

DUST: A Framework for Data-Driven Density Steering

Joshua Pilipovsky^{a,★} and Panagiotis Tsiotras^b

^a*D. Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA
(e-mail: jpilipovsky3@gatech.edu)*

^b*D. Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA
(e-mail: tsiotras@gatech.edu)*

Abstract

We consider the problem of data-driven stochastic optimal control of an unknown LTI dynamical system. Assuming the process noise is normally distributed, we pose the problem of steering the state's mean and covariance to a target Gaussian distribution, under noisy data collected from the underlying system, a problem commonly referred to as covariance steering (CS). A novel framework for Data-driven Uncertainty quantification and density STEering (DUST) is presented that simultaneously characterizes the noise affecting the measured data and designs an optimal affine-feedback controller to steer the density of the state to a prescribed terminal value. We use both indirect and direct data-driven design approaches based on the notions of persistency of excitation and subspace identification to exactly represent the mean and covariance dynamics of the state in terms of the data and noise realizations. Since both the mean and the covariance steering sub-problems are plagued with stochastic uncertainty arising from noisy data collection, we first estimate the noise realization from this dataset and subsequently compute tractable upper bounds on the estimation errors. The first and second moment steering problems are then solved to optimality using techniques from robust control and robust optimization. Lastly, we present an alternative control design approach based on the certainty equivalence principle and interpret the problem as one of CS under multiplicative uncertainty. We analyze the performance and efficacy of each of these data-driven approaches using a case study and compare them with their model-based counterparts.

Key words: data-driven control, stochastic optimal control, uncertainty quantification, stochastic system identification

1 Introduction

The pursuit of safe and reliable control under uncertainty stands as a fundamental challenge in control theory. Traditional model-based techniques, such as the linear-quadratic regulator (LQR) or model-predictive control (MPC) have been extensively explored, and have been shown to be extremely effective at controlling dynamical systems when the model accurately represents the actual physical system. When there are model inaccuracies due to errors during system identification, robust control techniques [1] have been used to combat these inaccuracies and ensure robust constraint satisfaction and optimality under worst-case conditions. By the same token, exogenous disturbances affecting the state of a system have been treated in a multitude of ways; when the uncertainties are bounded, the problem

falls under the realm of robust control [2]; when the uncertainties are probabilistic, the problem is treated using techniques from stochastic control [3].

Recently, there has been a paradigm shift from looking at control synthesis as an *indirect* design process of first estimating a model and subsequently solving an optimal control problem using the identified model, to a *direct* control design from raw data collected from the underlying physical system. This methodological shift has been inspired by the early works on behavioral system theory by Willems et. al. [4], which showed that one can completely characterize the trajectory space of an LTI system by solely using raw data, as long as this data is persistently exciting, a result known as the *Fundamental Lemma* (FL) [4]. This perspective has inspired a range of direct data-driven control methods (e.g., DeePC and data-driven LQR) [5–7]. In practice, however, measurement and process noise invalidate the exact FL parameterization and can lead to degraded performance [8].

★ Currently at RTX Technology Research Center (RTRC), East Hartford, CT (e-mail: joshua.y.pilipovsky@rtx.com)

This motivates the need for principled uncertainty quantification and robust data-driven synthesis for stochastic systems.

The setting of stochastic predictive control (SPC), however, is far less studied, and the problem of designing optimal controllers for *data-driven* stochastic LTI systems under noisy data is much more challenging. Existing works on data-driven stochastic control adopt rather restrictive assumptions, such as noise-free offline data [9] or exact polynomial chaos expansions of stochastic measurements [10,11]. A notable step towards a generalized framework for SPC is [12], where the authors explicitly quantify the statistical properties of error quantification from an indirect design based on noisy data, use a Kalman filter to estimate the initial condition, and reformulate the chance-constrained SPC problem as a semi-definite program (SDP).

