)\\
&\leq \frac{1}{2\eta}\log\left(\frac{N}{\delta}\right) - \frac{\tau}{8}\sum_{j=0}^{k-1} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau) && (76)\\
&\leq \frac{1}{2\eta}\log\left(\frac{N}{\delta}\right) - \left(\frac{1}{2\eta}+2\right)\log\left(\frac{N}{\delta}\right) && (77)\\
&= 2\log\left(\frac{N}{\delta}\right) && (78)
\end{aligned}$$

Thus, when (74) is satisfied Algorithm 1 terminates. This happens at the latest when (61) is satisfied and yields $\hat{\theta} = \theta_*$ with probability at least $1-\delta$. □

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 $u(t) = u_{\mathrm{d}}(t) + u_{\mathrm{r}}(t)$, where $u_d(t)$ is a deterministic sequence and $u_{\mathrm{r}}(t) \overset{i.i.d.}{\sim} \mathcal{N}(0, \Sigma_u)$. 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., $\mathcal{F}_k = \{x(0), u(0), \ldots, u(t_k - 1), x(t_k)\}$. Then the sequence*

$$S_i(k) = \exp\left(-\eta\left(\varepsilon_{\theta_i}(t_k) - \varepsilon_{\theta_*}(t_k) - \frac{1}{4}\sum_{j=0}^{k-1}\sum_{t=t_j}^{t_{j+1}-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}} | \mathcal{F}_j\right]\right)\right) \tag{79}$$

with $S_i(0) = 1$ is a supermartingale[5] for any $\eta > 0$ satisfying

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)} \tag{80}$$

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

with

$$\Sigma_{\Delta_i}(\tau) \coloneqq \Delta A_i \left(\sum_{s=0}^{\tau-1} A_*^{\tau-s} B \Sigma_u B^\top (A^{\tau-s})^\top + A_*^{\tau-s} \Sigma_w (A_*^{\tau-s})^\top \right) \Delta A_i^\top + \Delta B_i \Sigma_u \Delta B_i^\top. \tag{81}$$

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

$$\varepsilon_{\theta_i}(t_0, t_1) \coloneqq \sum_{t=t_0}^{t_1-1} \|x(t+1) - A_i x(t) - B_i u(t)\|^2_{\Sigma_w^{-1}}. \tag{82}$$

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

$$\begin{aligned} \mathbb{E}\left[S_i(k+1)|\mathcal{F}_k\right] =& S_i(k)\mathbb{E}\Bigg[\exp\Bigg(-\eta\Bigg(\varepsilon_{\theta_i}(t_k, t_{k+1}) \\ & - \varepsilon_{\theta_*}(t_k, t_{k+1}) - \frac{1}{4}\left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right]\right)\Bigg)\Bigg)\Bigg|\mathcal{F}_k\Bigg]. \end{aligned} \tag{83}$$

Thus, to show the result requires showing that

$$\mathbb{E}\left[e^{-\eta\left(\varepsilon_{\theta_i}(t_k,t_{k+1}) - \varepsilon_{\theta_*}(t_k,t_{k+1}) - \frac{\tau}{4}\left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right]\right)\right)}\Bigg|\mathcal{F}_k\right] \leq 1. \tag{84}$$

To this end, we define

$$\xi_{\theta_i}(t) \coloneqq \|x(t+1) - A_i x(t) - B_i u(t)\|^2_{\Sigma_w^{-1}}. \tag{85}$$

Note that

$$\begin{aligned} \mathbb{E}\left[e^{-\eta\left(\varepsilon_{\theta_i}(t_k,t_{k+1}) - \varepsilon_{\theta_*}(t_k,t_{k+1})\right)}\Big|\mathcal{F}_k\right] &= \mathbb{E}\left[e^{-\eta\sum_{t=t_k}^{t_{k+1}-1} \xi_{\theta_i}(t) - \xi_{\theta_*}(t)}\Big|\mathcal{F}_k\right] & (86)\\ &\leq \prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E}\left[e^{-\eta\tau(\xi_{\theta_i}(t) - \xi_{\theta_*}(t))}\Big|\mathcal{F}_k\right]\right)^{\frac{1}{\tau}} & (87)\\ &\leq \prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E}\left[e^{-\frac{\eta}{2}\tau\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}}\Big|\mathcal{F}_k\right]\right)^{\frac{1}{\tau}}, & (88) \end{aligned}$$

