Hidden Convexity in Active Learning: A Convexified Online Input Design for ARX Systems

Nicolas Chatzikiriakos, Bowen Song, Philipp Rank, and Andrea Iannelli

***Abstract*— The goal of this work is to accelerate the identification of an unknown ARX system from trajectory data through online input design. Specifically, we present an active learning algorithm that sequentially selects the input to excite the system according to an experiment design criterion using the past measured data. The adopted criterion yields a non-convex optimization problem, but we provide an exact convex reformulation allowing to find the global optimizer in a computationally tractable way. Moreover, we give sample complexity bounds on the estimation error due to the stochastic noise. Numerical studies showcase the effectiveness of our algorithm and the benefits of the convex reformulation.**

***Index Terms*— Input Design, Convex Optimization, Finite Sample Identification**

I. Introduction

Input design methods from system identification leverage the flexibility in choosing the excitation input to accelerate the learning process of identifying an unknown system from data. However, existing works suffer from relying on computationally expensive or intractable operations [1] or only provide asymptotic guarantees [2], [3]. To remedy this, in this work, we present a sample efficient algorithm that solves a computationally tractable problem to sequentially perform optimal experiments for fast learning of unknown autoregressive systems with exogenous inputs (ARX). Importantly, the algorithm relies on solving a convex optimization problem to design optimal excitations and is equipped with finite sample guarantees. Our work is closely connected to classical experiment design [3], [4] and recent works in finite sample system identification [1], [5]. Next, we review the relevant literature to paint a clearer picture of existing gaps.

Classical input design: The design of optimal experiments has a long history in the literature of system identification; see [2], [6] for an overview on the topic. One key problem in experiment design is that the design criteria, which are often a function of the Fisher-information matrix, naturally depend on the unknown system. To circumvent this problem, some works focus on a robust approach, typically using minimax objectives (see [7] and the references therein). Another research thread focused on adaptively designing inputs based on a running estimate [8]. In particular, [3] considered the problem of adaptive input design for ARX systems in an asymptotic setting, yielding parameter convergence and an asymptotically optimal covariance matrix. To tackle cases where the optimization criterion is non-convex, [4] leverages convex relaxations to convexify the design objective. However, even though in parts computationally attractive, these works only analyze the asymptotic behavior of the identification algorithms as the number of samples goes to infinity, whereas in practice, only finite data is available.

Finite sample identification and active learning: In recent years, due to novel results in high-dimensional statistics [9], there has been an increased interest in the finite sample analysis of system identification; see [5] for an overview of the topic. Primarily, the literature focused on analyzing system identification under random excitations for different classes of fully observed dynamical systems such as linear [10], [11], bilinear [12], [13] and certain classes of nonlinear systems [14], [15]. Further, there have been considerable efforts in the analysis of learning a system description from input-output data [5], [16]. Based on these findings, some works departed from purely random excitation and tried to leverage the excitation input to obtain a more sample-efficient identification. The first work to address this problem was [1], in which the authors provide a sequential algorithm resembling a finite-sample version of E-optimal experiment design. In [17], the same authors take downstream tasks and objectives into account, enabling task-specific learning. Subsequent works [18], [19] analyze algorithms for input design for non-linear systems under a finite sample perspective. However, even for the simplest case of fully observed linear systems, the optimal excitation is determined by solving a non-convex optimization problem. Thus, the global optimizer cannot be found efficiently, yielding a computationally unattractive algorithm. To remedy this, [17] proposes a convex relaxation of the problem. While this relaxation makes the algorithm computationally tractable, it is not guaranteed to recover the optimal solution of the original, non-convex problem. Finally, input design algorithms with finite sample guarantees are not available for ARX systems.

A. Contribution

We provide a new input design algorithm that sequentially generates periodic inputs to enable fast identification of unknown ARX systems affected by Gaussian process noise. Our analysis also gives an upper bound on the finite sample estimation error of the unknown parameters. To do so, we leverage that the covariates' evolution can be represented by a state-space system, enabling us to use existing tools for

NC and AI 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). BS acknowledges the support of the International Max Planck Research School for Intelligent Systems (IMPRS-IS).

The authors are with the University of Stuttgart, Institute for Systems Theory and Automatic Control, 70550 Stuttgart, Germany (e-mail: {nicolas.chatzikiriakos, bowen.song, andrea.iannelli}@ist.uni-stuttgart.de.)

the finite sample analysis. This results in an upper bound on the identification error of the parameters of the unknown ARX system. In addition, inspired by considerations in [1], [4], we propose a convex reformulation to leverage a hidden convexity in the input design objective. We do so by transforming the non-convex optimization problem into a convex one from which the *global* optimizer of the original non-convex objective can be reconstructed. Thus, the proposed algorithm only requires solving a convex optimization problem online. Importantly, the convex reformulation deployed is not restricted to ARX systems, but can also be used for fully observed linear systems and hence is of independent interest. With this, we provide a computationally tractable input design scheme equipped with finite sample guarantees, and therefore bridge some of the gaps between the active learning and classical input design literature. In concluding simulations, we demonstrate the effectiveness of our algorithm by comparing its estimation error to random excitations and the related approach from [17]. The experiments show that our algorithm achieves the lowest estimation error, thus enhancing learning, while being computationally attractive and giving a guaranteed estimation accuracy.