In this work, we take a slightly different approach to the data-driven stochastic optimal control problem. We pose the control problem as one of steering the probability density of the state, as opposed to just the (mean of the) state vector. The canonical problem in this regard is to steer the state density of a linear system to some target Gaussian distribution subject to independent Gaussian disturbances and chance constraints. This problem has been solved exactly in both discrete-time [13,14] and continuous-time [15] settings, with demonstrated success in many engineering applications, such as spacecraft rendezvous [16], interplanetary and cislunar trajectory optimization [17,?,?], and autonomous driving [18], to name a few. In the Gaussian setting, and since the state remains normally distributed for the entire horizon, the problem boils down to steering the first two moments of the state, hence the approach is often referred to as covariance steering (CS).

In this paper, we develop a general framework to steer the distribution of an *unknown* stochastic LTI system using raw data collected offline, instead of using a known system model. We are interested in the so-called Data-Driven Density Steering (DD-DS) control problem and we develop a generalized framework to solve this problem, henceforth referred to as **D**ata-driven **U**ncertainty quantification and density **S**teering (DUST). Since we will be dealing primarily with linear systems, we will, alternatively, investigate the Data-Driven Mean Steering (DD-MS) and Data-Driven Covariance Steering (DD-CS) problems. By combining model-based CS theory with behavioral systems theory and statistical learning, we provide a robust framework to steer both the mean and covariance of the state distribution to a desired terminal distribution. An illustrative flowchart of the proposed framework is shown in Figure 1.

We first decompose the problem into one of data-driven mean steering (DD-MS) and one of data-driven covariance steering (DD-CS), as this framework allows for a

separation principle for the individual moment trajectories, excluding the presence of constraints (Section 2). We then follow an indirect certainty-equivalence (CE) route to exactly parameterize the mean dynamics in terms of the unknown noise realizations in the data, resulting in an uncertain quadratic program (Section 3). For covariance control, we parameterize the feedback gains directly using the collected data and use established model-based CS theory to arrive at an uncertain SDP.

To handle the uncertainty during data collection, Section 4 develops a novel noise estimation scheme to quantify the noise using techniques from maximum likelihood estimation (MLE). We then provide connections with the corresponding indirect design techniques. Using the known statistical properties from MLE, as well as quantitative notions of persistence of excitation, we construct high-confidence noise estimation error bounds in Section 5. These uncertainty bounds are then used in Section 6 to construct uncertainty sets to tractably reformulate the uncertain MS and CS problems as robust control and robust optimization problems, respectively, which can be solved efficiently to optimality using standard off-the-shelf solvers.

Since the noise enters multiplicatively in the indirect design formulation, we alternatively formulate a parametric uncertainty DD-DS (PU-DD-DS) problem and use convex relaxations to tractably solve the original problem in Section 7. Lastly, to illustrate the proposed framework, in Section 8 we perform an in-depth study of the proposed control design methods, analyzing their efficacy and precision, and compare them with their model-based counterparts. We conclude with a discussion of the proposed framework, and we offer several avenues for future extensions. A full-length version of this paper (including appendices and additional derivations) will be made available on arXiv; we cite it as [?].

Contributions

This paper extends our previous work [19,20], where we studied the limiting case of this problem, namely, the dynamics were assumed to be deterministic and the uncertainty resided solely in the boundary conditions. In this case, Willems Fundamental Lemma holds exactly, and we can achieve an exact correspondence with model-based designs. In addition, [20] studied the full data-driven density steering problem with noise quantification. However, in that work, the uncertainty sets were unnecessarily conservative (see Section 5). In this paper, we remove this restrictive assumption.

In summary, the contributions of this work are as follows:

1. We introduce a unified data-driven framework for steering the full state distribution of an unknown

stochastic LTI system using offline data, by combining behavioral systems theory with model-based covariance steering ideas.

2. We develop data-driven noise estimation and uncertainty quantification tools to find the most likely noise realization, along with its covariance, and provide finite-sample/asymptotic, high-confidence ambiguity sets for the corresponding estimation errors.
3. We provide tractable (convex optimization) formulations for both the data-driven mean steering (DD-MS) and covariance steering (DD-CS) problems. Each problem is solved so as to guarantee terminal distribution constraint satisfaction with high probability.
4. We present a parametric-uncertainty (PU-DD-DS) formulation, interpreting the certainty equivalent (CE) problem formulation in terms of controlling a system subject to multiplicative uncertainty, and provide a tractable convex formulation for its solution, showcasing an alternative approach to the data-driven density steering problem.
5. We evaluate the performance and efficacy of the proposed data-driven designs against alternative CE-based approaches and against model-based counterparts.