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 (81) and $\mu_{\Delta_i}(t)$ depends on $x(t_k)$, the input sequence $u_{\mathrm{d}}(t)$, and the system matrices. Hence, conditioned on $\mathcal{F}_k$, by Proposition E.4 we have

$$\begin{aligned} \frac{1}{2}\Big(&\|\Delta A_i x(t_k + t) + \Delta B_i u(t_k + t)\|^2_{\Sigma_w^{-1}} \\ &- \mathbb{E}\Big[\|\Delta A_i x(t_k + t) + \Delta B_i u(t_k + t)\|^2_{\Sigma_w^{-1}}\Big]\Big) \sim \mathrm{subExp}\left(\nu(t)\right), \forall t \in [0, \tau - 1] \end{aligned} \tag{89}$$

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

$$\begin{aligned} &\prod_{t=t_k}^{t_{k+1}-1} \left(\mathbb{E}\left[e^{-\frac{\eta}{2}\tau\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}}\Big|\mathcal{F}_k\right]\right)^{\frac{1}{\tau}} & (90)\\ &\leq \left(\prod_{t=t_k}^{t_{k+1}-1} e^{\frac{\eta^2}{2}\tau^2\nu(t)^2} e^{-\frac{\eta}{2}\tau\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right]}\right)^{\frac{1}{\tau}} & (91)\\ &= \left(e^{-\frac{\eta}{2}\tau\left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] - \tau\eta\nu^2(t)\right)}\right)^{\frac{1}{\tau}} & (92)\\ &= e^{-\frac{\eta}{2}\left(\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] - \tau\eta\nu^2(t)\right)}. & (93) \end{aligned}$$

Imposing

$$\tau\eta\nu^2(t) \leq \frac{1}{2}\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \tag{94}$$

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

$$\mathbb{E}\left[e^{-\eta\left(\varepsilon_{\theta_i}(t_k,t_{k+1})-\varepsilon_{\theta_*}(t_k,t_{k+1})\right)}\Big|\mathcal{F}_k\right] \leq \exp(-\frac{\eta}{4}\sum_{t=t_k}^{t_{k+1}-1}\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}|\mathcal{F}_k\right]) \leq 1, \tag{95}$$

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

$$\begin{aligned}
&\mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}|\mathcal{F}_k\right] && (96)\\
&= \mathbb{E}\left[(\Delta A_i x(t) + \Delta B_i u(t))^\top \Sigma_w^{-1}(\Delta A_i x(t) + \Delta B_i u(t))|\mathcal{F}_k\right] && (97)\\
&= \operatorname{tr}\left(\Sigma_w^{-1}\mathbb{E}\left[(\Delta A_i x(t) + \Delta B_i u(t))(\Delta A_i x(t) + \Delta B_i u(t))^\top|\mathcal{F}_k\right]\right) && (98)\\
&= \operatorname{tr}\left(\Sigma_w^{-1}(\Sigma_{\Delta_i}(t) + \mu_{\Delta_i}(t)\mu_{\Delta_i}(t)^\top)\right) && (99)\\
&= \operatorname{tr}\left((\Sigma_{\Delta_i}(t)^{\frac{1}{2}})^\top\Sigma_w^{-1}\Sigma_{\Delta_i}(t)^{\frac{1}{2}} + \mu_{\Delta_i}(t)^\top\Sigma_w^{-1}\mu_{\Delta_i}(t)\right) && (100)\\
&\geq \lambda_{\max}\left((\Sigma_{\Delta_i}(t)^{\frac{1}{2}})^\top\Sigma_w^{-1}\Sigma_{\Delta_i}(t)^{\frac{1}{2}}\right) = \frac{\nu(t)}{16}, && (101)
\end{aligned}$$

hence condition (94) is implied by

$$\tau\eta\nu^2(t) \leq \frac{\nu(t)}{32} \overset{\nu(t)>0}{\iff} \tau\eta\nu(t) \leq \frac{1}{32}\forall t \in [0,\tau] \tag{102}$$

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

$$\tau\eta \leq \frac{1}{32\nu(\tau)} = \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^\top\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)}. \tag{103}$$

