Proceedings of Machine Learning Research vol 242:667–678, 2024

Data-Driven Robust Covariance Control for Uncertain Linear Systems

Joshua Pilipovsky JPILIPOVSKY3@GATECH.EDU
Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology

Panagiotis Tsiotras TSIOTRAS@GATECH.EDU
Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology

Abstract

The theory of covariance control and covariance steering (CS) deals with controlling the dispersion of trajectories of a dynamical system, under the implicit assumption that accurate prior knowledge of the system being controlled is available. In this work, we consider the problem of steering the distribution of a discrete-time, linear system subject to exogenous disturbances under an *unknown* dynamics model. Leveraging concepts from behavioral systems theory, the trajectories of this unknown, noisy system may be (approximately) represented using system data collected through experimentation. Using this fact, we formulate a direct data-driven covariance control problem using input-state data. We then propose a maximum likelihood uncertainty quantification method to estimate and bound the noise realizations in the data collection process. Lastly, we utilize robust convex optimization techniques to solve the resulting norm-bounded uncertain convex program. We illustrate the proposed end-to-end data-driven CS algorithm on a double integrator example and showcase the efficacy and accuracy of the proposed method compared to that of model-based methods.

Keywords: Data-driven control, distributional control, uncertainty quantification, system identification, robust convex optimization.

1. Introduction

Controlling the uncertainty is paramount for the development of safe and reliable systems. Originating from the pioneering contributions of Hotz and Skelton (1987), the theory of covariance control addresses the problem of asymptotically steering the *distribution* of a linear dynamical system from an initial to a terminal distribution when the system dynamics are corrupted by additive disturbances. When the time horizon is finite, the covariance control problem is often referred to as covariance steering (CS). This domain has evolved substantially in recent years, expanding to encompass more pragmatic scenarios. These extensions include incorporating probabilistic constraints on the state and the input (Pilipovsky and Tsiotras, 2021; Bakolas, 2016), the ability to steer more complex distributions (Sivaramakrishnan et al., 2022), and adaptations for receding horizon implementations (Saravanos et al., 2022) among many others. Covariance steering has also been successfully applied to diverse contexts such as urban air mobility (Rapakoulias and Tsiotras, 2023), vehicle path planning (Okamoto and Tsiotras, 2019), high-performance, aggressive driving (Knaup et al., 2023), spacecraft rendezvous (Renganathan et al., 2023), and interplanetary trajectory optimization (Ridderhof et al., 2020).

A common underlying assumption in all the previous methods is the availability of a model of the system dynamics, typically derived from physical first principles. Acquiring such a model often involves either direct data acquisition or a synthesis of empirical data with first principles, a process broadly categorized under system identification. These approaches, while effective, present some notable challenges. First principles modeling demands extensive domain-specific knowledge and effort to accurately represent the system. System identification offers a balance between accuracy and complexity but introduces complications in subsequent control design phases, such as determining the optimal model for control law synthesis. This *indirect* data-driven control approach leads to a complex bi-level optimization problem, involving both model identification and control design, which is generally inseparable (Markovsky et al., 2023). Consequently, research efforts have been devoted to the investigation of alternative strategies such as dual control (Feldbaum, 1963) and combined identification for control (Gevers, 2005).

Our work deviates from these traditional methodologies by delving into *direct* data-driven control design. The proposed approach bypasses the necessity for a parametric system model, instead deriving feedback control laws directly from empirical data. By applying principles from behavioral systems theory, namely *Willems' Fundamental Lemma* (Willems et al., 2004) we can effectively characterize the trajectories of a Linear Time-Invariant (LTI) system using the Hankel matrix of the input/state data (Markovsky and Dörfler, 2021; Willems et al., 2004). In cases when measurements are accurate, and in the absence of noise, this procedure exactly characterizes the system behavior, laying the groundwork for control design based on this data-driven paradigm. This method has demonstrated its effectiveness, e.g., in solving the Linear Quadratic Regulator (LQR) problem (De Persis and Tesi, 2020; Rotulo et al., 2020) without knowledge of the system matrices. It has been further adapted to situations with noisy data via regularization strategies (Dörfler et al., 2022, 2023). These developments successfully bridge the gap between *certainty-equivalence* (CE) and *robust* control paradigms. Additionally, recent advancements (Bisoffi et al., 2022; Pilipovsky and Tsiotras, 2023a) have extended this data-driven method to the design of feedback controllers that are robust to various forms of *bounded* disturbances and to uncertainties in the initial state distribution.

This paper introduces a novel approach to data-driven covariance control design in systems subjected to additive, potentially *unbounded*, Gaussian disturbances. Given that the collected data is influenced by noise, the resulting optimization program incorporates the unknown noise *realization* from the input/output dataset. We tackle this problem by first estimating the noise realization sequence and the underlying noise covariance matrix using maximum likelihood methods, while also establishing norm bounds for the estimation error. We then employ the concept of the robust counterpart of an uncertain convex program to enhance the tractability of the resulting convex optimization program. We illustrate the proposed framework on a double integrator system and compare the performance and robustness of the data-driven controller with that of a model-based controller.

