

Data-Driven Covariance Steering Control Design

Joshua Pilipovsky*

Panagiotis Tsiotras†

Abstract—This paper studies the problem of steering the distribution of a linear time-invariant system from an initial normal distribution to a terminal normal distribution under *no* knowledge of the system dynamics. This data-driven control framework uses data collected from the input and the state and utilizes the seminal work by Willems et al. to construct a data-based parametrization of the mean and the covariance control problems. These problems are then solved to optimality as convex programs using standard techniques from the covariance control literature. We also discuss the equivalence of indirect and direct data-driven covariance steering designs, as well as a regularized version of the problem that provides a balance between the two. We illustrate the proposed framework through a set of randomized trials. We then analyze the robustness properties of the data-free and data-driven covariance steering methods and demonstrate the trade-offs between performance and optimality among these methods in the presence of data corrupted with exogenous noise.

I. INTRODUCTION

Recently, there has been an emergence of an increasing reliance on data-driven methods for solving complex problems in science and engineering. The field of artificial intelligence has demonstrated the ability to solve extremely difficult problems using input and output data with the machinery of neural networks and learning-based algorithms [1], [2]. One of the major fundamental flaws of purely learning-based solutions is their lack of verifiability, that is, verifying that the networks will perform as expected given an input data stream. Many works have begun looking at robustness properties to verify these neural networks both in the deterministic [3], [4] and the probabilistic settings [5], [6]. The problem of analyzing the stability and robustness of a general learning-based solution is still intractable, however.

In the context of control theory, it is also often the case that we do not have prior knowledge of the system dynamics. Ideally, we would like to use the data collected to perform control designs that have guaranteed performance and robustness properties, albeit in a learning/data-driven context. To this end, another paradigm for control design is to use the input and output data streams to either estimate the model of the system or directly perform controller synthesis using these data. The former technique is referred to as *indirect* design, as it performs system identification (sysID) first, followed by controller synthesis, while the latter is referred to as a *direct* design, as it bypasses the sysID phase

completely and directly generates control commands from input and output data. Furthermore, each method may be classified into either a *certainty-equivalence* (CE) or *robust* approach, depending on whether uncertainties are taken into account.

The indirect or sysID approach has been long studied in a general setting [7], [8] with methods such as subspace identification with broader applications to filtering and state estimation [9]. For optimization-based control, such as the linear quadratic regular (LQR), there are many works that use the indirect approach [10]–[12]. Similarly, the direct approach to data-driven LQR has been approached using behavioral methods [13], gradient-based methods [14], [15], and Riccati-based methods [16]. More recently, this problem has been solved using concepts from behavioral systems theory and subspace methods [17] by making use of *Willems’ Fundamental Lemma*, which characterizes *all* trajectories of an LTI system through the range space of an input/output data matrix [18]. This gives rise to a parametrization of the feedback gains as a linear combination of the collected data and allows to perform *direct* data-driven control design through semi-definite programming (SDP) [19]. To this end, the authors in [20], [21] were able to come up with a CE and regularized design, thus bridging the gap between the indirect and direct approaches.

All these optimization-based approaches to direct data-driven LQR design are done in the deterministic setting, assuming a single realization of the state trajectories from x_0 to x_f . To the best of our knowledge, this paper is the first work that looks at the data-driven control design problem where the state is a stochastic process instead of a single deterministic trajectory. Since it is still unclear how to perform data-driven designs in the context of process noise, we limit our analysis to uncertainties in the initial and final states, leaving the case of process noise for future investigation. Specifically, we assume that the initial and final values of the state must follow a normal distribution, and the objective is to steer the entire *distribution* of states from an initial to a final one. This problem is referred to in the literature as *covariance steering* (CS) and has been extensively studied in the past couple of years in the Gaussian case [22]–[25], non-Gaussian case [26], nonlinear case [27], as well as in the presence of chance-constraints on the state and control [28]–[30]. The CS problem in the unconstrained case has an elegant separation property between the mean and covariance dynamics, and both problems turn into convex programs [22].