B. Notation

Given a matrix M we denote its spectral radius, smallest eigenvalue, and spectral norm by $\rho(M)$, $\lambda_{\min}(M)$, and $\|M\|$, respectively. We denote universal constants by C. The set of all non-negative integers (the set of all non-negative integers smaller than k) is indicated by $\mathbb{N}$ ($[k]$). We use S_+^n to denote the cone of positive semi-definite matrices of dimension n. We use $[a]_i$ ($[A]_{ij}$) when referring to the i-th (ij-th) element of a vector a (matrix A). Given matrices A_i, $i \in [1,p]$ we use $A_{1:p} = \begin{bmatrix} A_1 & \dots & A_p \end{bmatrix}$. We write A^{H} to denote the conjugate transpose of a complex matrix A. Given a periodic time-domain signal $\{x_t\}_{t=1}^k$ we express its Fourier coefficients with $\check{x}_t$ and the Fourier transform by $\mathcal{F}\{\cdot\}$.

II. Preliminaries and problem statement

We consider the problem of identifying an unknown ARX system of known orders $p, q > 0$ described by

$$y_t = \sum_{i=1}^{p} A_i^* y_{t-i} + \sum_{j=1}^{q} B_j^* u_{t-j} + \Sigma_w^{\frac{1}{2}} w_t, \tag{1}$$

where $y_t \in \mathbb{R}^{n_y}$ and $u_t \in \mathbb{R}^{n_u}$ are the output and input of the system at timestep t, and $w_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{n_x})$ is normalized process noise. We assume that the noise covariance matrix Σ_w is known. As initial conditions we assume $y_{-i} = 0$, $u_{-i} = 0$, for integers $i \in [1,q]$. The goal of this work is to identify the unknown system matrices $\{A_i^*\}_{i=1}^p$, $\{B_j^*\}_{j=1}^q$ from data $\{y_t\}_{t=0}^T$, $\{u_t\}_{t=0}^{T-1}$ collected from a single trajectory. To this end, set $n_x \coloneqq (pn_y + qn_u)$ and define $\theta^* \coloneqq \begin{bmatrix} A_{1:p}^* & B_{1:p}^* \end{bmatrix} \in \mathbb{R}^{n_y \times n_x}$ and $x_t \coloneqq \begin{bmatrix} y_{t-1:t-p}^\top & u_{t-1:t-q}^\top \end{bmatrix}^\top \in \mathbb{R}^{n_x}$ to rewrite (1) as

$$y_t = \theta^* x_t + \Sigma_w^{\frac{1}{2}} w_t. \tag{2}$$

To identify the unknown matrix θ^* we use the ordinary least squares (OLS) estimator

$$\hat{\theta}_T = \underset{\theta \in \mathbb{R}^{n_y \times n_x}}{\arg\min} \frac{1}{T} \sum_{t=1}^{T} \|y_t - \theta x_t\|_2^2. \tag{3}$$

Assuming that the inverse of the empirical covariance matrix

$$\hat{\Sigma}_T \coloneqq \frac{1}{T} \sum_{t=1}^{T} x_t x_t^\top \tag{4}$$

exists, the OLS (3) admits a well-known, unique closed-form solution, and its estimation error is given by

$$\hat{\theta}_T - \theta^* = \sum_{t=1}^{T} w_t x_t^\top \Big(\sum_{t=1}^{T} x_t x_t^\top \Big)^{-\frac{1}{2}} \Big(\sum_{t=1}^{T} x_t x_t^\top \Big)^{-\frac{1}{2}}. \tag{5}$$

As derived in [5], the evolution of the covariates can be represented by the non-minimal process

$$x_{t+1} = \mathcal{A}^* x_t + \mathcal{B}_u u_t + \mathcal{B}_w w_t, \tag{6}$$

where $\mathcal{A}^* \coloneqq \begin{bmatrix} \mathcal{A}_{11}^* & \mathcal{A}_{12}^* \\ 0 & \mathcal{A}_{22} \end{bmatrix}$ and

$$\mathcal{A}_{11}^* \coloneqq \begin{bmatrix} A_1^* & A_2^* & \dots & A_p^* \\ I & 0 & \dots & 0 \\ 0 & I & \dots & 0 \\ 0 & \dots & I & 0 \end{bmatrix}, \tag{7a}$$

$$\mathcal{A}_{12}^* \coloneqq \begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix}^\top \otimes B_{1:q}^*, \tag{7b}$$

$$\mathcal{A}_{22} \coloneqq \begin{bmatrix} 0 & \dots & \dots & 0 \\ I_{n_u} & 0 & \dots & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & \dots & I_{n_u} & 0 \end{bmatrix}, \tag{7c}$$

$$\mathcal{B}_u \coloneqq \begin{bmatrix} 0_{n_u \times pn_y} & I_{n_u} & 0_{n_u \times (q-1)n_u} \end{bmatrix}^\top, \tag{7d}$$

$$\mathcal{B}_w \coloneqq \begin{bmatrix} \Sigma_w^{\frac{1}{2}} & n_y \times 0_{((p-1)n_y + qn_u)} \end{bmatrix}^\top. \tag{7e}$$

Subsequently, when writing $\hat{\mathcal{A}}_T$ we refer to the matrix $\mathcal{A}$ where all unknown quantities have been replaced by the estimate $\hat{\theta}_T$. Note that the matrices $\mathcal{B}_u$ and $\mathcal{B}_w$ are knowns. In line with the relevant literature [1], [3] we assume that the ARX system (1) is asymptotically stable.