2. Problem Statement

2.1. Notation

Real-valued vectors are denoted by lowercase letters, $u \in \mathbb{R}^m$, matrices are denoted by uppercase letters, $V \in \mathbb{R}^{n\times m}$, and random vectors are denoted by boldface letters, $\boldsymbol{w} \in \mathbb{R}^p$. $\chi^2_{p,q}$ denotes the inverse cumulative distribution function (CDF) of the chi-square distribution with p degrees of freedom and quantile q. The Kronocker product is denoted as $\otimes$ and the vectorization of a matrix A is denoted as $\mathrm{vec}(A) = [a_1^\intercal, \dots, a_M^\intercal]^\intercal$, where a_i is the ith column of A. Given two matrices A and B having the same number of columns, the matrix $[A; B]$ denotes the stacking of the two matrices columnwise. The set $\mathbb{N}_{[a,b]}$ with $a < b$, denotes the set of natural numbers between $a, b \in \mathbb{N}$. We denote the two-norm by $\|\cdot\|$ and the matrix Frobenius norm by $\|\cdot\|_{\mathrm{F}}$. Lastly, we denote a discrete-time signal $z_0, z_1, \dots, z_{T-1}$ by $\{z_k\}_{k=0}^{T-1}$.

2.2. Data-Driven Covariance Steering

We consider the following discrete-time, stochastic, time-invariant system

$$\boldsymbol{x}_{k+1} = A\boldsymbol{x}_k + B\boldsymbol{u}_k + D\boldsymbol{w}_k, \quad k \in \mathbb{N}_{[0,N-1]}, \tag{1}$$

where $\boldsymbol{x} \in \mathbb{R}^n, \boldsymbol{u} \in \mathbb{R}^m$, and $\boldsymbol{w} \sim \mathcal{N}(0, I_d)$ for all $k = 0, 1, \ldots, N-1$, where N represents the time horizon. We assume the noise is i.i.d. at each time step and is uncorrelated with the initial state, that is, $\mathbb{E}[\boldsymbol{w}_k \boldsymbol{w}_j^\intercal] = \mathbb{E}[\boldsymbol{x}_0 \boldsymbol{w}_k^\intercal] = 0$, for all $j \neq k$. The system matrices A, B, D are assumed to be constant, but unknown. The initial uncertainty in the system resides in the initial state $\boldsymbol{x}_0$, which is a random n-dimensional vector drawn from the normal distribution

$$\boldsymbol{x}_0 \sim \mathcal{N}(\mu_i, \Sigma_i), \tag{2}$$

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. The objective is to steer the trajectories of (1) from the initial distribution (2) to the terminal distribution

$$\boldsymbol{x}_N = \boldsymbol{x}_f \sim \mathcal{N}(\mu_f, \Sigma_f), \tag{3}$$

where μ_f and $\Sigma_f \in \mathbb{R}^{n \times n} \succ 0$ are the desired state mean and covariance at time N, respectively. Without loss of generality, we may assume that $\mu_f = 0$. The cost function to be minimized is

$$J(u_0, \ldots, u_{N-1}) \coloneqq \mathbb{E}\left[\sum_{k=0}^{N-1} \boldsymbol{x}_k^\intercal Q_k \boldsymbol{x}_k + \boldsymbol{u}_k^\intercal R_k \boldsymbol{u}_k\right], \tag{4}$$

where $Q_k \succeq 0$ and $R_k \succ 0$ for all $k = 0, \ldots, N-1$. Let $\mathbb{D} := \{x_k^{(d)}, u_k^{(d)}, x_T^{(d)}\}_{k=0}^{T-1}$ be a dataset collected from the system with an experiment over the time horizon T. In general, $N \neq T$.

Problem 1 (Data-Driven CS Controller Design) *Given the unknown linear system (1), find the optimal control sequence $\{\boldsymbol{u}_k\}_{k=0}^{N-1}$ that minimizes the objective function (4), subject to the initial (2) and terminal (3) state distributions using the dataset $\mathbb{D}$.*

2.3. Problem Reformulation

Borrowing from the work of Liu et al. (2023), we adopt the control policy

$$\boldsymbol{u}_k = K_k(\boldsymbol{x}_k - \mu_k) + v_k, \tag{5}$$

where $K_k \in \mathbb{R}^{m \times n}$ is the feedback gain that controls the covariance of the state and $v_k \in \mathbb{R}^m$ is the feed-forward term that controls 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 (Löfberg, 2004). With no constraints present, the state distribution remains Gaussian at all time steps, completely characterized by its first two moments. Consequently, we may decompose the system dynamics (1) into the mean dynamics and covariance dynamics. Substituting the control law (5) into the dynamics (1) yields the decoupled mean and covariance dynamics

$$\mu_{k+1} = A\mu_k + Bv_k, \tag{6a}$$
$$\Sigma_{k+1} = (A + BK_k)\Sigma_k(A + BK_k)^\intercal + DD^\intercal. \tag{6b}$$

In the sequel, and similar to the recent work by Rapakoulias and Tsiotras (2023), we treat the moments of the intermediate states $\{\Sigma_k, \mu_k\}_{k=0}^N$ over the steering horizon as decision variables in the resulting optimization problem.

Similar to the dynamics in (6), the cost function can be decoupled as $J = J_\mu(\mu_k, v_k) + J_\Sigma(\Sigma_k, K_k)$, where

$$J_\mu := \sum_{k=0}^{N-1} \left(\mu_k^\intercal Q_k \mu_k + v_k^\intercal R_k v_k\right), \tag{7a}$$

$$J_\Sigma := \sum_{k=0}^{N-1} \Big(\operatorname{tr}(Q_k \Sigma_k) + \operatorname{tr}(R_k K_k \Sigma_k K_k^\intercal)\Big). \tag{7b}$$

Lastly, the boundary conditions take the form

$$\mu_0 = \mu_i, \quad \mu_N = \mu_f, \tag{8a}$$

$$\Sigma_0 = \Sigma_i, \quad \Sigma_N = \Sigma_f, \tag{8b}$$

where $\Sigma_i, \Sigma_f \succ 0$. Problem 1 is now recast as the following two sub-problems.