In this work, we present a data-driven covariance control design that steers the mean dynamics through an indirect sysID approach, and steers the covariance dynamics through a direct approach, using the techniques outlined in [25] and

* J. Pilipovsky is a PhD student at the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150, USA. Email: jpilipovsky3@gatech.edu

† P. Tsiotras is the David & Lewis Chair and Professor at the School of Aerospace Engineering and the Institute for Robotics & Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332-0150, USA. Email: tsiotras@gatech.edu

[19]. For direct covariance steering the certainty-equivalence approach is used, which adds an orthogonality constraint to the resulting optimization problem. We also analyze the regularized approach, which is a hybrid of the two. We finally study the robustness of these approaches through a set of randomized simulations with noisy data and compare the various covariance steering frameworks.

II. PROBLEM STATEMENT

We consider the following discrete-time *deterministic* time-invariant system

$$x_{k+1} = Ax_k + Bu_k, \quad (1)$$

where $x \in \mathbb{R}^n, u \in \mathbb{R}^m$, with time steps $k = 0, \dots, N-1$, where N representing the finite horizon. The system matrices A and B are assumed to be unknown. The uncertainty in the system resides in the initial state x_0 , which is a random n -dimensional vector drawn from the normal distribution

$$x_0 \sim \mathcal{N}(\mu_i, \Sigma_i), \quad (2)$$

where $\mu_i \in \mathbb{R}^n$ is the initial state mean and $\Sigma_i \in \mathbb{R}^{n \times n} \succ 0$ is the initial state covariance. Thus, although the system dynamics (1) is deterministic, the uncertainty in the initial state makes the state evolution $\{x_k\}_{k=1}^N$ a random process. The objective is to steer the trajectories of (1) from the initial distribution (2) to the terminal distribution

$$x_N = x_f \sim \mathcal{N}(\mu_f, \Sigma_f), \quad (3)$$

where $\mu_f \in \mathbb{R}^n$ and $\Sigma_f \succ 0$ are the desired state mean and covariance at time N , respectively. The cost function to be minimized is

$$J(u_0, \dots, u_{N-1}) := \mathbb{E} \left[\sum_{k=0}^{N-1} x_k^\top Q_k x_k + u_k^\top R_k u_k \right], \quad (4)$$

where $Q_k \succeq 0$ and $R_k \succ 0$ for all $k = 0, \dots, N-1$.

Problem 1: Given the unknown linear system (1), find the optimal control sequence $\{u_k\}_{k=0}^{N-1}$ that minimizes the objective function (4), subject to the initial state (2) and terminal state (3).

III. PROBLEM REFORMULATION

Borrowing from the work in [25], we adopt the control policy

$$u_k = K_k(x_k - \mu_k) + v_k, \quad (5)$$

where $K_k \in \mathbb{R}^{m \times n}$ are the feedback gains that control the evolution of the covariance of the state, and $v_k \in \mathbb{R}^m$ is the feedforward term that controls the evolution of the mean of the state. Under the control law (5) it is possible to re-write Problem 1 as a convex program, which can be solved to optimality using off-the-shelf solvers. Since the state distribution remains Gaussian at all time steps, and since a normal distribution is completely characterized by its first two moments, we decompose the system dynamics (1) into the mean dynamics and the covariance dynamics.

Plugging in the control law (5) into the dynamics (1) yields the decoupled dynamics

$$\mu_{k+1} = A\mu_k + Bv_k, \quad (6a)$$

$$\Sigma_{k+1} = (A + BK_k)\Sigma_k(A + BK_k)^\top. \quad (6b)$$

In the sequel, and similar to [31], we treat the moments of the intermediate states $\{\Sigma_k, \mu_k\}_{k=1}^{N-1}$ in the steering horizon as decision variables in the resulting optimization problem.

Similar to the dynamics, the cost function can be decoupled and written in terms of the first two moments as follows

$$J = J_\mu(\mu_k, v_k) + J_\Sigma(\Sigma_k, K_k), \quad (7a)$$

$$J_\mu := \sum_{k=0}^{N-1} (\mu_k^\top Q_k \mu_k + v_k^\top R_k v_k), \quad (7b)$$

$$J_\Sigma := \sum_{k=0}^{N-1} \left(\text{tr}(Q_k \Sigma_k) + \text{tr}(R_k K_k \Sigma_k K_k^\top) \right). \quad (7c)$$

Lastly, the two boundary conditions are written as

$$\mu_0 = \mu_i, \quad \mu_N = \mu_f, \quad (8a)$$

$$\Sigma_0 = \Sigma_i, \quad \Sigma_N = \Sigma_f. \quad (8b)$$

Problem 1 is now recast as the following two sub-problems.