□

Lemma E.3. *Define*

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

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

$$\mathbb{E}\left[e^{-\eta(\xi_{\theta_i}(t)-\xi_{\theta_*}(t))}\Big|\mathcal{F}_k\right] \leq \mathbb{E}\left[\exp\left(-\frac{\eta}{2}\left\|\Delta A_i x(t) - \Delta B_i u(t)\right\|^2_{\Sigma_w^{-1}}\right)\Big|\mathcal{F}_k\right] \quad \forall k = 1,2,\ldots,t. \tag{105}$$

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

$$\begin{aligned}
\xi_{\theta_i}(t) &= \|\Delta A_i x(t) - \Delta B_i u(t) + w(t)\|^2_{\Sigma_w^{-1}} && (106)\\
&= \|\Delta A_i x(t) - \Delta B_i u(t)\|^2_{\Sigma_w^{-1}} + 2w(t)^\top\Sigma_w^{-1}(A_i x(t) - \Delta B_i u(t)) + w(t)^\top\Sigma_w^{-1}w(t) && (107)
\end{aligned}$$

and hence

$$\xi_{\theta_i}(t) - \xi_{\theta_*}(t) = \|\Delta A_i x(t) - \Delta B_i u(t)\|^2_{\Sigma_w^{-1}} + 2w(t)^\top\Sigma_w^{-1}(A_i x(t) - \Delta B_i u(t)). \tag{108}$$

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 $w(t)^\top\Sigma_w^{-\frac{1}{2}} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{n_x})$. Recall that the definition of the moment generating function for a random variable $w \sim \mathcal{N}(0, I_{n_x})$ is given by

$$M_w(\lambda) = \mathbb{E}\left[e^{w^\top\lambda}\right] = \exp\left(\frac{1}{2}\lambda^\top\lambda\right). \tag{109}$$

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

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

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 $\frac{1}{2}\left(\|z\|^2_M - \mathbb{E}\left[\|z\|^2_M\right]\right)$ is sub-exponential[6] with parameter ν, where*

$$\nu = 16\lambda_{\max}\left(\Sigma_z^{\frac{1}{2}^\top} M \Sigma_z^{\frac{1}{2}}\right). \tag{110a}$$

Proof. Define $\zeta = \Sigma_z^{-\frac{1}{2}^\top}(z - \mu_z)$, yielding $\|z\|_M = \|\zeta + \Sigma_z^{-\frac{1}{2}^\top}\mu_z\|_{\Sigma_z^{\frac{1}{2}^\top} M \Sigma_z^{\frac{1}{2}}}$ 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 (Wainwright, 2019, Theorem 2.26)

$$\|\zeta + \mu_z\|_M \sim \mathrm{subG}\left(\lambda_{\max}\left(\Sigma_z^{\frac{1}{2}^\top} M \Sigma_z^{\frac{1}{2}}\right)\right). \tag{111}$$

Applying (Rigollet and Hütter, 2023, Lemma 1.12) we obtain the result. □

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 $u(t) \overset{i.i.d.}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$ then Algorithm 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \leq \delta$ and with probability at least $1 - \delta$ terminates no later than when k satisfies*

$$k\left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\mathrm{mean}}\left(\sum_{i=1}^N p_i W_i(\tau)\right) + \min_{i \in [1,N]} \frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^i_{I_{n_x}} S_w(\tau)\right)\right) \tag{112}$$

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

Further, if the optimal oracle excitation input U^* defined in (21) is applied the estimate $\hat{\theta}$ satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$, and with probability $1 - \delta$ Algorithm 1 terminates no later than when k satisfies

$$k\left(\gamma_u^2 \min_{p \in \Delta_N} \lambda_{\max}\left(\sum_{i=1}^N p_i W_i(\tau)\right) + \min_{i \in [1,N]} \frac{1}{\tau}\mathrm{tr}\left(S_w(\tau)^\top Q^i_{I_{n_x}} S_w(\tau)\right)\right) \tag{114}$$

$$\geq 8\left(\frac{2}{\tau} + \frac{1}{2\eta\tau}\right)\log\left(\frac{N}{\delta}\right). \tag{115}$$

For both results we have

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

where $\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)$ is defined in (81).