Problem 2 (Data-Driven Mean Steering (DD-MS)) *Given the mean dynamics (6a), find the optimal mean trajectory $\{\mu_k\}_{k=0}^N$ and feed-forward control $\{v_k\}_{k=0}^{N-1}$ that minimize the mean cost (7a) subject to the boundary conditions (8a) using the dataset $\mathbb{D}$.*

Problem 3 (Data-Driven Covariance Steering (DD-CS)) *Given the covariance dynamics (6b), find the optimal covariance trajectory $\{\Sigma_k\}_{k=0}^N$ and feedback gains $\{K_k\}_{k=0}^{N-1}$ that minimize the covariance cost (7b) subject to the boundary conditions (8b) using the dataset $\mathbb{D}$.*

Note that both the mean and covariance steering problems rely on the system matrices A and B through the system dynamics (6). Thus, the two problems stated above, are not yet amenable to a data-driven solution. 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 the dataset $\mathbb{D}$.

3. Data-Driven Parameterization

We use the concept of persistence of excitation, along with Willems' Fundamental Lemma (Willems et al., 2004) to parameterize the decision variables of the control policy. First, recall the following definitions from De Persis and Tesi (2020).

Definition 1 *Given a signal $\{z_k\}_{k=0}^{T-1}$ where $z \in \mathbb{R}^\sigma$, its Hankel matrix is given by*

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

where $i \in \mathbb{N}_0$ and $\ell, j \in \mathbb{N}$. For shorthand of notation, if $\ell = 1$, we will denote the Hankel matrix by

$$Z_{i,1,j} \equiv Z_{i,j} := [z_i \; z_{i+1} \; \ldots \; z_{i+j-1}]. \tag{10}$$

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

Given the dataset $\mathbb{D}$, let the corresponding Hankel matrices for the input sequence, state sequence, and shifted state sequence (with $\ell = 1$) be $U_{0,T} := [u_0^{(d)}\ u_1^{(d)}\ \ldots\ u_{T-1}^{(d)}]$, $X_{0,T} := [x_0^{(d)}\ x_1^{(d)}\ \ldots\ x_{T-1}^{(d)}]$, and $X_{1,T} := [x_1^{(d)}\ x_2^{(d)}\ \ldots\ x_T^{(d)}]$. The next result characterizes the rank of the stacked Hankel matrices of the input and output data, and is central to our approach to formulate a tractable data-driven covariance steering problem.

Assumption 1 *We assume that the data is persistently exciting, i.e., the Hankel matrix of input/state data is full row rank*

$$\operatorname{rank}\begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} = n + m. \tag{11}$$

Remark 1 *In this work, the full rank condition in (11) is an assumption, rather than a derived condition based on notions of persistency of excitation. Indeed, when the system is purely deterministic, (11) holds if and only if the input* $\{u_k\}_{k=0}^{T-1}$ *is persistently exciting of order* $n+1$*, a result now called the fundamental Lemma Willems et al. (2004). However, in the present case, it can be shown that if in addition the disturbance realization* $W_{0,T} \triangleq [\boldsymbol{w}_0^{(d)}, \ldots, \boldsymbol{w}_{T-1}^{(d)}]$ *is persistently exciting of order* $n+1$*, then Assumption 1 holds Bisoffi et al. (2022).*

Assumption 1 implies that *any* arbitrary input-state sequence can be expressed as a linear combination of the collected input-state data. Furthermore, as shown in the next section, this idea has been used by De Persis and Tesi (2020) to parameterize arbitrary feedback interconnections as well. In the following section, based on the work of Rotulo et al. (2020), we parameterize the feedback gains in terms of the input-state data and reformulate Problem 3 as a semi-definite program (SDP).

3.1. 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

$$\begin{bmatrix} K_k \\ I_n \end{bmatrix} = \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} G_k, \quad k = 0, \ldots, N-1, \tag{12}$$

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.

Theorem 1 *Using the data-driven parameterization of the feedback gains (12), Problem 3 may be relaxed as the convex program*

$$\min_{\Sigma_k, S_k Y_k} \bar{J}_\Sigma = \sum_{k=0}^{N-1} \Big(\operatorname{tr}(Q_k \Sigma_k) + \operatorname{tr}(R_k U_{0,T} Y_k U_{0,T}^\intercal) \Big), \tag{13a}$$

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

$$C_k^{(1)} := \begin{bmatrix} \Sigma_k & S_k^\intercal \\ S_k & Y_k \end{bmatrix} \succeq 0, \tag{13b}$$

$$C_k^{(2)} := \begin{bmatrix} \Sigma_{k+1} - \Sigma_{\boldsymbol{\xi}} & (X_{1,T} - \Xi_{0,T}) S_k \\ S_k^\intercal (X_{1,T} - \Xi_{0,T})^\intercal & \Sigma_k \end{bmatrix} \succeq 0, \tag{13c}$$

$$G_k^{(1)} := \Sigma_k - X_{0,T} S_k = 0, , \quad G_k^{(2)} := \Sigma_N - \Sigma_f = 0, \tag{13d}$$

where $\Xi_{0,T} := [\boldsymbol{\xi}_0, \ldots, \boldsymbol{\xi}_{T-1}] \in \mathbb{R}^{n \times T}$ *is the Hankel matrix of the (unknown) disturbances, and* $\boldsymbol{\xi}_k \sim \mathcal{N}(0, \Sigma_{\boldsymbol{\xi}})$*, where* $\Sigma_{\boldsymbol{\xi}} := DD^\intercal$ *for all* $k = 0, \ldots, T-1$.