Problem 2: Given the mean dynamics (6a), find the optimal mean trajectory $\{\mu_k\}_{k=1}^{N-1}$ and feedforward control $\{v_k\}_{k=0}^{N-1}$ that minimize the mean cost (7b) subject to the boundary conditions (8a).

Problem 3: Given the covariance dynamics (6b), find the optimal covariance trajectory $\{\Sigma_k\}_{k=1}^{N-1}$ and feedback gains $\{K_k\}_{k=0}^{N-1}$ that minimize the covariance cost (7c) subject to the boundary conditions (8b).

Remark 1: Problem 2 is a standard quadratic program with linear constraints that can be solved analytically given knowledge of the system matrices [22]. As such, we will perform an indirect design by first estimating the A and B matrices to solve this problem in a data-driven fashion. Problem 3, however, is a non-linear program due to the cost term $\text{tr}(RK_k \Sigma_k K_k^\top)$ and the covariance dynamics.

In the following section, we review the main concepts from behavioral systems theory that will allow us to parametrize the decision variables in Problems 2 and 3 in terms of input and output data streams.

IV. DATA-DRIVEN PARAMETERIZATION

In order to incorporate data into the problem formulation, we use the concept of persistence of excitation, along with Willems' Fundamental Lemma [18] to parametrize the feedforward and feedback gains of the control policy. First, recall the following definitions.

Definition 1: Given a signal $\{z_k\}_{k=0}^{T-1}$ where $z_k \in \mathbb{R}^\sigma$, we denote its Hankel matrix by

$$Z_{i,\ell,j} := \begin{bmatrix} z_i & z_{i+1} & \dots & z_{i+j-1} \\ z_{i+1} & z_{i+2} & \dots & z_{i+j} \\ \vdots & \vdots & \ddots & \vdots \\ z_{i+\ell-1} & z_{i+\ell} & \dots & z_{i+\ell+j-2} \end{bmatrix} \in \mathbb{R}^{\sigma \ell \times j}, \quad (9)$$

where $i \in \mathbb{Z}$ and $\ell, j \in \mathbb{N}$. For shorthand notation, if $\ell = 1$, we denote the Hankel matrix by

$$Z_{i,1,j} \equiv Z_{i,j} := [z_i \ z_{i+1} \ \dots \ z_{i+j-1}]. \quad (10)$$

Definition 2: The signal $\{z_k\}_{k=0}^{T-1} : [0, T-1] \cap \mathbb{Z} \rightarrow \mathbb{R}^\sigma$ is *persistently exciting* of order ℓ if the matrix $Z_{0,\ell,j}$ with $j = T - \ell + 1$ has rank $\sigma\ell$.

Corollary 1: In order for a signal to be persistently exciting of order ℓ , it must be sufficiently long, i.e., it must hold that $T \geq (\sigma + 1)\ell - 1$.

Suppose we carry out an experiment of duration $T \in \mathbb{N}$ where we collect input and state data $\{u_k^{(d)}\}_{k=0}^{T-1}$ and $\{x_k^{(d)}\}_{k=0}^T$, respectively. Let the corresponding Hankel matrices for the input sequence, state sequence, and shifted state sequence (with $\ell = 1$) be

$$U_{0,T} := [u_0^{(d)} \ u_1^{(d)} \ \dots \ u_{T-1}^{(d)}], \quad (11a)$$

$$X_{0,T} := [x_0^{(d)} \ x_1^{(d)} \ \dots \ x_{T-1}^{(d)}], \quad (11b)$$

$$X_{1,T} := [x_1^{(d)} \ x_2^{(d)} \ \dots \ x_T^{(d)}]. \quad (11c)$$

The next result characterizes the rank of the stacked Hankel matrices of the input and output data, and is central to the method used to formulate a tractable data-driven covariance steering problem.