Assumption 1: System (1) is asymptotically stable, i.e., $\rho(\mathcal{A}_{11}) < 1$.

While Assumption 1 might seem restrictive, related works providing a finite sample analysis of system identification with random inputs are restricted to marginally stable systems [5], [10], and learning unstable systems using the OLS has been shown to be statistically inconsistent [20]. The slightly stronger requirement of asymptotic stability is necessary for the proposed active learning algorithm (see [1, Remark B.5]) and cannot easily be circumvented.

We now define the controllability Gramians with respect to the control input and the process noise, respectively

$$\Gamma_t(\mathcal{A}, \mathcal{B}_u) \coloneqq \sum_{s=0}^{t-1} \mathcal{A}^s \mathcal{B}_u \mathcal{B}_u^\top \mathcal{A}^{s\top}, \tag{8a}$$

$$\Gamma_t(\mathcal{A}, \mathcal{B}_w) \coloneqq \sum_{s=0}^{t-1} \mathcal{A}^s \mathcal{B}_w \mathcal{B}_w^\top \mathcal{A}^{s\top}. \tag{8b}$$

Given a periodic input u_t of period k and bounded energy $\sum_{t=0}^{k-1} u_t^\top u_t \leq \gamma^2 k$ we denote the normalized steady-state covariance when $w_t \equiv 0$, as

$$\Gamma_k^{\check{u}}(\mathcal{A}, \mathcal{B}_u) := \lim_{T\to\infty} \frac{1}{T\gamma^2} \sum_{t=1}^{T} x_t x_t^\top \tag{9}$$
$$= \frac{1}{\gamma^2 k^2} \sum_{\ell=0}^{k-1} (e^{j\frac{2\pi\ell}{k}} I - \mathcal{A})^{-1} \mathcal{B}_u \check{u}_\ell \check{u}_\ell^{\mathsf{H}} \mathcal{B}_u^{\mathsf{H}} (e^{j\frac{2\pi\ell}{k}} I - \mathcal{A})^{-\mathsf{H}},$$

where the equality holds due to Parseval's Theorem. Given matrices $\{\check{U}_\ell\}_{\ell\in[k]} \in \mathrm{S}_+^{n_u}$, with abuse of notation, we define

$$\Gamma_k^{\check{U}}(\mathcal{A}, \mathcal{B}_u) \tag{10}$$
$$:= \frac{1}{\gamma^2 k^2} \sum_{\ell=0}^{k-1} (e^{j\frac{2\pi\ell}{k}} I - \mathcal{A})^{-1} \mathcal{B}_u \check{U}_\ell \mathcal{B}_u^{\mathsf{H}} (e^{j\frac{2\pi\ell}{k}} I - \mathcal{A})^{-\mathsf{H}}.$$

It is important to observe at this point that (9) can be equivalently represented by (10) by taking $\check{U}_\ell = \check{u}_\ell \check{u}_\ell^{\mathsf{H}}$, but the opposite is only true if $\mathrm{rank}(\check{\mathrm{U}}_\ell) = 1$.

III. Active learning for ARX systems

In [5], it was shown that isotropic Gaussian excitations are sufficient to identify the unknown ARX system (1) from trajectory data. However, in most cases isotropic Gaussian excitations are not optimal in terms of the resulting sample complexity. To improve the sample complexity of identifying (1) we propose to sequentially generate informative data using Algorithm 1, which is inspired by the algorithm proposed in [1] for fully observed linear systems.

Algorithm 1 Active Learning Algorithm

Require: γ^2, T_0, k_0

Collect dataset $\mathcal{D}_0 = \left\{\{y_t\}_{t=0}^{T_0}, \{u_t\}_{t=0}^{T_0}\right\}$ using $u_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma^2}{n_u} I_{n_u})$
Estimate $\hat{\theta}_0$ from $\mathcal{D}_0$ by solving (3)
Set $T = T_0$, $\mathcal{D} = \mathcal{D}_0$
for i = 1,2,... **do**
  Set $T_i = 3T_{i-1}$, $k_i = 2k_{i-1}$
  $u^* = \textsc{OptInput}(T, T_i, \hat{\theta}_i, k_i, \frac{\gamma^2}{2}, \mathcal{D})$
  Set $u_t^\eta = u_t^* + \eta_t$, where $\eta_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma^2}{2n_u} I_{n_u})$
  Collect $\mathcal{D}_i = \left\{\{y_t\}_{t=T}^{T+T_i}, \{u_t^\eta\}_{t=T}^{T+T_i}\right\}$ using u_t^η
  Estimate $\hat{\theta}_i$ from $\mathcal{D} = \bigcup_{j=0}^{i} \mathcal{D}_j$ by solving (3)
  Set $T = T + T_i$