Proof Please see Appendix A (Pilipovsky and Tsiotras, 2023b). ■

3.2. Indirect Data-Driven Mean Steering

Given the mean dynamics (6a) in terms of the open-loop control v_k, Assumption 1 can also be used to provide a system identification type-of-result using the following theorem.

Theorem 2 *Suppose the signal* $u_k^{(d)}$ *is persistently exciting of order* $n+1$*. Then, system* (6a) *has the following equivalent representation*

$$\mu_{k+1} = (X_{1,T} - \Xi_{0,T}) \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger \begin{bmatrix} v_k \\ \mu_k \end{bmatrix}. \tag{14}$$

Proof See De Persis and Tesi (2020) for details. ■

Using Theorem 2, we can express Problem 2 as the following convex problem

$$\min_{\mu_k, v_k} J_\mu = \sum_{k=0}^{N-1} (\mu_k^\intercal Q_k \mu_k + v_k^\intercal R_k v_k), \tag{15a}$$

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

$$F_\mu(\Xi_{0,T})\mu_k + F_v(\Xi_{0,T})v_k - \mu_{k+1} = 0, \quad \mu_N - \mu_f = 0, \tag{15b}$$

where $F_\mu \in \mathbb{R}^{n \times n}$ and $F_v \in \mathbb{R}^{n \times m}$ result from the partition of

$$F := (X_{1,T} - \Xi_{0,T}) \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger = \begin{bmatrix} F_v(\Xi_{0,T}) & F_\mu(\Xi_{0,T}) \end{bmatrix}. \tag{16}$$

Given knowledge of $\Xi_{0,T}$, the solution of (15) yields an indirect design for the feed-forward control that solves the mean steering problem. In (Pilipovsky and Tsiotras, 2023a) it was shown that, in the absence of noise, the previous optimization problem can be solved as a convex program.

4. Uncertainty Quantification

4.1. Maximum Likelihood Estimation of Noise Realization

The optimization problem (13) as well as (15) depend on the unknown noise realization. In this section, we propose a maximum likelihood estimation (MLE) scheme to estimate the noise realization from the collected data. First, we encode the stochastic linear system dynamics by enforcing a *consistency* condition on the realization data. To this end, we first notice that there exists a noise realization such that the dynamics (1) satisfy

$$X_{1,T} = AX_{0,T} + BU_{0,T} + \Xi_{0,T}. \tag{17}$$

For notational convenience, define the augmented Hankel matrix $S := [U_{0,T}; X_{0,T}] \in \mathbb{R}^{(m+n) \times T}$, from which we may re-write (17) as $X_{1,T} = [B\ A]S + \Xi_{0,T}$. Additionally, noting that the pseudoinverse satisfies the property $SS^\dagger S = S$, (17) is equivalently written as $X_{1,T} - \Xi_{0,T} = [B\ A]SS^\dagger S$, which, using the relation $X_{1,T} - \Xi_{0,T} = [B\ A]S$, yields

$$(X_{1,T} - \Xi_{0,T})(I_T - S^\dagger S) = 0. \tag{18}$$

Equation (18) is a *model-free* condition that must be satisfied for all noisy linear system data realizations, and hence is a consistency relation for any feasible set of data. Given the constraint (18), the MLE problem then becomes

$$\max_{\Xi_{0,T},\Sigma_{\boldsymbol{\xi}}} \mathcal{J}_{\mathrm{MLE}}(\Xi_{0,T},\Sigma_{\boldsymbol{\xi}} \mid \mathbb{D}) \triangleq \sum_{k=0}^{T-1} \log \rho_{\boldsymbol{\xi}}(\xi_k), \quad \text{s.t.} \quad (X_{1,T} - \Xi_{0,T})(I_T - S^{\dagger}S) = 0, \tag{19}$$

where $\rho_{\boldsymbol{\xi}}(x)$ is the probability density function (PDF) of the noise random vector $\boldsymbol{\xi}$.

Theorem 3 *Assuming $\Sigma_{\boldsymbol{\xi}} \succ 0$, the MLE optimization problem (19) may be solved as the following difference-of-convex (DC) program*

$$\min_{\Xi_{0,T},\Sigma_{\boldsymbol{\xi}},U} \left[\frac{1}{2}\mathrm{tr}(U) - \left(-\frac{T}{2}\log\det\Sigma_{\boldsymbol{\xi}}\right)\right] \tag{20a}$$

$$\text{s.t.} \qquad (X_{1,T} - \Xi_{0,T})(I_T - S^{\dagger}S) = 0, \qquad \begin{bmatrix} \Sigma_{\boldsymbol{\xi}} & \Xi_{0,T} \\ \Xi_{0,T}^{\intercal} & U \end{bmatrix} \succeq 0. \tag{20b}$$

The previous DC program may be solved using a successive convexification procedure known as the convex-concave procedure (CCP), which is guaranteed to converge to a feasible point (Yuille and Rangarajan, 2001; Lipp and Boyd, 2016). Oftentimes, the disturbance matrix D is known beforehand, as is the case when one knows how the disturbances affect the system state variables. In such cases, $\Sigma_{\boldsymbol{\xi}} = DD^{\intercal}$ is no longer a decision variable, and the MLE problem (19) may be solved analytically as given by the following Corollary 1.