Notation

Real-valued vectors are denoted by lowercase letters, $u \in \mathbb{R}^m$, matrices are denoted by upper-case letters, $V \in \mathbb{R}^{n \times m}$, and random vectors are denoted by boldface letters, $\mathbf{w} \in \mathbb{R}^p$. χ_p^2 denotes the chi-square distribution with p degrees of freedom, and $Q_{\chi_p^2}(q)$ denotes its associated quantile function (i.e., the inverse of the cumulative distribution function) at the quantile $q \in (0, 1)$. The Kronecker product is denoted as $\otimes$ and the vectorization of a matrix A is denoted as $\text{vec}(A) = [a_1^\top, \dots, a_m^\top]^\top$, where a_i is the i th column of A . We use the shorthand notation $[A; B]$ to denote the vertical stacking of two matrices or vectors of compatible dimension. We define the set $\llbracket T \rrbracket \triangleq \{1, \dots, T\}$ and similarly $\llbracket T \rrbracket_0 \triangleq \{0\} \cup \llbracket T \rrbracket$, for any natural number $T \in \mathbb{N}$. Given two matrices A and B of compatible dimensions, the inner product between A and B is defined as $\langle A, B \rangle \triangleq \text{tr}(A^\top B)$. We denote the minimum non-zero eigenvalue of a matrix by $\lambda_{\min}^+$. For a matrix A , we denote its Moore-Penrose pseudoinverse by $A^\dagger$. We denote the induced matrix two-norm (or spectral norm) by $\|\cdot\|$ and the matrix Frobenius norm by $\|\cdot\|_F$. We denote the identity matrix of size $n \times n$ as I_n and the zero matrix of size $m \times n$ as $0_{m \times n}$. The set of $n \times n$ symmetric matrices is denoted by $\mathbb{S}^n$, the set of positive definite matrices by $\mathbb{S}_{++}^n$, and the set of positive semi-definite matrices by $\mathbb{S}_+^n$. For simplicity, we denote the m -long vector of zeros as 0_m . Often, we will drop the subscript in these matrices if the dimension is clear from the context. Lastly, we succinctly denote a discrete-time signal $z_0, z_1, \dots, z_T$ by $\{z_k\}_{k=0}^T$.

2 Problem Formulation

2.1 Problem Statement

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

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k + D\mathbf{w}_k, \quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^n$ is the state, $\mathbf{u}_k \in \mathbb{R}^m$ is the control input, and $\mathbf{w}_k \sim \mathcal{N}(0, I_d)$ are i.i.d. Gaussian disturbances, with time steps $k \in \llbracket N-1 \rrbracket_0$, where N represents the horizon length. Alternatively, we may re-write (1) as

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k + \boldsymbol{\xi}_k, \quad (2)$$

where $\boldsymbol{\xi}_k \sim \mathcal{N}(0, \Sigma_\xi)$ are i.i.d. random vectors for all $k \in \llbracket T-1 \rrbracket_0$, and where $\Sigma_\xi \triangleq DD^\top$ is the covariance of the noise. The constant system matrices A, B, D are assumed to be *unknown*. The initial state $\mathbf{x}_0$ is a random n -dimensional vector drawn from the normal distribution

$$\mathbf{x}_0 \sim \mathcal{N}(\mu_i, \Sigma_i), \quad (3)$$

where $\mu_i \in \mathbb{R}^n$ is the initial state mean and $\Sigma_i \in \mathbb{S}_{++}^n$ is the initial state covariance. It is assumed that the initial state is independent of the noise sequence, that is, $\mathbb{E}[\mathbf{x}_0 \mathbf{w}_k^\top] = 0$, for all $k \in \llbracket N-1 \rrbracket_0$.