Lemma 1 (Willems' Fundamental Lemma [18]):

Suppose that system (1) is controllable. If the input signal $\{u_k\}_{k=0}^{T-1}$ is persistently exciting of order $n + 1$, then

$$\text{rank} \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} = n + m. \quad (12)$$

Remark 2: In order to ensure that the input $u_k \in \mathbb{R}^m$ is persistently exciting of order $n + 1$ to satisfy Willems' Fundamental Lemma, it is sufficient to check that $T \geq (m + 1)n + m$. In practice, this can always be achieved in real-time during data collection.

Lemma 1 implies that *any* arbitrary input-state sequence can be expressed as a linear combination of the collected input-state data. Furthermore, this can be extended [13] to parameterizing any arbitrary feedback interconnection as well. In the following section, based on the work in [19], we parameterize the feedback gains in terms of the input-state data and reformulate the covariance steering problem as a semi-definite program (SDP).

A. Direct Data-Driven Covariance Steering

Assuming the signal $\{u_k\}_{k=0}^{T-1}$ is persistently exciting of order $n + 1$, we can express the feedback gains as follows

$$\begin{bmatrix} K_k \\ I_n \end{bmatrix} = \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} G_k, \quad (13)$$

where $G_k \in \mathbb{R}^{T \times n}$ are newly defined decision variables that provide the link between the feedback gains and the input-state data. Furthermore, we can re-write the covariance dynamics (6b) as

$$\begin{aligned} \Sigma_{k+1} &= [B \ A] \begin{bmatrix} K_k \\ I_n \end{bmatrix} \Sigma_k \begin{bmatrix} K_k \\ I_n \end{bmatrix}^\top [B \ A]^\top \\ &= X_{1,T} G_k \Sigma_k G_k^\top X_{1,T}^\top, \end{aligned} \quad (14)$$

where we use the fact that $X_{1,T} = AX_{0,T} + BU_{0,T}$. Similarly, the covariance cost (7c) can be re-written as

$$J_{\Sigma,k} = \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} G_k \Sigma_k G_k^\top U_{0,T}^\top). \quad (15)$$

To remedy the nonlinearity $G_k \Sigma_k G_k^\top$ in the covariance dynamics and the cost, and assuming that $\Sigma_k \succ 0$, define the new decision variables $S_k := G_k \Sigma_k \in \mathbb{R}^{T \times n}$, which yields the covariance dynamics

$$\Sigma_{k+1} = X_{1,T} S_k \Sigma_k^{-1} S_k^\top X_{1,T}^\top, \quad (16)$$

and the covariance cost

$$J_{\Sigma,k} = \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} S_k \Sigma_k^{-1} S_k^\top U_{0,T}^\top). \quad (17)$$

This problem is still non-convex due to the nonlinear term $S_k \Sigma_k^{-1} S_k^\top$. To this end, let us relax the covariance dynamics by defining a new decision variable $Y_k \succeq S_k \Sigma_k^{-1} S_k^\top$, which yields the relaxed optimization problem

$$\min_{\Sigma_k, S_k, Y_k} \bar{J}_\Sigma = \sum_{k=0}^{N-1} \left(\text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} Y_k U_{0,T}^\top) \right), \quad (18a)$$

such that, for all $k = 0, \dots, N - 1$,

$$C_k := S_k \Sigma_k^{-1} S_k^\top - Y_k \preceq 0, \quad (18b)$$

$$G_k^{(1)} := X_{1,T} Y_k X_{1,T}^\top - \Sigma_{k+1} = 0, \quad (18c)$$

$$G_k^{(2)} := \Sigma_k - X_{0,T} S_k = 0, \quad (18d)$$

with the boundary conditions (8b). The last equality constraint (18d) comes from the second block in (13) by multiplying Σ_k on the right. The relaxed problem (18) is convex, since the constraint (18b) can be written using the Schur complement as the linear matrix inequality (LMI)

$$\begin{bmatrix} \Sigma_k & S_k^\top \\ S_k & Y_k \end{bmatrix} \succeq 0. \quad (19)$$

The equality constraint (18c) and cost (18a), on the other hand, are simply linear in all the decision variables, and hence convex.