Corollary 1 *Suppose $\Sigma_{\boldsymbol{\xi}}$ is known. Then the MLE problem (19) has the exact solution*

$$\Xi_{0,T}^{\star} = X_{1,T}(I_T - S^{\dagger}S). \tag{21}$$

4.2. Uncertainty Error Bounds

We are interested in deriving bounds to ensure robust satisfaction (with high probability) of the terminal CS constraints (13c) for the DD-CS problem. To this end, we use the statistical properties of $\Xi_{0,T}$ and generate an ellipsoidal uncertainty set based on some degree of confidence $\delta \in [0.5, 1)$. For simplicity, assume $\Sigma_{\boldsymbol{\xi}} \succ 0$ is known. First, we re-write the MLE problem (19) in terms of the vectorized parameters to be estimated $\xi := \mathrm{vec}(\Xi_{0,T}) = [\xi_0^{\intercal}, \dots, \xi_{T-1}^{\intercal}]^{\intercal} \in \mathbb{R}^{nT}$ as

$$\min_{\xi} \mathcal{J}_{\mathrm{MLE}}(\xi \mid \mathbb{D}) = \frac{1}{2}\xi^{\intercal}(I_T \otimes \Sigma_{\boldsymbol{\xi}}^{-1})\xi \quad \text{s.t.} \quad C(\xi) := (\Gamma \otimes I_n)\xi - \lambda = 0, \tag{22}$$

where $\Gamma := I_T - S^{\dagger}S \in \mathbb{R}^{T\times T}, \Lambda := X_{1,T}\Gamma \in \mathbb{R}^{n\times T}$, and $\lambda = \mathrm{vec}(\Lambda)$. It can be shown that (Newey and McFadden, 1994), as the number of samples grows, the MLE noise estimate $\hat{\xi}$ converges to a normal distribution as $\sqrt{T}(\xi - \hat{\xi}) \xrightarrow{d} \mathcal{N}(0, \mathcal{I}^{-1})$, where $\mathcal{I} = \mathbb{E}_{\xi}\left[\frac{\partial^2}{\partial \xi^2}\mathcal{J}_{\mathrm{MLE}}(\xi \mid \mathbb{D})\right]$ is the *Fisher Information Matrix* (FIM), which is given by $\mathcal{I} = I_T \otimes \Sigma_{\boldsymbol{\xi}}^{-1}$ in the unconstrained case. For a constrained MLE problem, it can similarly be shown, as by Stoica and Ng (1998), that the asymptotic distribution of the estimates has covariance $\Sigma_{\Delta} = \mathcal{I}^{-1} - \mathcal{I}^{-1}J^{\intercal}(J\mathcal{I}^{-1}J^{\intercal})^{-1}J\mathcal{I}^{-1}$, where $J := \frac{\partial}{\partial \xi}C(\xi) = \Gamma \otimes I_n$ is the Jacobian of the constraints. We next present a few technical results that are useful for computing the uncertainty covariance of the MLE estimates and the associated confidence sets. The proofs of these results are given in the Appendix (Pilipovsky and Tsiotras, 2023b).

Lemma 1 *The error of the constrained MLE estimator* (19) *for the unknown noise realization* $\xi = \mathrm{vec}(\Xi_{0,T})$ *converges to the normal distribution* $\mathcal{N}(0, \Sigma_\Delta)$*, where* $\Sigma_\Delta = S^\dagger S \otimes \Sigma_{\boldsymbol{\xi}}$.

Proposition 1 *Assume the uncertainty error estimate is normally distributed as* $\Delta\xi \sim \mathcal{N}(0, \Sigma_\Delta)$*. Then, given some level of risk* $\delta \in [0.5, 1)$*, the uncertainty set* $\Delta := \{\|\Delta\Xi_{0,T}\| \leq \rho\}$*, contains the* $(1-\delta)$*-quantile of* $\Delta\Xi_{0,T}$*, with* $\rho = \frac{\chi_{nT,1-\delta}}{\sqrt{\lambda_{\min}(\Sigma_\Delta^{-1})}}$*, where* $\chi_{p,q}$ *is the square root of the inverse CDF of the* $\chi^2_{p,q}$ *distribution.*

Corollary 2 *For the MLE problem* (19)*, the associated* $(1-\delta)$*-quantile uncertainty set* $\Delta :=$ $\{\|\Delta\Xi_{0,T}\| \leq \rho\}$ *has the bound* $\rho = \|\Sigma_{\boldsymbol{\xi}}^{1/2}\|\chi_{nT,1-\delta}$*, where* $\Delta\Xi_{0,T} = \Xi_{0,T} - \hat{\Xi}_{0,T}$.