The objective is to steer the trajectories of (1) from the initial distribution (3) to the terminal distribution

$$\mathbf{x}_N = \mathbf{x}_f \sim \mathcal{N}(\mu_f, \Sigma_f), \quad (4)$$

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

$$J(\mathbf{u}) \triangleq \mathbb{E} \left[\sum_{k=0}^{N-1} (\mathbf{x}_k - \mathbf{x}_k^r)^\top Q_k (\mathbf{x}_k - \mathbf{x}_k^r) + \mathbf{u}_k^\top R_k \mathbf{u}_k \right], \quad (5)$$

where $\{\mathbf{x}_k^r\}_{k=0}^{N-1}$ is a given reference trajectory, and $Q_k \succeq 0$ and $R_k \succ 0$ for all $k \in \llbracket N-1 \rrbracket_0$.

Remark 1 We assume that the system (1) is controllable, that is, for any $\mathbf{x}_0, \mathbf{x}_f \in \mathbb{R}^n$, and no noise ($\mathbf{w}_k \equiv 0$, $k \in \llbracket N-1 \rrbracket_0$), there exists a sequence of control inputs $\{\mathbf{u}_k\}_{k=0}^{N-1}$ that steer the system from $\mathbf{x}_0$ to $\mathbf{x}_f$.

It is assumed that we are given a T -long trajectory dataset $\mathbb{D} \triangleq \{\mathbf{x}_k^{(d)}, \mathbf{u}_k^{(d)}, \mathbf{x}_T^{(d)}\}_{k=0}^{T-1}$ for control design. Although it is also possible to synthesize the controllers using multiple episodic datasets $\mathbb{D}_\ell$ ($\ell = 1, \dots, M$), for simplicity, in the present work we assume only a single dataset. Borrowing from the work in [?], we assume affine state feedback policies of the form,

$$\mathbf{u}_k = K_k(\mathbf{x}_k - \mu_k) + v_k, \quad (6)$$




Fig. 1. Breakdown of DUST framework into data-collection, noise estimation, uncertainty set construction, and robust control. The noisy dataset $\mathbb{D}$ is used to estimate the past noise realization $\hat{\Xi}_{0,T}$, which is subsequently used to generate norm-bounded uncertainty sets Δ_{model} for the (indirect) DD-MS and Δ_{noise} for the (direct) DD-CS problems. The end result is optimal moment trajectories that satisfy terminal distributional constraints with high probability.

parameterized by an open-loop control sequence $v = \{v_k\}_{k=0}^{N-1}$ and a feedback gain sequence $K = \{K_k\}_{k=0}^{N-1}$, where $\mu_k \triangleq \mathbb{E}[\mathbf{x}_k]$ is the mean state.

In summary, the data-driven distribution steering (DD-DS) problem is stated below.

Problem 1 (DD-DS) *Given the dataset $\mathbb{D}$, find the feedforward $\{v_k\}_{k=0}^{N-1}$ and feedback gain $\{K_k\}_{k=0}^{N-1}$ sequences that minimize the objective function (5), subject to the initial state (3) and terminal state (4) boundary conditions.*

We note that under the control law (6), and with complete system knowledge, it is possible to re-write Problem 1 as a convex program [?], which can be solved to optimality using off-the-shelf solvers [21]. Next, we proceed to solve Problem 1 in the case where the system matrices are not known.

2.2 Problem Reformulation

Since the state distribution remains Gaussian at all time steps, we decompose the system dynamics (1) into the mean dynamics and covariance dynamics. Plugging in the control law (6) into the dynamics (1) yields the decoupled dynamics

$$\mu_{k+1} = A\mu_k + Bv_k, \quad (7a)$$

$$\Sigma_{k+1} = (A + BK_k)\Sigma_k(A + BK_k)^\top + DD^\top, \quad (7b)$$

where the state covariance $\Sigma_k \triangleq \mathbb{E}[(\mathbf{x}_k - \mu_k)(\mathbf{x}_k - \mu_k)^\top]$.

Note that $v_k \in \mathbb{R}^m$ controls the state mean and $K_k \in \mathbb{R}^{m \times n}$ controls the state covariance.

Contrary to the approaches in [14,13] that formulate a convex program in the lifted space of state and control trajectories, in this work, and similarly to [?], we treat the moments of the intermediate states $\{\mu_k, \Sigma_k\}_{k=0}^N$ over the steering horizon as decision variables subject to the dynamic constraints (7).