In essence, Algorithm 1 sequentially generates periodic signals that excite the true system (1) according to some optimality criterion. Hereby, the length of each episode T_i is increased exponentially to put more weight on inputs generated with more accurate estimates. Similar strategies have also been used successfully in other sequential learning problems, such as best-arm identification [21]. In our work, we consider periodic input signals since, as shown in [1], this class is general enough to achieve an optimal rate. The number of excitation frequencies k_i is increased sequentially, although at a slower rate than T_i. The obtained periodic signal is repeated until the end of the episode. For technical reasons, the algorithm requires to check whether all frequencies are safe to plan with, which is why in the original formulation an uncertainty estimate ϵ_i is obtained in each round. Empirically, this step is not necessary (see [1, Remark 2.5] for a detailed discussion) and hence is omitted here for clarity of exposition. However, it is important to note that this step is necessary for the proof of Theorem 1. As information criterion driving input selection, we take the viewpoint of E-optimality [2], [7]. That is, the input sequence is selected to maximize the smallest eigenvalue of the empirical covariance matrix (4), while satisfying an energy constraint on the input. As shown in [1], this translates to optimizing the Fourier coefficients of the optimal input sequence via the non-convex optimization problem

$$\check{u}^* \in \underset{\{\check{u}_\ell\}_{\ell\in[k_i]}}{\arg\max}\ \lambda_{\min}\Big(\frac{\gamma^2}{2} T_i \Gamma_{k_i}^{\check{u}}(\hat{\mathcal{A}}_i, \mathcal{B}_u) + \sum_{t=1}^{T} x_t x_t^\top\Big) \tag{11a}$$
$$\text{s.t.} \quad \sum_{\ell\in[k_i]} \mathrm{Tr}(\check{u}_\ell \check{u}_\ell^{\mathsf{H}}) \leq \frac{k_i^2 \gamma^2}{2}. \tag{11b}$$

The matrix $\Gamma_{k_i}^{\check{u}}(\hat{\mathcal{A}}_i, \mathcal{B}_u)$ in the optimality criterion (11a) describes the predicted future effect of the designed excitation input on the covariance (9) and $\sum_{t=1}^{T} x_t x_t^\top$ is the empirical covariance based on the past measurements collected in $\mathcal{D}$. The constraint (11b) is necessary to ensure the energy bound for the input. We observe that (11) is non-convex in the Fourier coefficients of the unknown input signal. After solving (11) the designed input is computed using the inverse Fourier transform

$$\textsc{OptInput}(T, T_i, \hat{\theta}_i, k_i, \tfrac{\gamma^2}{2}, \mathcal{D}) = \mathcal{F}^{-1}\{\check{u}^*\}. \tag{12}$$

Since θ_* is unknown to the learning algorithm, the estimate $\hat{\theta}_i$ is used to predict the future excitation in the optimization.

In the subsequent Section III-A, we show that Algorithm 1 enjoys finite sample guarantees for the identification of ARX systems when the input is chosen as the solution of (11)-(12), which as noted earlier is a non-convex optimization problem. This is followed by Section III-B, where we provide an exact convex reformulation of (11), so that the optimal excitation can be determined efficiently and reliably.

A. Finite sample identification of ARX Systems

To simplify our presentation, we make the following additional assumption.

Assumption 2: The hyperparameters T_0 and k_0 are chosen such that $T_0 \geq f_{\theta_*}(k_0)$ for some known function $f_{\theta_*}(k_0)$.

We do not provide the exact expression for f_{θ_*} here to streamline the discussion and refer to [1, Appendix B] for the definition. Note that, even though f_{θ_*} depends on θ_* Assumption 2 is not restrictive, since the sequence $\{T_i\}_{i\geq 0}$ increases faster than $\{k_i\}_{i\geq 0}$. Thus, even if Assumption 2 is not satisfied for $i = 0$ it will eventually be satisfied; see [1, Appendix B] for a more detailed discussion. The sample complexity analysis of Algorithm 1 relies on the

system to reach steady-state for the given input sequence. To characterize the transient behavior of process (6) consider

$$\beta(\mathcal{A}^*) := \sup_{k\in\mathbb{N}} \|\mathcal{A}^{*k}\| \left(\tfrac{1}{2} + \tfrac{\rho(\mathcal{A}^*)}{2}\right)^{-k}. \tag{13}$$

As discussed in [1], $\beta(\mathcal{A}^*)$ is finite when Assumption 1 holds. We are now ready to provide our first main result.

Theorem 1: Let Assumptions 1 and 2 hold and fix some $\delta \in (0,1)$. Then Algorithm 1 will produce inputs that satisfy $\mathbb{E}\left[\sum_{t=1}^{T}(u_t^\eta)^\top u_t^\eta\right] \leq \gamma^2 T$. If further $T \geq \bar{T}(\theta_*, \frac{1}{\delta}, k(T))$, where $\bar{T}(\theta_*, \frac{1}{\delta}, k(T))$ is defined in [1, Appendix B], the estimate $\hat{\theta}_T$ produced by Algorithm 1 will satisfy (14), where

$$\bar{\Gamma}_T = \mathcal{O}\left(\frac{\beta(\mathcal{A}^*)^2\gamma^2 T}{(1-\rho(\mathcal{A}^*))^2}\right) I_{n_x} \tag{15}$$

and $\check{u}^* = \mathcal{F}\{\textsc{OptInput}(T, \theta^*, k(T), \gamma^2, \emptyset)\}$.