5. Robust DD-CS

The problem outlined in (13) is categorized as an uncertain convex program (Ben-Tal et al., 2009), where the LMI constraints (13c) are required to hold across all realizations of the uncertain parameter $\Delta\Xi_{0,T}$ within the set Δ defined in Proposition 1. This formulation results in a semi-infinite problem, making the original constraints intractable. To address this, we focus on the robust counterpart (RC) of a class of uncertain LMIs (uLMI), which simplifies the problem to a robust feasibility problem. Letting $\mathcal{A}(Z_k, \Xi_{0,T}) \succeq 0$ denote the LMI in (13c) for the decision variables $Z_k = \{S_k, \Sigma_k, \Sigma_{k+1}\}$, we can ensure tractability by enforcing the condition $\sup_{\Delta\Xi_{0,T} \in \Delta} \mathcal{A}(Z_k, \Xi_{0,T}) \succeq 0$. Adhering to this robust counterpart ensures that the original constraints are robustly met with high probability, characterized by the risk value δ. In what follows, we will form the RC of the uncertain CS problem in (13). We first state the following theorem on the equivalence of the RC of a uLMI in which the uncertainty appears linearly in the constraints.

Theorem 4 (Ben-Tal et al. (2009), Proposition 6.4.1) *The RC of the uncertain LMI*

$$\mathcal{A}(y, \Pi) := \hat{\mathcal{A}}(y) + L^\intercal(y)\Pi R + R^\intercal \Pi^\intercal L(y) \succeq 0, \tag{23}$$

with unstructured norm-bounded uncertainty set $\mathcal{Z} = \{\Pi \in \mathbb{R}^{p\times q} : \|\Pi\| \leq \rho\}$*, can be equivalently represented by the LMI*

$$\begin{bmatrix} \lambda I_p & \rho L(y) \\ \rho L^\intercal(y) & \hat{\mathcal{A}}(y) - \lambda R^\intercal R \end{bmatrix} \succeq 0, \tag{24}$$

in the decision variables y, λ.

5.1. Problem Formulation and Robust LMI Counterpart

Using (20) to obtain a noise realization estimate $\hat{\Xi}_{0,T}$ and Corollary 2 to obtain the associated uncertainty set Δ, the original covariance LMI constraints (13c), may be decomposed using $\Xi_{0,T} = \hat{\Xi}_{0,T} + \Delta\Xi_{0,T}$ as the semi-infinite uLMIs $\hat{G}_k^\Sigma + \Delta G_k^\Sigma(\Delta\Xi_{0,T}) \succeq 0$, for all $\|\Delta\Xi_{0,T}\| \leq \rho$, where,

$$\hat{G}_k^\Sigma = \begin{bmatrix} \Sigma_{k+1} - \Sigma_{\boldsymbol{\xi}} & (X_{1,T} - \hat{\Xi}_{0,T})S_k \\ S_k^\intercal (X_{1,T} - \hat{\Xi}_{0,T})^\intercal & \Sigma_k \end{bmatrix}, \tag{25}$$

is the nominal LMI, and

$$\Delta G_k^\Sigma = \begin{bmatrix} 0_n & -\Delta\Xi_{0,T} S_k \\ -S_k^\intercal \Delta\Xi_{0,T}^\intercal & 0_n \end{bmatrix} \succeq 0, \tag{26}$$

is the perturbation to the covariance LMI. Next, we represent the matrix (26) as

$$\Delta G_k^{\Sigma} = L^{\intercal}(S_k)\Delta\Xi_{0,T}^{\intercal}R + R^{\intercal}\Delta\Xi_{0,T}L(S_k), \tag{27}$$

where, $L^{\intercal}(S_k) = [0_{n,T}; -S_k^{\intercal}]$, $R^{\intercal} = [I_n; 0_n]$. Finally, using Theorem 4, we may equivalently represent the RC of the uLMI (27) as the LMI

$$\begin{bmatrix} \lambda I_T & \rho L(S_k) \\ \rho L^{\intercal}(S_k) & \hat{G}_k^{\Sigma}(Z_k, S_k) - \lambda R^{\intercal}R \end{bmatrix} \succeq 0, \tag{28}$$

in terms of the decision variables $\{S_k, \lambda\}$ and $Z_k = \{\Sigma_k, \Sigma_{k+1}\}$.

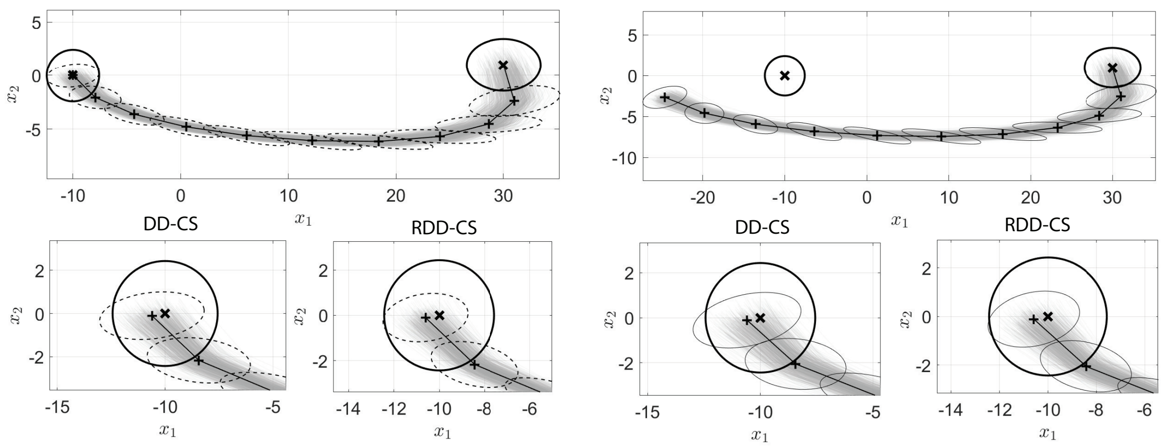

(*a*) MB-CS (top) and R/DD-CS (bottom) for nominal dynamics model $\{A, B\}$.