B. Indirect Data-Driven Mean Steering

Given the mean dynamics (6a) in terms of the open-loop control v_k , Lemma 1 also provides a system identification type of result using the following theorem.

Theorem 1: Suppose the signal u_k is persistently exciting of order $n + 1$. Then, the system (6a) has the following equivalent representation

$$\mu_{k+1} = X_{1,T} \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger \begin{bmatrix} v_k \\ \mu_k \end{bmatrix}. \quad (20)$$

Proof: See [13] for details. \blacksquare

Remark 3: Theorem 1 gives a data-based open-loop representation of a (noise-less) linear system. One may equivalently interpret equation (20) as the solution to the least-squares problem

$$\min_{B,A} \left\| X_{1,T} - [B \ A] \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} \right\|_F, \quad (21)$$

where $\|\cdot\|_F$ is the Frobenius norm. Thus, equation (20) provides the solution for the system matrices that best approximates the system dynamics. Using Theorem 1, we can express the mean steering problem as the following convex problem

$$\min_{\mu_k, v_k} J_\mu = \sum_{k=0}^{N-1} (\mu_k^\top Q_k \mu_k + v_k^\top R_k v_k), \quad (22a)$$

such that, for all $k = 0, \dots, N-1$,

$$H_k^{(1)} := F_\mu \mu_k + F_v v_k - \mu_{k+1} = 0, \quad (22b)$$

with the boundary conditions (8a), where $F_\mu \in \mathbb{R}^{n \times n}$ and $F_v \in \mathbb{R}^{n \times m}$ result from the partition of

$$F := X_{1,T} \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger = \begin{bmatrix} F_v \\ F_\mu \end{bmatrix} \in \mathbb{R}^{n \times (m+n)}. \quad (23)$$

V. CERTAINTY EQUIVALENCE AND REGULARIZED DATA-DRIVEN METHODS

In this section, we establish the link between the direct CS design in Section IV-A and the indirect design, as well as briefly outline a regularized design based on [21] that trade-offs the two frameworks.

A. Certainty-Equivalence Design

For notational simplicity, let the matrix

$$W_0 := \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}.$$

In the direct data-driven covariance control design, the set of optimal solutions G_k^* to (18) coincides with the set of solutions to (13), that is,

$$\{G_k^* : (\Sigma_k^*, S_k^*, Y_k^*) \in \text{argmin}(18)\} = W_0^\dagger \begin{bmatrix} K_k^* \\ I_n \end{bmatrix} + G_{\text{hom}}, \quad (24)$$

where G_{hom} is any matrix in the null space of W_0 . Let $\Pi_{W_0} := I_T - W_0^\dagger W_0$ be the orthogonal projection on the nullspace of W_0 . It has been shown in [20] that by introducing the extra *orthogonality* constraint $\Pi_{W_0} G_k = 0$, for all $k = 1, \dots, N$, results in an optimization problem (18) that is *equivalent* to the corresponding indirect design. Since S_k is a decision variable in (18), this amounts to adding the additional equality constraints $\Pi_{W_0} S_k = 0$ to (18). In the context of covariance steering, the indirect design is equivalent to the *bi-level* program

$$\min_{\Sigma_k, U_k Y_k} \bar{J}_\Sigma = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k Y_k), \quad (25a)$$

such that, for all $k = 0, \dots, N-1$,

$$C_k^{(1)} := \begin{bmatrix} \Sigma_k & U_k^\top \\ U_k & Y_k \end{bmatrix} \succeq 0, \quad (25b)$$

$$C_k^{(2)} := \begin{bmatrix} \Sigma_k & U_k^\top \hat{B}^\top \\ \hat{B} U_k & \hat{\Gamma}_k \end{bmatrix} \succeq 0, \quad (25c)$$

$$[\hat{B} \ \hat{A}] = \text{argmin}_{B, A} \|X_1 - [B \ A] W_0\|_F, \quad (25d)$$

where $\Gamma_k := \Sigma_{k+1} - \hat{A} \Sigma_k \hat{A}^\top - \hat{A} U_k^\top \hat{B}^\top - \hat{B} U_k \hat{A}^\top$. See [25] for details on this derivation.