Proof: First, note that by (3) the estimation error can be bounded as

$$\begin{aligned}\|\hat{\theta}_T - \theta^*\| \leq & \lambda_{\min}\Big(\sum_{t=1}^{T} x_t x_t^\top\Big)^{-\frac{1}{2}} \\ & \times \left\|\sum_{t=1}^{T} w_t x_t^\top \Big(\sum_{t=1}^{T} x_t x_t^\top\Big)^{-\frac{1}{2}}\right\|.\end{aligned} \tag{16}$$

Recall that the evolution of the covariates follows (6). Since $\mathcal{A}^*$ is block-diagonal we have

$$\rho(\mathcal{A}^*) = \max\Big(\rho(\mathcal{A}_{11}^*), \rho(\mathcal{A}_{22})\Big) = \rho(\mathcal{A}_{11}^*) < 1, \tag{17}$$

where the second equality follows since $\rho(\mathcal{A}_{22}) = 0$ by the lower triangular structure and the inequality follows from Assumption 1. With this, we can use the results from [1] to obtain Theorem 1. Specifically, the analysis and the active learning algorithm use (6) to quantify the influence of the control input and process noise on the covariates. ■

Remark 1: For clarity in Theorem 1 we only provide a finite sample error upper bound for Algorithm 1. However, the optimality analysis in [1] can also be adapted for ARX systems by using the same rationale as in the proof of Theorem 1.

B. Convexified input design

Up to this point, Algorithm 1 relied on solving the non-convex optimization problem (11). However, this is often intractable and might result only in *local* optima. This is an issue because the theoretical guarantees in Theorem 1 only hold when the solution of (11)-(12) is used, and the empirical performance of Algorithm 1 degrades when suboptimal solutions are used. To circumvent this, we propose an exact convex reformation of (11) and show that the optimal input can be obtained through a convex optimization problem, precisely, a Semidefinite Program.

Theorem 2: Let $\check{U}_\ell^*$, $\ell \in [k_i]$ be the solution of

$$\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in \mathbb{S}_+^{d_u}} \lambda_{\min}\left(\tfrac{\gamma^2}{2} T_i \Gamma_{k_i}^{\check{U}}(\hat{\mathcal{A}}_i, \mathcal{B}_u) + \sum_{t=1}^{T} x_t x_t^\top\right) \tag{18a}$$

$$\text{s.t.} \quad \sum_{\ell\in[k_i]} \operatorname{Tr}(\check{U}_\ell) \leq \frac{{k_i}^2\gamma^2}{2} \tag{18b}$$

and let $\check{u}_\ell^*$ be the eigenvector corresponding to the largest eigenvalue of $\check{U}_\ell^*$ for each $\ell \in [k_i]$. Then $\check{u}_\ell^*$, $\ell \in [k_i]$ is the solution of (11).

Proof: In this proof, we first show that (11) can be equivalently formulated as (18) with an additional non-convex rank constraint. We proceed by analyzing the objective arising in both optimization problems. Then, we analytically derive the optimizers of (11) and (18) and use them to show the result.

Part 1: Equivalent reformulation of (11)*:* Given a vector $\check{u}_\ell$, we define $\check{U}_\ell := \check{u}_\ell \check{u}_\ell^{\mathsf{H}}$ which is rank-1. This allows us to rewrite the optimization problem (11) as

$$\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in \mathbb{S}_+^{d_u}} \lambda_{\min}\Big(\frac{\gamma^2}{2} T_i \Gamma_{k_i}^{\check{U}}(\hat{\mathcal{A}}_i, \mathcal{B}_u) + \sum_{t=1}^{T} x_t x_t^\top\Big) \tag{19a}$$

$$\text{s.t.} \quad \sum_{\ell\in[k_i]} \operatorname{Tr}(\check{U}_\ell) \leq \frac{{k_i}^2\gamma^2}{2} \tag{19b}$$

$$\operatorname{rank}(\check{U}_\ell) = 1, \quad \forall \ell \in [k_i]. \tag{19c}$$

Note that the only difference between (19) and (18) is the non-convex rank constraint (19c). In particular, the costs (18a) and (19a) are identical and can be analyzed jointly.

Part 2: Analyzing the unconstrained problem: We define

$$M(\check{U}) := \frac{\gamma^2}{2} T_i \Gamma_{k_i}^{\check{U}}(\hat{\mathcal{A}}_i, \mathcal{B}_u) + \sum_{t=1}^{T} x_t x_t^\top \succeq 0 \tag{20}$$

and recall that $\lambda_{\min}(M(\check{U})) = \min_{\|v\|=1} v^{\mathsf{H}} M(\check{U}) v$. Defining $F_\ell := (e^{j\frac{2\pi\ell}{k_i}} I - \hat{\mathcal{A}}_i)^{-1}\mathcal{B}_u$ for $\ell \in [k_i]$, $c_i := \frac{T_i}{2k_i^2}$ and using elementary reformulations we obtain

$$\begin{aligned}&\lambda_{\min}(M(\check{U})) \\ &\quad = \min_{\|v\|=1} c_i v^{\mathsf{H}}\Big(\sum_{\ell\in[k_i]} F_\ell \check{U}_\ell F_\ell^{\mathsf{H}}\Big) v + v^{\mathsf{H}} \sum_{t=1}^{T} x_t x_t^\top v.\end{aligned} \tag{21}$$

With this we can rewrite (19a) as

$$\begin{aligned}\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in \mathbb{S}_+^{d_u}} \min_{\|v\|=1} \Bigg[& c v^{\mathsf{H}}\Big(\sum_{\ell\in[k_i]} F_\ell \check{U}_\ell F_\ell^{\mathsf{H}}\Big) v \\ & + v^{\mathsf{H}} \sum_{t=1}^{T} x_t x_t^\top v\Bigg].\end{aligned} \tag{22}$$