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

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

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 $\mathbb{P}\left[\hat{\theta} \neq \theta_*\right] \leq \delta$ and with probability at least $1-\delta$ terminates no later than when k satisfies*

$$\sum_{j=1}^{k} c_{u_j}^{\mathrm{Alg}}(\tau) + c_w(\tau) \geq 8\left(\frac{2}{\tau} + \frac{1}{2\eta\tau}\right)\log\left(\frac{N}{\delta}\right), \tag{117}$$

where

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)}, \tag{118}$$

with $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ as defined in (81) and $c_{u_j}^{\mathrm{Alg}}(\tau)$ and $c_w(\tau)$ as defined in (29).

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

F Asymptotic optimality of Algorithm 2

The theorem below is the full version of Theorem 4.3.

Theorem F.1. *Consider Algorithm 1 where the excitation input is derived by Algorithm 2, with*

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^{\top}\Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)} \tag{119}$$

where $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (81). Let $\rho_0 = c_0 \in (0,1)$, $\lim_{k\to\infty}\rho_k = c_\infty \in [0, c_0]$ and $\rho_{k-1} \geq \rho_k \geq (1-\varepsilon)\rho_{k-1}$, where

$$\varepsilon < \frac{\left(1 - e^{-\frac{\eta\tau}{4}c_w(\tau)}\right)(1-\rho_0)}{\rho_0 e^{-\frac{\eta\tau}{4}c_w(\tau)}}. \tag{120}$$

Then for any $\delta \in (0,1)$ there exists a k_δ, such that with probability at least $1-\delta$ Algorithm 1 terminates after at most k_δ epochs with an estimate $\hat{\theta}$ that satisfies $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and

$$\lim_{\delta\to 0}\frac{8\left(\frac{1}{2\eta\tau} + \frac{1}{\tau}\right)\log\left(\frac{N}{\delta}\right)}{k_\delta} \leq c_w(\tau) + \chi\gamma_u^2(1-c_\infty)c_u^{\mathrm{opt}}(\tau), \tag{121}$$

where χ is a constant that can be chosen arbitrarily close to 1 as $\delta \to 0$.

Proof. First note that by Propositions F.3 and F.4 we know that for any $k_0 > 0$ there exists a $\chi(k_0) > 0$ such that

$$\forall k \geq k_0 : (1-\rho_k)(1-p_k)c_u^{\mathrm{opt}} + \rho_k c_u^{\mathrm{rand}} \geq \chi(k_0)(1-c_\infty)c_u^{\mathrm{opt}}, \tag{122}$$

i.e., the excitation provided by Algorithm 2 is $\chi(k_0)$-optimal. Importantly, $\chi(k_0) \to 1$ as k_0 grows, hence we can push $\chi(k_0)$ arbitrarily close to 1 by increasing k_0. Note that, correctness of Algorithm 1 follows immediately from the proof of Theorem 4.2. Whenever $k \geq k_0$ we can use (122) to bound the termination time of Algorithm 1 in (74) by

$$8\left(\frac{1}{2\eta} + 1\right)\log\left(\frac{N}{\delta}\right) \leq \tau\left(kc_w + \sum_{j=1}^{k_0-1} c_{u_j}^{\mathrm{Alg}}(\tau)\gamma_u^2 + \sum_{j=k_0}^{k}\chi(k_0)(1-c_\infty)u_u^{\mathrm{opt}}\gamma_u^2\right) \tag{123}$$

$$= \tau\left(kc_w + \gamma_u^2\left((k-k_0)\chi(k_0)(1-c_\infty)c_u^{\mathrm{opt}} + \sum_{j=1}^{k_0-1} c_{u_j}^{\mathrm{Alg}}(\tau)\right)\right) \tag{124}$$

Taking the limit $\delta \to 0$, we realize that $\delta \to 0$ requires $k \to \infty$ since all quantities on the right-hand side besides k are fixed. As $k \to \infty$, the term $\sum_{j=1}^{k_0-1} c_{u_j}^{\mathrm{Alg}}(\tau)$ is dominated by the rest of the inequality, yielding the result. □

F.1 Intermediate Results for the Proof of Theorem F.1

Corollary F.2. *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 $c_{w_j}(\tau)$. Let η chosen to satisfy*

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)^\top \Sigma_w^{-1}\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)} \tag{125}$$

where $\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)$ is defined in (81). Then the estimate $\hat{\theta}_k$ drawn in Line 2 of Algorithm 2 satisfies

$$\mathbb{P}\left[\hat{\theta}_k \neq \theta_i\right] \leq N \exp\left(-\frac{\eta\tau}{4}\sum_{j=1}^{k} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau)\right). \tag{126}$$