B. Regularized Design

By adding the constraint $\Pi_{W_0} S_k = 0$ to the objective function, we arrive at a regularized direct data-driven covariance steering formulation. Letting $\lambda \geq 0$ be a tunable hyperparameter that balances indirect with direct designs, the regularized problem becomes

$$\min_{\Sigma_k, S_k Y_k} \bar{J}_\Sigma = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} Y_k U_{0,T}^\top) + \lambda \|\Pi_{W_0} S_k\|, \quad (26a)$$

such that, for all $k = 0, \dots, N-1$,

$$C_k := S_k \Sigma_k^{-1} S_k^\top - Y_k \preceq 0, \quad (26b)$$

$$G_k^{(1)} := X_{1,T} Y_k X_{1,T}^\top - \Sigma_{k+1} = 0, \quad (26c)$$

$$G_k^{(2)} := \Sigma_k - X_{0,T} S_k = 0. \quad (26d)$$

It can be shown [21] that for λ sufficiently large, the regularized design (26) coincides with the certainty-equivalence design. This regularized design is useful for balancing the robustness of the indirect design with the performance of the direct design, as we shall see in Section VI.

VI. NUMERICAL EXAMPLE

To illustrate the proposed data-driven method, we run a set of 100 trials on the double integrator system

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad (27)$$

with initial distribution $\mu_0 = [20 \ -2]^\top$, $\Sigma_0 = \text{diag}(1, 0.5)$ and terminal distribution $\mu_f = [0 \ 8]^\top$, $\Sigma_f = 0.5I_2$. The state and control cost weights are $Q_k = 0.01I_2$ and $R_k = 1$, for all $k = 0, 1, \dots, N$, respectively. We pick a control horizon $N = 10$ and data collection horizon $T = 15$ to ensure that $T \geq (m+1)n + m = 5$. The data is generated for every trial by randomly sampling the initial state and input over the collection horizon T from a standard normal distribution. The following set of simulations was run on a 32 GB Intel i7-10850H @ 2.60 GHz computer.






Fig. 1: Monte Carlo trajectories of data-driven solution.

Figure 1 shows the trajectories of the system using the data-driven framework on a set of 100 Monte Carlo runs. As



(a) Absolute difference in feedforward control.



(b) Two-norm difference in feedback control.

Fig. 2: Comparison of differences between model-based and data-driven covariance steering solutions in the noiseless case ($\beta = 0$) with mean errors (solid line) and 3σ error bars.

mentioned, the terminal covariance is indeed less than the desired one, as denoted by the solid black line. Nevertheless, the control law successfully steers the system between the two distributions with no knowledge of the system matrices.

It is also fruitful to compare the gains and feedforward control to that of the *model-based* covariance steering solution, as outlined [25]. Figure 2 shows the difference between the two solutions with the corresponding mean and 3σ errors over the set of trial runs. Figure 2 shows that the two solutions are nearly exact, with the feedforward control having an error of within 10^{-8} and the feedback control within 10^{-2} . Next, we compare the robustness properties of the various data-driven control designs as outlined in Section V. To this end, we add noise into the data collection and simulations but keep the designs as if there is no noise in the system. The question then becomes how these methods will perform when the data is corrupted by noise. For the data collection, we add an extra term Dw_k into the dynamics, where $w_k \sim \mathcal{N}(0, I_n)$ and $D = \beta I_n$, with $\beta > 0$ a tunable parameter for the noise intensity. Figure 3 shows the effect of increasing levels of noise on the accuracy of the feedforward control data-driven solution. Even for moderately large noise levels ($\beta = 10^{-2}$), the 3σ variance of the error is within 0.04. Thus, indirect mean steering design is a fruitful avenue for designing robust nominal controllers against noise. Even for moderately large noise levels ($\beta = 10^{-2}$), the 3σ variance of the error is within 0.04. Thus, indirect mean steering design is a fruitful avenue for designing robust nominal controllers against noise.




Fig. 3: Absolute value error in data-driven feedforward control for different levels of exogenous noise.