(*b*) MB-CS (top) and R/DD-CS (bottom) for perturbed dynamics model $\{A+\Delta A, B+\Delta B\}$.

Figure 1: Comparison of model-based and data-driven covariance steering controllers for (a) perfect system knowledge, and (b) imperfect system knowledge. The robust data-driven solutions are not only adaptable to any linear dynamics due to sampling from the true system but also achieve a desirable terminal covariance.

6. Numerical Example

To illustrate the proposed robust DD-CS (RDD-CS) framework, we consider the ground truth double integrator model

$$x_{k+1} = \begin{bmatrix} 1 & \Delta T \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ \Delta T \end{bmatrix} u_k + Dw_k, \tag{29}$$

with $\Delta T = 1$ and $D = 0.1 I_2$, with initial conditions $\mu_0 = [30, 1]^{\intercal}, \Sigma_0 = \mathrm{diag}(1, 0.5)$ and terminal conditions $\mu_f = [-10, 0]^{\intercal}, \Sigma_f = 0.5 I_2$. The planning horizon is chosen to be $N = 10$ and the data-collection horizon $T = 15$ which satisfies the persistence of excitation criterion. For the uncertainty set, we choose a risk threshold $\delta = 0.001$. Figure 1 shows the performance of the model-based and data-driven covariance controllers. In Figure 1(*a*), we see that MB-CS achieves exact uncertainty control, as expected since it has access to the *true* model. The DD-CS method with MLE noise estimation also achieves similar terminal behavior, although the vanilla DD-CS solution (bottom-left) violates the terminal covariance constraints $\Sigma_N \preceq \Sigma_f$, due to the misalignment in the mean steering from noisy data. The robust DD-CS (RDD-CS) solution (bottom-right), on the other hand, achieves a more precise terminal covariance and is fully contained within the desired terminal covariance, even with the slight mean misalignment. Specifically, the robust solution achieves a terminal covariance

$$\Sigma_N^{\mathrm{RDD-CS}} = \begin{bmatrix} 0.2612 & 0.0252 \\ 0.0252 & 0.0941 \end{bmatrix} \prec \Sigma_f.$$

In Figure 1(*b*), we illustrate the performance of the MB and DD controllers in the case where the *true* nominal dynamics model is perturbed with $\Delta A = \tau e_1 e_2^\intercal$ and $\Delta B = \tau e_2$, with $\tau = 0.05$ and e_i is the unit vector along the ith axis. In this case, the model-based design completely fails, as it takes the original model $\{A, B\}$ as the ground truth, while the data-driven design can automatically adapt to any model as it samples data $\mathbb{D}$ from the true underlying system.

Next, we include the estimation of the disturbance covariance $\Sigma_{\boldsymbol{\xi}}$ into the DD-CS algorithm. The estimate $\hat{\Sigma}_{\boldsymbol{\xi}}$ is subsequently used in the resulting RDD-CS program. The results are akin to those of Figure 1 and hence are omitted. We then perturb the disturbance matrix from the ground truth model with $\Delta D = 0.2 I_2$, representing cases in which the modeler cannot effectively quantify the intensity of the disturbance. Figure 2 shows that under this new disturbance structure, the model-based design fails as it anticipates weaker disturbances than the actual system experiences, while the data-driven solution achieves the desired terminal covariance, albeit the terminal mean state is more heavily perturbed due to the increased uncertainty. Designing a robust open-loop control to also take into account the uncertainty in the mean motion is a topic of interest for future work.

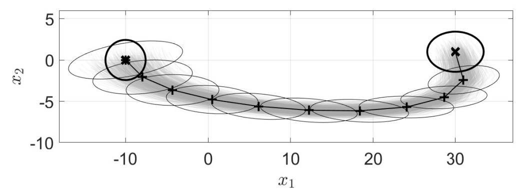

(*a*) MB-CS solution.

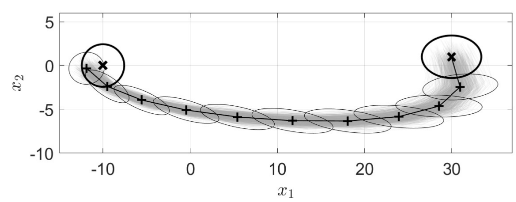

(*b*) RDD-CS solution.

Figure 2: Model-based (a) and robust data-driven (b) covariance control designs for perturbed disturbance model $D + \Delta D$.

7. Conclusion

We have presented a data-driven uncertainty control method to steer the distribution of an unknown linear dynamical system subject to Gaussian disturbances (DD-CS). Since the underlying collected data is noisy, an exact system representation is infeasible, and thus we have proposed a maximum likelihood estimation (MLE) scheme to solve for the underlying noise realization from the data as a difference-of-convex program. Using the statistical properties of MLE, we derived bounds on the $(1-\delta)$-quantile of the uncertainty estimates, which is subsequently used to solve a robust DD-CS problem. Numerical examples show that the robust DD-CS method achieves desirable levels of performance for an underlying unknown linear system compared to that of the model-based counterpart and any disturbance intensity satisfying the theoretically achievable bound (Liu et al., 2023). Future work will investigate robust solutions to the mean steering problem in the case of noisy data, as well as extensions to output-feedback systems, moving-horizon implementations, and non-linear systems.