First, and similar to the decoupled dynamics (7), the cost function (5) can be decoupled in terms of the mean and covariance as follows

$$J = J_\mu(\mu_k, v_k) + J_\Sigma(\Sigma_k, K_k), \quad (8)$$

where,

$$J_\mu \triangleq \sum_{k=0}^{N-1} \left((\mu_k - x_k^r)^\top Q_k (\mu_k - x_k^r) + v_k^\top R_k v_k \right), \quad (9a)$$

$$J_\Sigma \triangleq \sum_{k=0}^{N-1} \left(\text{tr}(Q_k \Sigma_k) + \text{tr}(R_k K_k \Sigma_k K_k^\top) \right). \quad (9b)$$

The two boundary conditions (3) and (4) are written as

$$\mu_0 = \mu_i, \quad \mu_N = \mu_f, \quad (10a)$$

$$\Sigma_0 = \Sigma_i, \quad \Sigma_N = \Sigma_f, \quad (10b)$$

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

Problem 2 (DD-MS) Given the (unknown) mean dynamics (7a), find the optimal feed-forward control sequence $\{v_k\}_{k=0}^{N-1}$ and the corresponding optimal mean trajectory $\{\mu_k\}_{k=0}^N$ that minimize the mean cost (9a) subject to the boundary conditions (10a).

Problem 3 (DD-CS) Given the (unknown) covariance dynamics (7b), find the optimal feedback gain sequence $\{K_k\}_{k=0}^{N-1}$ and the corresponding covariance trajectory $\{\Sigma_k\}_{k=0}^N$ that minimize the covariance cost (9b) subject to the boundary conditions (10b).

3 Data-Driven Parameterization

In this section, we review the main concepts from behavioral systems theory [22] that will allow us to parametrize the decision variables in Problems 2 and 3 in terms of the collected input and output data streams.

First, recall the following definitions.

Definition 2 Given a signal $\{z_k\}$ where $z_k \in \mathbb{R}^\sigma$, we denote the Hankel matrix of depth ℓ by

$$Z_{i,\ell,j} \triangleq \begin{bmatrix} z_i & z_{i+1} & \dots & z_{i+j-1} \\ z_{i+1} & z_{i+2} & \dots & z_{i+j} \\ \vdots & \vdots & \ddots & \vdots \\ z_{i+\ell-1} & z_{i+\ell} & \dots & z_{i+\ell+j-2} \end{bmatrix} \in \mathbb{R}^{\sigma\ell \times j}, \quad (11)$$

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

$$Z_{i,1,j} \equiv Z_{i,j} = [z_i \ z_{i+1} \ \dots \ z_{i+j-1}]. \quad (12)$$

Definition 3 The signal $\{z_k\}_{k=0}^{T-1} : \llbracket T-1 \rrbracket_0 \rightarrow \mathbb{R}^\sigma$ is persistently exciting of order ℓ if the matrix $Z_{0,\ell,j}$ with $j = T - \ell + 1$ has rank $\sigma\ell$.

Suppose we carry out an experiment of duration $T \in \mathbb{N}$, where we collect input and noisy 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

$$U_{0,T} \triangleq [u_0^{(d)} \ u_1^{(d)} \ \dots \ u_{T-1}^{(d)}], \quad (13a)$$

$$X_{0,T} \triangleq [x_0^{(d)} \ x_1^{(d)} \ \dots \ x_{T-1}^{(d)}], \quad (13b)$$

$$X_{1,T} \triangleq [x_1^{(d)} \ x_2^{(d)} \ \dots \ x_T^{(d)}]. \quad (13c)$$

Using (13), the system dynamics (2) take the form

$$X_{1,T} = AX_{0,T} + BU_{0,T} + \Xi_{0,T}, \quad (14)$$

where $\Xi_{0,T} \triangleq [\xi_0^{(d)} \ \dots \ \xi_{T-1}^{(d)}] \in \mathbb{R}^{n \times T}$ is the Hankel matrix of the (unknown) disturbances, and $\xi_k^{(d)} \sim \mathcal{N}(0, \Sigma_\xi)$.