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

$$\varepsilon_{\theta_i}(t_0, t_1) \coloneqq \sum_{t=t_0}^{t_1-1} \|x(t+1) - A_i x(t) - B_i u(t)\|^2_{\Sigma_w^{-1}}. \tag{127}$$

First note, that by similar arguments as in (Muehlebach et al., 2025, Lemma B.1) we have

$$\mathbb{P}\left[\hat{\theta}_k = \theta_i\right] \leq \mathbb{E}\left[e^{-\eta\left(\varepsilon_{\theta_i}(0,k\tau) - \varepsilon_{\theta_*}(0,k\tau)\right)}\right]. \tag{128}$$

By using Proposition E.2 we obtain

$$\mathbb{P}\left[\hat{\theta}_k = \theta_i\right] \leq \mathbb{E}\left[e^{-\eta\left(\varepsilon_{\theta_i}(0,k\tau) - \varepsilon_{\theta_*}(0,k\tau)\right)}\right] \leq \exp\left(-\frac{\eta\tau}{4}\sum_{j=1}^{k} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau)\right). \tag{129}$$

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

□

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 Muehlebach et al. (2025).

Proposition F.3. *Consider Algorithm 2 with $\rho_0 = c_0 \in (0, 1)$ and $\lim_{k\to\infty} \rho_k = c_\infty \in [0, c_0]$. Let η be chosen to satisfy*

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)^\top \Sigma_w^{-1}\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)\right)} \tag{130}$$

where $\Sigma^{\frac{1}{2}}_{\Delta_i}(\tau)$ is defined in (81). Then

$$\lim_{k\to\infty} c_{u_k}^{\text{Alg}} \geq (1 - c_\infty)c_u^{\text{opt}}.$$

Proof. This proof is carried out by contradiction. Suppose there exists some $\tilde{c} \in [0, 1)$ such that

$$\lim_{k\to\infty} c_{u_k}^{\text{Alg}} = \tilde{c}(1 - c_\infty)c_u^{\text{opt}} < (1 - c_\infty)c_u^{\text{opt}}. \tag{131}$$

plugging in the definition of $c_{u_k}^{\text{Alg}}$ we see this requires

$$\lim_{k\to\infty}(1 - p_k) = \tilde{c} \Leftrightarrow \lim_{k\to\infty} p_k = 1 - \tilde{c} \tag{132}$$

Invoking Corollary F.2 we that

$$p_k \leq N \exp\left(-\frac{\eta\tau}{4}\left(kc_w(\tau) + \sum_{j=1}^{k} c_{u_j}^{\text{Alg}}(\tau)\gamma_u^2\right)\right). \tag{133}$$

Hence, (132) requires

$$\lim_{k\to\infty} kc_w + \sum_{j=1}^{k} c_{u_j}^{\text{Alg}}(\tau)\gamma_u^2 = \frac{4}{\tau\eta}\log\left(\frac{1-\tilde{c}}{N}\right). \tag{134}$$

Since $c_w \geq 0$ we require $c_w = 0$. This is unlikely but for the remainder, we consider this case to be true. In any case, we additionally require

$$\lim_{k\to\infty} c_{u_k}^{\text{Alg}} = 0 \tag{135}$$

since if this is not true, the sum in (134) does not converge as $k \to \infty$. Going back to the definition of $c_{u_k}^{\text{Alg}}$ we see that (134) can only be satisfied if the conditions

- $\lim_{k\to\infty} \rho_k = 0$
- $\lim_{k\to\infty}(1-\rho_k)(1-p_k) = 0$

are both met. Clearly, if $c_\infty \neq 0$, the first condition is not met, yielding the result for all $c_\infty \in (0, c_0]$. For the case where $c_\infty = 0$ the second condition requires that $\lim_{k\to\infty} p_k = 1$. Going back to (133), we see that for this to happen, we require

$$\lim_{k\to\infty} \sum_{j=1}^{k} c_{u_j}^{\text{Alg}}\gamma_u^2 = 0. \tag{136}$$