Let v_* be the vector that achieves the minimum in (22). Note that for a fixed v_*, the maximum in (22) is independent of the second term of the sum in (22). Defining $H_\ell := F_\ell^{\mathsf{H}} v_* v_*^{\mathsf{H}} F_\ell$, the optimal solution to (22) is identical to the solution to

$$\begin{aligned}\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in \mathbb{S}_+^{d_u}} & \sum_{\ell\in[k_i]} v_*^{\mathsf{H}} F_\ell \check{U}_\ell F_\ell^{\mathsf{H}} v_* \\ & = \max_{\{\check{U}_\ell^*\}_{\ell\in[k_i]}\in \mathbb{S}_+^{d_u}} \sum_{\ell\in[k_i]} \operatorname{Tr}\left(H_\ell \check{U}_\ell\right),\end{aligned} \tag{23}$$

where the equality follows from the cyclic property of the trace. Note that each matrix $H_\ell \succeq 0$, $\ell \in [k_i]$, admits an eigendecomposition $H_\ell = V_\ell \Lambda_\ell V_\ell^{\mathsf{H}}$ with $V_\ell V_\ell^{\mathsf{H}} = I$. With

$$\mathbb{P}\left[\|\hat{\theta}_T-\theta_*\|\leq C\frac{\sqrt{\log\frac{1}{\delta}+n_x+\log\det\left(\bar{\Gamma}_T\left(\Gamma_{k(T)}(\mathcal{A}^*,\mathcal{B}_w^*)+\frac{\gamma^2}{p}\Gamma_{k(T)}(\mathcal{A}^*,\mathcal{B}_u^*)\right)^{-1}+I_{n_x}\right)}}{\sqrt{T\lambda_{\min}\left(\Gamma_{k(T)}(\mathcal{A}^*,\mathcal{B}_w^*)+\gamma^2\Gamma^{\check{u}^*}_{k(T)}(\mathcal{A}^*,\mathcal{B}_u^*)\right)}}\right]\geq 1-\delta \tag{14}$$

this, it holds that

$$\begin{aligned}&\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in\mathbb{S}_+^{d_u}}\sum_{\ell\in[k_i]}\operatorname{Tr}\left(H_\ell\check{U}_\ell\right)\\&\quad=\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in\mathbb{S}_+^{d_u}}\sum_{\ell\in[k_i]}\operatorname{Tr}\left(V_\ell\Lambda_\ell V_\ell^{\mathsf{H}}\check{U}_\ell\right).\end{aligned}\tag{24}$$

Part 3: Analytical solution of (19)*:* By the constraint (19c) there exist vectors u_ℓ such that $\check{U}_\ell=u_\ell u_\ell^{\mathsf{H}}$ for all $\ell\in[k_i]$, so that we can rewrite (24) as

$$\begin{aligned}&\max_{\{\check{U}_\ell\}_{\ell\in[k_i]}\in\mathbb{S}_+^{d_u}}\sum_{\ell\in[k_i]}\operatorname{Tr}\left(V_\ell\Lambda_\ell V_\ell^{\mathsf{H}}\check{U}_\ell\right)\\&\quad=\max_{\{u_\ell\}_{\ell\in[k_i]}}\sum_{\ell\in[k_i]}u_\ell^{\mathsf{H}}V_\ell\Lambda_\ell V_\ell^{\mathsf{H}}u_\ell.\end{aligned}\tag{25}$$

Applying the change of variables $z_\ell=V_\ell^{\mathsf{H}}u_\ell$ we equivalently rewrite (25) as

$$\max_{\{z_\ell\}_{\ell\in[k_i]}}\sum_{\ell\in[k_i]}z_\ell^{\mathsf{H}}\Lambda_\ell z_\ell=\max_{\{z_\ell\}_{\ell\in[k_i]}}\sum_{\ell\in[k_i]}\sum_{j=1}^{n_u}[z_\ell]_j[z_\ell^{\mathsf{H}}]_j[\Lambda_\ell]_{jj}.\tag{26}$$

Since the sum is separable, the optimal strategy under the energy constraint (19b) is to allocate all the weight to the largest eigenvalue across all Λ_ℓ, $\ell\in[k_i]$. Note that if the largest eigenvalue has multiplicity $m>1$, there exist multiple optimal solutions that distribute the energy across the largest eigenvalues. *Part 4: Analytical solution of* (18)*:* Using the cyclic property of the trace and the change of variables $Z_\ell=V_\ell^{\mathsf{H}}\check{U}_\ell V_\ell,\forall\ell\in[k_i]$, we can equivalently rewrite (24) as

$$\begin{aligned}&\max_{\{Z_\ell\}_{\ell\in[k_i]}\in\mathbb{S}_+^{n_u}}\sum_{\ell\in[k_i]}\operatorname{Tr}(H_\ell V_\ell Z_\ell V_\ell^{\mathsf{H}})\\&\quad=\max_{\{Z_\ell\}_{\ell\in[k_i]}\in\mathbb{S}_+^{n_u}}\sum_{\ell\in[k_i]}\sum_{j=1}^{n_u}\left([\Lambda_\ell]_{jj}[Z_\ell]_{jj}\right),\end{aligned}\tag{27}$$

where the second equality holds because Λ_ℓ is diagonal. Under the energy constraint (18b), the optimal allocation strategy remains to allocate all weight to the largest eigenvalue across all Λ_ℓ with $\ell\in[k_i]$, or if multiple largest eigenvalues exist, to distribute the weight among them.If the latter is true, multiple optimal solutions exist. Thus, the optimal solutions Z_ℓ^* for $\ell\in[k_i]$ are diagonal matrices with rank from 1 to at most m, where m is the number of the largest eigenvalues across all Λ_ℓ for $\ell\in[k_i]$. It is important to emphasize that given v_* a rank-1 solution always exists. *Part 5: Conclusion:* Let z_ℓ^* and Z_ℓ^* be the optimal solutions to (26) and (27), respectively. Specially, let Z_ℓ^1 denote the optimal rank-1 solution to (27). Then, we have

$$\sum_{\ell\in[k_i]}z_\ell^{*\mathsf{H}}\Lambda_\ell z_\ell^*=\sum_{\ell\in[k_i]}\operatorname{Tr}\left(\Lambda_\ell Z_\ell^1\right)=\sum_{\ell\in[k_i]}\operatorname{Tr}\left(\Lambda_\ell Z_\ell^*\right).$$