8. 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

E. Bakolas. Optimal covariance control for discrete-time stochastic linear systems subject to constraints. In *55th IEEE Conference on Decision and Control (CDC)*, pages 1153–1158, Las Vegas, NV, Dec. 12–14, 2016.

A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski. *Robust Optimization*, volume 28 of *Princeton Series in Applied Mathematics*. Princeton University Press, 2009.

A. Bisoffi, C. De Persis, and P. Tesi. Data-driven control via Petersen's lemma. *Automatica*, 145: 110537, 2022. ISSN 0005-1098.

C. De Persis and P. Tesi. Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control*, 65(3):909–924, 2020.

F. Dörfler, P. Tesi, and C. De Persis. On the role of regularization in direct data-driven LQR control. In *61st IEEE Conference on Decision and Control (CDC)*, pages 1091–1098, Cancún, Mexico, Dec. 6–9, 2022.

F. Dörfler, P. Tesi, and C. De Persis. On the certainty-equivalence approach to direct data-driven LQR design. *IEEE Transactions on Automatic Control*, pages 1–8, 2023.

A. Feldbaum. Dual control theory problems. *IFAC Proceedings Volumes*, 1(2):541–550, 1963.

M. Gevers. Identification for control: From the early achievements to the revival of experiment design. *European Journal of Control*, 11(4):335–352, 2005.

A. F. Hotz and R. E. Skelton. Covariance control theory. *International Journal of Control*, 46(1): 13–32, 1987.

J. Knaup, K. Okamoto, and P. Tsiotras. Safe high-performance autonomous off-road driving using covariance steering stochastic model predictive control. *IEEE Transactions on Control Systems Technology*, 31(5):2066–2081, 2023.

T. Lipp and S. Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17(2):263–287, 2016.

F. Liu, G. Rapakoulias, and P. Tsiotras. Optimal covariance steering for discrete-time linear stochastic systems, 2023. arXiv: 2211.00618.

J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *IEEE International Symposium on Computer Aided Control Systems Design*, pages 284–289, Taipei, Taiwan, 2004.

I. Markovsky and F. Dörfler. Behavioral systems theory in data-driven analysis, signal processing, and control. *Annual Reviews in Control*, 52:42–64, 2021.

I. Markovsky, L. Huang, and F. Dörfler. Data-driven control based on the behavioral approach: From theory to applications in power systems. *IEEE Control Systems Magazine*, 43(5):28–68, 2023.

W. K. Newey and D. McFadden. Chapter 36 large sample estimation and hypothesis testing. In *Handbook of Econometrics*, volume 4, pages 2111–2245. Elsevier, 1994.

K. Okamoto and P. Tsiotras. Optimal stochastic vehicle path planning using covariance steering. *IEEE Robotics and Automation Letters*, 4(3):2276–2281, 2019.

J. Pilipovsky and P. Tsiotras. Covariance steering with optimal risk allocation. *IEEE Transactions on Aerospace and Electronic Systems*, 57(6):3719–3733, 2021.

J. Pilipovsky and P. Tsiotras. Data-driven covariance steering control design. In *62nd IEEE Conference on Decision and Control*, Marina Bay, Singapore, Dec. 13–15, 2023a.

J. Pilipovsky and P. Tsiotras. Data-driven robust covariance control for uncertain linear systems, 2023b. arXiv.

G. Rapakoulias and P. Tsiotras. Discrete-time optimal covariance steering via semidefinite programming. In *62nd IEEE Conference on Decision and Control*, Marina Bay, Singapore, Dec. 13–15, 2023.

V. Renganathan, J. Pilipovsky, and P. Tsiotras. Distributionally robust covariance steering with optimal risk allocation. In *American Control Conference (ACC)*, pages 2607–2614, San Diego, CA, May 31 – June 2, 2023.

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.

M. Rotulo, C. De Persis, and P. Tesi. Data-driven linear quadratic regulation via semidefinite programming. *IFAC-PapersOnLine*, 53(2):3995–4000, 2020. 21st IFAC World Congress.

A. D. Saravanos, I. M. Balci, E. Bakolas, and E. A. Teodorou. Distributed model predictive covariance steering, 2022. arXiv:2212.00398.

V. Sivaramakrishnan, J. Pilipovsky, M. M. K. Oishi, and P. Tsiotras. Distribution steering for discrete-time linear systems with general disturbances using characteristic functions. In *American Control Conference (ACC)*, pages 4183–4190, Atlanta, Georgia, June 8–10, 2022.

P. Stoica and B. C. Ng. On the Cramer-Rao bound under parametric constraints. *IEEE Signal Processing Letters*, 5(7):177–179, 1998.

J. C. Willems, P. Rapisarda, and B. Markovsky, I. andDe Moor. A note on persistency of excitation. In *43rd IEEE Conference on Decision and Control (CDC)*, volume 3, pages 2630–2631, Atlantis, Paradise, Bahamas, Dec 14-17, 2004.

A. L. Yuille and A. Rangarajan. The concave-convex procedure (CCCP). In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001.