Since we selected $\rho_0 = c_0 > 0$ we obtain $c_{u_0}^{\text{Alg}} \geq c_0 c_u^{\text{rand}} > 0$ which contradicts the condition. Hence, there cannot exist a $\tilde{c} \in [0,1)$ for which

$$\lim_{k\to\infty} c_{u_k}^{\text{Alg}} = \tilde{c}(1-c_\infty)c_u^{\text{opt}} < (1-c_\infty)c_u^{\text{opt}}. \tag{137}$$

which completes the proof. □

Proposition F.4. *Consider Algorithm 2 with $\rho_0 = c_0 \in (0,1)$ and $\rho_{k-1} \geq \rho_k \geq (1-\varepsilon)\rho_{k-1}$. Let η be chosen to satisfy*

$$\eta\tau \leq \frac{1}{512\lambda_{\max}\left(\left(\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)^\top \Sigma_w^{-1}\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)\right)} \tag{138}$$

where $\Sigma_{\Delta_i}^{\frac{1}{2}}(\tau)$ is defined in (81). If $\varepsilon < \frac{\left(1-e^{-\frac{\eta\tau}{4}c_w(\tau)}\right)(1-\rho_0)}{\rho_0 e^{-\frac{\eta\tau}{4}c_w(\tau)}}$, then $c_{u_k}^{\text{Alg}} \geq c_{u_{k-1}}^{\text{Alg}}$ for all $k \geq 1$.

Proof. Plugging in the definitions of the PE coefficients, the statement of the theorem can be rewritten as

$$(1-\rho_k)(1-p_k)c_u^{\text{opt}} + \rho_k c_u^{\text{rand}} > (1-\rho_{k-1})(1-p_{k-1})c_u^{\text{opt}} + \rho_{k-1}c_u^{\text{rand}}. \tag{139}$$

Rearanging terms, we obtain

$$\begin{aligned} & c_u^{\text{opt}}\Big((1-\rho_k)(1-p_k) - (1-\rho_{k-1})(1-p_{k-1})\Big) && > c_u^{\text{rand}}(\rho_{k-1}-\rho_k) && (140)\\ \Leftrightarrow\quad & c_u^{\text{opt}}\Big(1-\rho_k-p_k+p_k\rho_k - (1-\rho_{k-1}-p_{k-1}+p_{k-1}\rho_{k-1})\Big) && > c_u^{\text{rand}}(\rho_{k-1}-\rho_k) && (141)\\ \Leftrightarrow\quad & c_u^{\text{opt}}(\rho_{k-1}-\rho_k+p_{k-1}-p_k+p_k\rho_k-p_{k-1}\rho_{k-1}) && > c_u^{\text{rand}}(\rho_{k-1}-\rho_k) && (142)\end{aligned}$$

Clearly, since ρ_k is a non-increasing sequence, (142) is satisfied whenever

$$p_{k-1} - p_k + p_k\rho_k - p_{k-1}\rho_{k-1} > 0 \tag{143}$$

Note that, by Corollary F.2 we have $p_k = N\exp\left(-\frac{\eta\tau}{4}\sum_{j=1}^{k} c_{u_j}(\tau)\gamma_u^2 + c_w(\tau)\right)$ and hence

$$p_k = p_{k-1}e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2 + c_w(\tau))}. \tag{144}$$

Using this, and the fact that $\rho_k \geq \rho_{k-1}(1-\varepsilon)$ we have that (143) is implied by

$$p_{k-1}\left(1 - e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2 + c_w(\tau))} + \rho_{k-1}(1-\varepsilon)e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2 + c_w(\tau))} - \rho_{k-1}\right) > 0, \tag{145}$$

which can only be satisfied if $1-e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}+\rho_{k-1}(1-\varepsilon)e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}-\rho_{k-1}>0$. Using standard algebraic manipulations, we can rewrite this as

$$\left(1-e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}\right)(1-\rho_{k-1})>\rho_{k-1}\varepsilon e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}. \tag{146}$$

Solving this for ε we obtain

$$\varepsilon<\frac{\left(1-e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}\right)(1-\rho_{k-1})}{\rho_{k-1}e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}}. \tag{147}$$

Since we require an upper bound for all $k\geq 1$, we use the (loose) lower bound to the right-hand side

$$\frac{\left(1-e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}\right)(1-\rho_{k-1})}{\rho_{k-1}e^{-\frac{\eta\tau}{4}(c_{u_k}(\tau)\gamma_u^2+c_w(\tau))}}>\frac{\left(1-e^{-\frac{\eta\tau}{4}c_w(\tau)}\right)(1-\rho_0)}{\rho_0 e^{-\frac{\eta\tau}{4}c_w(\tau)}}, \tag{148}$$

to obtain that if $\varepsilon<\frac{\left(1-e^{-\frac{\eta\tau}{4}c_w(\tau)}\right)(1-\rho_0)}{\rho_0 e^{-\frac{\eta\tau}{4}c_w(\tau)}}$ we have $c_{u_k}^{\mathrm{Alg}}>c_{u_{k-1}}^{\mathrm{Alg}}$ for all $k\geq 1$. □