Assuming that the data is *persistently exciting* (PE), the block Hankel matrix of the input and state data has full row rank

$$\text{rank} \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} = n + m. \quad (15)$$

This PE assumption is crucial for direct data-driven control design, and is generally a mild assumption in practice, especially when noisy data is used [7]. The condition in (15) implies that *any* arbitrary input-state sequence of (1) can be expressed as a linear combination of the collected input-state data. Furthermore, this idea can be used to parameterize any arbitrary feedback interconnection [23]. In the following section, we parameterize the feedback gains in terms of the input-state data and reformulate the covariance steering problem as a semi-definite program (SDP). First, we focus on the mean steering problem, by first identifying the system matrices using the collected input-state data. Because of this extra system identification step, the data-driven mean steering problem is solved using an indirect approach.

3.1 Indirect Data-Driven Mean Steering (DD-MS)

Given the mean dynamics (7a) in terms of the open-loop control v_k , the PE condition (15) provides a system identification type of result using the following theorem.

Theorem 4 Suppose $\mathbb{D} = \{x_k^{(d)}, u_k^{(d)}, x_T^{(d)}\}_{k=0}^{T-1}$ is a dataset collected from the underlying system (1) such that the rank condition (15) holds. Then, system (7a) has the following equivalent representation

$$\mu_{k+1} = (X_{1,T} - \Xi_{0,T}) \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger \begin{bmatrix} v_k \\ \mu_k \end{bmatrix}. \quad (16)$$

PROOF. See [24] for details.

Theorem 4 provides a data-based representation of a linear system. Assuming exact knowledge of the noise realization $\Xi_{0,T}$, one may equivalently interpret equation (16) as the solution to the least-squares problem

$$\min_{B,A} \left\| X_{1,T} - \Xi_{0,T} - [B \ A] \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} \right\|_F, \quad (17)$$

where $\|\cdot\|_F$ is the Frobenius norm. Data-driven indirect designs based on the certainty-equivalence (CE) principle compute an approximate system description by solv-

ing (17) assuming no noise (i.e., $\Xi_{0,T} = 0$). Using Theorem 4, we can express the mean steering problem as the following convex optimization problem

$$\min_{\mu_k, v_k} J_\mu = \sum_{k=0}^{N-1} \left((\mu_k - x_k^r)^\top Q_k (\mu_k - x_k^r) + v_k^\top R_k v_k \right), \quad (18a)$$

such that, for all $k \in \llbracket N-1 \rrbracket_0$,

$$F_v(\Xi_{0,T})v_k + F_\mu(\Xi_{0,T})\mu_k - \mu_{k+1} = 0, \quad (18b)$$

with the boundary conditions (10a), where $F_\mu \in \mathbb{R}^{n \times n}$ and $F_v \in \mathbb{R}^{n \times m}$ result from the partition of the matrix F as follows

$$F \triangleq (X_{1,T} - \Xi_{0,T}) \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix}^\dagger = \begin{bmatrix} F_v(\Xi_{0,T}) & F_\mu(\Xi_{0,T}) \end{bmatrix}. \quad (18c)$$

3.2 Direct Data-Driven Covariance Steering (DD-CS)

To solve the DD-CS problem, first notice that from the rank condition (15), we can express the feedback gains as follows

$$\begin{bmatrix} K_k \\ I_n \end{bmatrix} = \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} G_k, \quad (19)$$

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, using this parameterization, we can re-write the covariance dynamics (7b) as

$$\begin{aligned} \Sigma_{k+1} &= [B \ A] \begin{bmatrix} K_k \\ I_n \end{bmatrix} \Sigma_k \begin{bmatrix} K_k \\ I_n \end{bmatrix}^\top [B \ A]^\top + DD^\top \\ &= (X_{1,T} - \Xi_{0,T}) G_k \Sigma_k G_k^\top (X_{1,T} - \Xi_{0,T})^\top + \Sigma_\xi. \end{aligned} \quad (20)$$

Similarly, the covariance cost (9b) can be re-written as

$$J_{\Sigma,k} = \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} G_k \Sigma_k G_k^\top U_{0,T}^\top). \quad (21)$$