Applying the inverse change of variables $u_\ell^*=V_\ell z_\ell$, $\check{U}_\ell^*=V_\ell Z_\ell^* V_\ell^{\mathsf{H}}$ and $\check{U}_\ell^1=V_\ell Z_\ell^1 V_\ell^{\mathsf{H}}$, we have

$$\sum_{\ell\in[k_i]}u_\ell^{*\mathsf{H}}H_\ell u_\ell^*=\sum_{\ell\in[k_i]}\operatorname{Tr}\left(H_\ell\check{U}_\ell^1\right)=\sum_{\ell\in[k_i]}\operatorname{Tr}\left(H_\ell\check{U}_\ell^*\right).$$

As established in Part 4, any solution $\{\check{U}_\ell^*\}_{\ell\in[k_i]}$ of the convex problem (18) consists of matrices with rank at most m. Take an orthogonal eigendecomposition of $\{\check{U}_\ell^*\}_{\ell\in[k_i]}$ and denote any eigenvector corresponding to a nonzero eigenvalue by $\{\bar{u}_\ell^*\}_{\ell\in[k_i]}$. Then the optimal solution $\sqrt{\operatorname{Tr}(\check{U}_\ell^*)}\bar{u}_\ell^*$, $\forall\ell\in[k_i]$ achieves the same objective as in (19). ■

Theorem 2 shows that the non-convex optimization problem (11) used in line 6 of Algorithm 1 has a hidden convexity property that we leverage to obtain an exact convex reformulation. It is worth observing that Theorem 2 also applies to fully observed state space systems, hence, the convex reformulation can also be used for the algorithm presented in [1], yielding computationally attractive input design algorithms for both setups.

IV. Numerical Example

To showcase the effectiveness of the proposed algorithm, we compare its estimation error with the error resulting from purely isotropic Gaussian excitations and the approach proposed in [17], which also relies on solving a convex problem, but is not guaranteed to recover the solution of (11). The key difference of our approach compared to [17, Algorithm 3], which also proposes a convex input design, is that the latter sequentially applies all eigenvectors of the solutions of the convex problem (18) instead of just the largest eigenvector. As shown in [17, Proposition 6.5], this reproduces any optimal solution to (18) as $T\to\infty$, i.e., only as the number of samples goes to infinity. We consider the following ARX system

$$A_1^*=\begin{bmatrix}0.7&0.1\\0&0.9\end{bmatrix},\ A_2^*=\begin{bmatrix}-0.5&0\\0.1&-0.2\end{bmatrix},\ B_1^*=\begin{bmatrix}0.1&0\\0&5\end{bmatrix},$$

which satisfies Assumption 1. Further, we fix $\Sigma_w=I_{n_x}$ and $\gamma=10$ and select the hyperparameters $T_0=200$ and $k_0=10$. We initially collect 50 datapoints with random excitations and then compare the average estimation error over 200 runs resulting from the following excitations: [1]

1) $u_t\overset{\text{i.i.d.}}{\sim}\mathcal{N}(0,\frac{\gamma^2}{n_u}I_{n_u})$.
2) u_t computed by Algorithm 1 with [17, Algorithm 3]
3) u_t computed by Algorithm 1 with OptInput $=\mathcal{F}^{-1}\{\check{u}^*\}$ and $\check{u}^*$ as in Theorem 2.

[1]The Python code for the numerical example can be accessed at: https://github.com/col-tasas/2025-hidden-convexity-active-learning

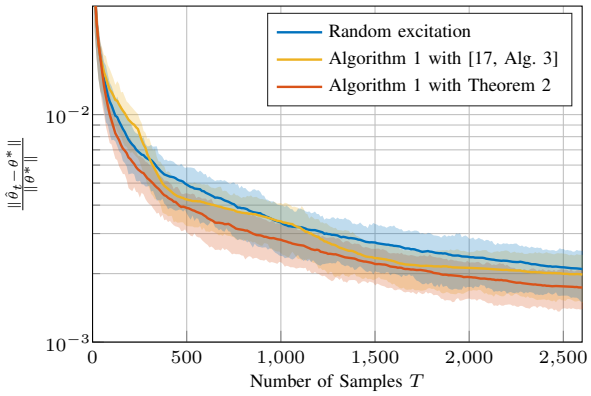

Fig. 1. Comparison of the average estimation error over 200 runs for different excitations. Shaded regions indicate the 25% and 75% percentiles.