Note that by Proposition F.4 as long as $\rho_0\neq 1$ we can show that the sequence $\{c_{u_k}^{\mathrm{Alg}}\}_{k\geq 0}$ increases at every timestep. However, even if we initially collect some data using random excitations (i.e., we select $\rho_0=1$), we can still guarantee that after a potential initial decay $\{c_{u_k}^{\mathrm{Alg}}\}_{k\geq 1}$ increases at every timestep.

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 $u(t)=(1-\rho)u_{\mathrm{d}}(t)+\rho u_{\mathrm{p}}(t)$, with $\rho\in[0,1]$, $u_{\mathrm{p}}(t)\overset{i.i.d.}{\sim}\mathcal{N}(0,\sigma_u^2 I_{n_u})$ and a deterministic control input $u_{\mathrm{d}}(t)$. Then for any fixed matrix $M\in\mathbb{S}_{++}^{n_x}$, any $i\in[1,N]$, and any $\tau\geq 0$ it holds that*

$$\begin{aligned}\sum_{t=0}^{\tau}\mathbb{E}\left[\|\Delta A_i x(t)+\Delta B_i u(t)\|_M^2\right]&=(1-\rho)^2U_{\mathrm{d}}^\top W_i(\tau)U_{\mathrm{d}}+2(1-\rho)U_{\mathrm{d}}^\top m_i(x_0)+c_i(x(0))\\&\quad+\sigma_u^2\rho^2\mathrm{tr}\left(W_i(\tau)\right)+\mathrm{tr}\left(S_w(\tau)^\top Q_M^i(\tau)S_w(\tau)\right),\end{aligned} \tag{149}$$

where

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

Furthermore,

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

$$\begin{aligned}\sum_{t=0}^{\tau}\mathbb{E}\left[\|\Delta A_i x(t)+\Delta B_i u(t)\|_M^2\right]&\geq(1-\rho)^2U_{\mathrm{d}}^\top W_i(\tau)U_{\mathrm{d}}+\sigma_u^2\rho^2\mathrm{tr}\left(W_i(\tau)\right)\\&\quad+\mathrm{tr}\left(S_w(\tau)^\top Q_M^i(\tau)S_w(\tau)\right)\end{aligned} \tag{150}$$

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

$$\sum_{t=0}^{\tau}\mathbb{E}\left[\|\Delta A_i x(t)+\Delta B_i u(t)\|_M^2\right]\geq\sigma_u^2\mathrm{tr}\left(W_i(\tau)\right)+\mathrm{tr}\left(S_w(\tau)^\top Q_M^i(\tau)S_w(\tau)\right) \tag{151}$$

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

$$\Delta A_i x(t) + \Delta B_i u(t) = \Delta A_i x(t) + (1-\rho)\Delta B_i u_{\mathrm{d}}(t) + \rho \Delta B_i u_{\mathrm{p}}(t) \tag{152}$$

$$= \underbrace{\Delta A_i \left(\sum_{s=0}^{t-1} (1-\rho) A_*^{s-t-1} B_* u_{\mathrm{d}}(s) + A_*^s x(0) \right) + (1-\rho)\Delta B_i u_{\mathrm{d}}(t)}_{\text{deterministic}} \tag{153}$$

$$+ \underbrace{\rho \Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_{\mathrm{p}}(t) + \rho \Delta B_i u_{\mathrm{p}}(t)}_{\text{rand. vec., dep. on } u_{\mathrm{p}}} + \underbrace{\Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} w(t)}_{\text{rand. vec., dep. on } w}.$$