For the feedback design, we compare the direct solution as outlined in Section IV-A with the certainty-equivalence and regularized designs. To this end, we fix the noise intensity to be $\beta = 10^{-3}$ and choose the regularization hyperparameter as $\lambda = 2 \times 10^{-2}$. Figure 4 shows the mean and 3σ variance of the normed error of the corresponding data-driven methods when corrupted with noise. Interestingly, the certainty-equivalence (i.e., indirect) approach is the *most* robust among the three, achieving an error within 10^{-2} , while the direct approach is the least robust. Table I shows




Fig. 4: Norm error in data-driven feedback control for direct, regularized, and certainty-equivalence methods for a noise level $\beta = 10^{-3}$ and regularization hyperparameter $\lambda = 2 \times 10^{-2}$.

the mean and variance of the empirical error for the k th trial

$$\mathcal{E}_k := \frac{|J^* - \hat{J}_k|}{J^*}, \quad (28)$$

between the optimal model-based cost and the data-driven costs. We see that although the direct data-driven solution is the *least* robust, it is also the *most* optimal in the sense that it achieves the closest cost with the true, model-based solution. On the other hand, the CE approach has the best

TABLE I: Mean cost error among different data-driven covariance steering approaches for corrupted data.

$\times 10^4$	Direct	Regularized	CE
$\mathbb{E}[\mathcal{E}_k]$	3.9886	4.3614	12.0494
$\sigma^2[\mathcal{E}_k]$	0.0009	0.0013	0.0261

robustness but worst performance. Thus, the regularized form of the problem is a useful design for balancing robustness with optimality by varying λ .

VII. CONCLUSION

In this work, we have presented a tractable data-driven solution to the Louisville problem of steering the distribution of a deterministic linear system from one normal distribution to another. The problem was solved by decoupling the dynamics into mean and covariance dynamics and then performing an indirect data-driven design for the mean motion, and a direct data-driven design for the covariance motion through a convex relaxation. The data-driven solution matches almost exactly with its model-based counterpart. Further work in this direction will look at how to incorporate noise into the system dynamics. This problem is much harder to solve, as the introduction of random noise (a) cannot be measured in a real-time data-driven scenario, and (b) yields input-state data that can have multiple realizations.

VIII. ACKNOWLEDGMENT

This work has been supported by NASA University Leadership Initiative award 80NSSC20M0163 and ONR award N00014-18-1-2828. The article solely reflects the opinions and conclusions of its authors and not any NASA entity.