Approaches 2) and 3) require solving problem (18), for which we use CVXPY [22], [23] with the solver MOSEK [24]. As shown in Figure 1, our proposed input design algorithm outperforms both isotropic random excitations and the approach proposed in [17]. Specifically, achieving the same accuracy with random excitations requires over 50% more samples than with our algorithm. Note that in the approach presented in [17], the input sequence associated with each eigenvector is applied for T_i time steps. This results in an episode length that scales with n_u. Crucially, our approach maintains the original feedback frequency, computing new optimal excitations every T_i rather than every $n_u T_i$ samples. This is beneficial since the estimate $\hat{\theta}_i$ is used to derive the optimal excitation instead of the true system matrix θ^*, and further, the true system is affected by noise. On the other hand, when the model used to predict the future evolution of the covariates is exact and there is no noise, our approach will result in the same input sequences as [17, Algorithm 3]. We do not provide a comparison with an input obtained by solving with numerical methods (11). This is because solving the non-convex problem reliably required significantly more run-time and proved to be very sensitive to the provided initial guess, which cannot be easily obtained in practice.

V. Conclusion

We present an input design algorithm that accelerates identification of ARX systems affected by Gaussian Process noise and provide a finite sample identification error bound. Furthermore, we propose a convex reformulation of the originally non-convex input design problem, which enables finding the optimal excitation by solving a convex optimization problem. Notably, this convex reformulation can also be applied to fully observed state-space systems and hence is of independent interest. With this, we obtain a computationally attractive input design algorithm with finite sample guarantees. A concluding example showcases the effectiveness of our approach. Interesting future research directions include the development of finite-sample versions of other classical optimality criteria to complement the existing literature.

References

[1] A. Wagenmaker and K. Jamieson, "Active learning for identification of linear dynamical systems," in *Conf. on Learning Theory*. PMLR, 2020, pp. 3487–3582.

[2] X. Bombois, M. Gevers, R. Hildebrand, and G. Solari, "Optimal experiment design for open and closed-loop system identification," *Commun. in Information and Systems*, vol. 11, no. 3, pp. 197–224, 2011.

[3] L. Gerencsér, H. Hjalmarsson, and J. Mårtensson, "Identification of arx systems with non-stationary inputs—asymptotic analysis with application to adaptive input design," *Automatica*, vol. 45, no. 3, pp. 623–633, 2009.

[4] I. R. Manchester, "Input design for system identification via convex relaxation," in *49th IEEE Conf. on Decision and Control (CDC)*, 2010.

[5] I. Ziemann, A. Tsiamis, B. Lee, Y. Jedra, N. Matni, and G. J. Pappas, "A tutorial on the non-asymptotic theory of system identification," in *62nd IEEE Conf. on Decision and Control (CDC)*, 2023.

[6] G. Goodwin and R. Payne, *Dynamic System Identification: Experiment Design and Data Analysis*, ser. Mathematics in Science and Engineering. Academic Press, 1977.

[7] C. R. Rojas, J. S. Welsh, G. C. Goodwin, and A. Feuer, "Robust optimal experiment design for system identification," *Automatica*, vol. 43, no. 6, pp. 993–1008, 2007.

[8] L. Gerencsér and H. Hjalmarsson, "Adaptive input design in system identification," in *44th IEEE Conf. on Decision and Control (CDC)*, 2005.

[9] M. J. Wainwright, *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge university press, 2019, vol. 48.

[10] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *Conf. On Learning Theory*. PMLR, 2018, pp. 439–473.

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

[12] Y. Sattar, S. Oymak, and N. Ozay, "Finite sample identification of bilinear dynamical systems," in *61st IEEE Conf. on Decision and Control (CDC)*, 2022.

[13] N. Chatzikiriakos, R. Strässer, F. Allgöwer, and A. Iannelli, "End-to-end guarantees for indirect data-driven control of bilinear systems with finite stochastic data," arXiv:2409.18010, 2024.

[14] D. Foster, T. Sarkar, and A. Rakhlin, "Learning nonlinear dynamical systems from a single trajectory," in *Proc. of the 2nd Conf. on Learning for Dynamics and Control*. PMLR, 2020.

[15] Y. Sattar and S. Oymak, "Non-asymptotic and accurate learning of nonlinear dynamical systems," *J. of Machine Learning Research*, vol. 23, no. 140, pp. 1–49, 2022.

[16] S. Oymak and N. Ozay, "Non-asymptotic identification of LTI systems from a single trajectory," in *2019 Amer. Control Conf. (ACC)*. IEEE, 2019, pp. 5655–5661.

[17] A. Wagenmaker, M. Simchowitz, and K. Jamieson, "Task-optimal exploration in linear dynamical systems," in *Int. Conf. on Machine Learning*. PMLR, 2021, pp. 10 641–10 652.

[18] H. Mania, M. I. Jordan, and B. Recht, "Active learning for nonlinear system identification with guarantees," *J. of Machine Learning Research*, vol. 23, no. 32, pp. 1–30, 2022.

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

[20] T. Sarkar and A. Rakhlin, "Near optimal finite time identification of arbitrary linear dynamical systems," in *Int. Conf. on Machine Learning*. PMLR, 2019, pp. 5610–5618.

[21] Z. Karnin, T. Koren, and O. Somekh, "Almost optimal exploration in multi-armed bandits," in *Int. Conf. on Machine Learning*. PMLR, 2013, pp. 1238–1246.

[22] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[23] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," *Journal of Control and Decision*, vol. 5, no. 1, pp. 42–60, 2018.

[24] MOSEK ApS, *MOSEK Optimizer API for Python 11.0*, 2025. [Online]. Available: https://docs.mosek.com/11.0/pythonapi.pdf