Using that all, random vectors are independent and zero mean, and thus, cross terms are zero, it follows by plugging in (153) that

$$\begin{aligned}
&\sum_{t=0}^{\tau} \mathbb{E}[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2] \\
&\quad = \sum_{t=0}^{\tau} \mathbb{E}\left[(\Delta A_i x(t) + \Delta B_i u(t))^\top M (\Delta A_i x(t) + \Delta B_i u(t)) \right] \\
&\quad = \sum_{t=0}^{\tau} \Bigg(\left\| \Delta A_i \left((1-\rho) \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_{\mathrm{d}}(s) + A_*^s x(0) \right) + (1-\rho)\Delta B_i u_{\mathrm{d}}(t) \right\|_M^2 \\
&\qquad + \rho^2 \mathbb{E}\left[\left\| \Delta B_i u_{\mathrm{p}}(t) + \Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_{\mathrm{p}}(s) \right\|_M^2 \right] \\
&\qquad + \mathbb{E}\left[\left\| \Delta A_i \sum_{s=0}^{t-1} A_*^{s-t-1} w(s) \right\|_M^2 \right] \Bigg).
\end{aligned} \tag{154}$$

We carry on by rewriting the first term in (154)

$$\begin{aligned}
&\sum_{t=0}^{\tau} \left\| \Delta A_i \left((1-\rho) \sum_{s=0}^{t-1} A_*^{s-t-1} B_* u_{\mathrm{d}}(s) + A_*^s x(0) \right) + (1-\rho)\Delta B_i u_{\mathrm{d}}(t) \right\|_M^2 \\
&\quad = (1-\rho)^2 U_{\mathrm{d}}^\top W_i(\tau) U_{\mathrm{d}} + \underbrace{\sum_{t=0}^{\tau} \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right)^\top M \left(\Delta A_i \sum_{s=0}^{t-1} A_*^s x(0) \right)}_{:= c_i(x(0)) \geq 0} \\
&\qquad + 2(1-\rho) U_{\mathrm{d}}^\top \underbrace{\left(S_u(\tau)^\top Q_M^i(\tau) S_w(\tau) X(0) + N_M^i(\tau) S_w(\tau) X(0) \right)}_{:= m_i(x_0)},
\end{aligned} \tag{155}$$

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 (155) to obtain

$$\begin{aligned}
\sum_{t=0}^{\tau} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2 \right] &= (1-\rho)^2 U_{\mathrm{d}}^\top W_i(\tau) U_{\mathrm{d}} + 2(1-\rho) U_{\mathrm{d}}^\top m_i(x_0) + c_i(x(0)) \\
&\quad + \rho^2 \mathbb{E}\left[U_{\mathrm{p}} W_i(\tau) U_{\mathrm{p}} \right] + \mathbb{E}\left[W^\top S_w(\tau)^\top Q_M^i(\tau) S_w(\tau) W \right]
\end{aligned} \tag{156}$$

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

$$\begin{aligned}
\sum_{t=0}^{\tau} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|_M^2 \right] &= (1-\rho)^2 U_{\mathrm{d}}^\top W_i(\tau) U_{\mathrm{d}} + 2(1-\rho) U_{\mathrm{d}}^\top m_i(x_0) + c_i(x(0)) \\
&\quad + \sigma_u^2 \rho^2 \mathrm{tr}\,(W_i(\tau)) + \mathrm{tr}\left(S_w(\tau)^\top Q_M^i(\tau) S_w(\tau) \right)
\end{aligned} \tag{157}$$

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 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

$$\tilde{A}_* = \begin{bmatrix} 0 & 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.9 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.9 \end{bmatrix}, \qquad \tilde{B}_* = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

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

Further, we fix $\gamma_u = 1$ and consider the process noise $w(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_w^2 I_6)$, 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 (28) 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.

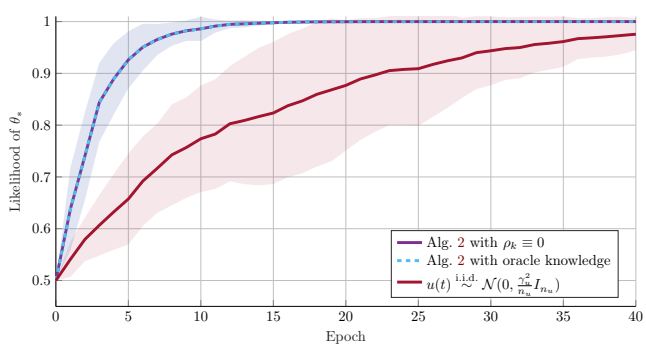

Figure 2: Mean and σ-band of the likelihood of θ_* given the data for different data collection strategies for the motivating example.