REFERENCES

- [1] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, and et al., “State-of-the-art speech recognition with sequence-to-sequence models,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4774–4778, 2018.
- [2] Y. Song, M. Steinweg, E. Kaufmann, and D. Scaramuzza, “Autonomous drone racing with deep reinforcement learning,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 1205–1212.
- [3] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, “Reluplex: An efficient SMT solver for verifying deep neural networks,” in *Computer Aided Verification*. Cham: Springer International Publishing, 2017, pp. 97–117.
- [4] M. Fazlyab, M. Morari, and G. J. Pappas, “Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming,” *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 1–15, 2022.
- [5] J. Pilipovsky, V. Sivaramakrishnan, M. M. K. Oishi, and P. Tsiotras, “Probabilistic verification of ReLU neural networks via characteristic functions,” in *Proceedings of the 5th Conference on Learning for Dynamics and Control*, 2023, accepted.
- [6] M. Fazlyab, M. Morari, and G. J. Pappas, “Probabilistic verification and reachability analysis of neural networks via semidefinite programming,” in *IEEE 58th Conference on Decision and Control (CDC)*, Nice, France, Dec 11 - 13 2019, pp. 2726–2731.
- [7] K. J. Keesman, *System Identification: An Introduction*. Springer, 2011.
- [8] L. Ljung, *System Identification: Theory for the User*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1987.
- [9] M. Verhaegen and V. Verdult, *Filtering and System Identification: A Least Squares Approach*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [10] G. Shi and R. E. Skelton, “Markov Data-Based LQG Control,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 122, no. 3, pp. 551–559, 10 1998.
- [11] H. Mania, S. Tu, and B. Recht, “Certainty equivalence is efficient for linear quadratic control,” 2019, arXiv: 1902.07826.
- [12] M. Ferizbegovic, J. Umenberger, H. Hjalmarsson, and T. B. Schön, “Learning robust lq-controllers using application oriented exploration,” *IEEE Control Systems Letters*, vol. 4, no. 1, pp. 19–24, 2020.
- [13] C. De Persis and P. Tesi, “Formulas for data-driven control: Stabilization, optimality, and robustness,” *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 909–924, 2020.
- [14] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, “Global convergence of policy gradient methods for the linear quadratic regulator,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 10–15 Jul 2018, pp. 1467–1476.
- [15] H. Mohammadi, M. Soltanolkotabi, and M. R. Jovanović, “On the linear convergence of random search for discrete-time LQR,” *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 989–994, 2021.
- [16] H. J. van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, “Data informativity: A new perspective on data-driven analysis and control,” *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4753–4768, 2020.
- [17] I. Markovsky and F. Dörfler, “Behavioral systems theory in data-driven analysis, signal processing, and control,” *Annual Reviews in Control*, vol. 52, pp. 42–64, 2021.
- [18] J. C. Willems, P. Rapisarda, I. Markovsky, and B. D. Moor, “A note on persistency of excitation,” vol. 3, Atlantis, Paradise, Bahamas, Dec 14–17 2004, pp. 2630–2631.
- [19] M. Rotulo, C. D. Persis, and P. Tesi, “Data-driven linear quadratic regulation via semidefinite programming,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 3995–4000, 2020, 21st IFAC World Congress.
- [20] F. Dörfler, P. Tesi, and C. De Persis, “On the certainty-equivalence approach to direct data-driven LQR design,” *IEEE Transactions on Automatic Control*, pp. 1–8, 2023.
- [21] —, “On the role of regularization in direct data-driven lqr control,” in *2022 61th IEEE Conference on Decision and Control (CDC)*, 2022, pp. 1091–1098.
- [22] M. Goldshtein and P. Tsiotras, “Finite-horizon covariance control of linear time-varying systems,” in *56th IEEE Conference on Decision and Control*, Melbourne, Australia, Dec 12–15 2017, pp. 3606–3611.
- [23] E. Bakolas, “Finite-horizon separation-based covariance control for discrete-time stochastic linear systems,” in *57th IEEE Conference on Decision and Control*, Miami Beach, FL, Dec 17–19, 2018, pp. 3299–3304.
- [24] K. Okamoto and P. Tsiotras, “Optimal stochastic vehicle path planning using covariance steering,” *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2276–2281, 2019.
- [25] F. Liu, G. Rapakoulias, and P. Tsiotras, “Optimal covariance steering for discrete-time linear stochastic systems,” 2023, arXiv: 2211.00618.
- [26] V. Sivaramakrishnan, J. Pilipovsky, M. Oishi, and P. Tsiotras, “Distribution steering for discrete-time linear systems with general disturbances using characteristic functions,” in *American Control Conference (ACC)*, Atlanta, Georgia, June 8 - 10 2022, pp. 4183–4190.
- [27] J. Ridderhof, K. Okamoto, and P. Tsiotras, “Nonlinear uncertainty control with iterative covariance steering,” in *58th IEEE Conference on Decision and Control*, Nice, France, Dec 11–13 2019, pp. 3484–3490.
- [28] J. Pilipovsky and P. Tsiotras, “Chance-constrained optimal covariance steering with iterative risk allocation,” in *American Control Conference*, New Orleans, LA, May 26–28 2021, pp. 2011–2016.
- [29] E. Bakolas, “Optimal covariance control for discrete-time stochastic linear systems subject to constraints,” in *55th IEEE Conference on Decision and Control*, Las Vegas, NV, Dec 12–14, 2016, pp. 1153–1158.
- [30] J. Ridderhof, J. Pilipovsky, and P. Tsiotras, “Chance-constrained covariance control for low-thrust minimum-fuel trajectory optimization,” in *AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, Aug 9–13 2020.
- [31] G. Rapakoulias and P. Tsiotras, “Discrete-time optimal covariance steering via semidefinite programming,” in *62nd IEEE Conference on Decision and Control (CDC)*, Marina Bay Sands, Singapore, Dec. 13–15, 2023.