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

$$\Sigma_{k+1} = (X_{1,T} - \Xi_{0,T}) S_k \Sigma_k^{-1} S_k^\top (X_{1,T} - \Xi_{0,T})^\top + \Sigma_\xi, \quad (22)$$

and

$$J_{\Sigma,k} = \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} S_k \Sigma_k^{-1} S_k^\top U_{0,T}^\top). \quad (23)$$

This problem is still non-convex due to the nonlinear term $S_k \Sigma_k^{-1} S_k^\top$. To this end, we define a new set of decision variables $Y_k \succeq S_k \Sigma_k^{-1} S_k^\top$ and relax the objective

function with these new variables. Similarly, we relax the covariance propagation constraints to soft (inequality) constraints, which yields the relaxed optimization problem

$$\min_{\Sigma_k, S_k, Y_k} \bar{J}_\Sigma = \sum_{k=0}^{N-1} \left(\text{tr}(Q_k \Sigma_k) + \text{tr}(R_k U_{0,T} Y_k U_{0,T}^\top) \right), \quad (24a)$$

such that, for all $k \in \llbracket N-1 \rrbracket_0$,

$$S_k \Sigma_k^{-1} S_k^\top - Y_k \preceq 0, \quad (24b)$$

$$(X_{1,T} - \Xi_{0,T}) S_k \Sigma_k^{-1} S_k^\top (X_{1,T} - \Xi_{0,T})^\top + \Sigma_\xi - \Sigma_{k+1} \preceq 0, \quad (24c)$$

$$\Sigma_k - X_{0,T} S_k = 0, \quad (24d)$$

$$\Sigma_N - \Sigma_f = 0. \quad (24e)$$

The equality constraint (24d) comes from the second block in (19) by multiplying Σ_k on the right. The relaxed optimization problem (24) is convex, since the constraints (24b) and (24c) can be written using the Schur complement in terms of the linear matrix inequalities (LMI)

$$\begin{bmatrix} \Sigma_k & S_k^\top \\ S_k & Y_k \end{bmatrix} \succeq 0, \quad (25a)$$

$$G_\Sigma^k \triangleq \begin{bmatrix} \Sigma_{k+1} - \Sigma_\xi & (X_{1,T} - \Xi_{0,T}) S_k \\ S_k^\top (X_{1,T} - \Xi_{0,T})^\top & \Sigma_k \end{bmatrix} \succeq 0. \quad (25b)$$

The cost (24a) and the equality constraints (24d)-(24e), on the other hand, are linear in all the decision variables, and hence are trivially convex.

4 Noise Estimation Algorithms

The optimization problems (18) and (24), as they stand, albeit convex, are still intractable because we know neither the disturbance covariance Σ_ξ nor the noise realization history $\Xi_{0,T}$. To this end, in this section we estimate the noise realization matrix $\Xi_{0,T}$ and the disturbance matrix Σ_ξ from the collected data $\mathbb{D}$ using maximum likelihood estimation (MLE). We then briefly connect this estimator to the corresponding indirect (certainty-equivalence) identification approach.

4.1 ML Noise Estimation

To encode the stochastic linear system dynamics (2) as a constraint that we can use in the MLE scheme, we need to enforce *consistency* of the realization data. To this end, any realization of the dynamics (2) must satisfy

(14). For notational convenience, henceforth, we denote the augmented Hankel matrix in (13) as

$$\mathcal{S} \triangleq \begin{bmatrix} U_{0,T} \\ X_{0,T} \end{bmatrix} \in \mathbb{R}^{(m+n) \times T}, \quad (26)$$

from which we may re-write the dynamics realization (14) as

$$X_{1,T} = [B \ A]\mathcal{S} + \Xi_{0,T}. \quad (27)$$

Additionally, noting that the matrix pseudoinverse satisfies the property $\mathcal{S}\mathcal{S}^\dagger\mathcal{S} = \mathcal{S}$, equation (27) can be written, equivalently, as

$$X_{1,T} - \Xi_{0,T} = [B \ A]\mathcal{S}\mathcal{S}^\dagger\mathcal{S}. \quad (28)$$

Inserting the relation $X_{1,T} - \Xi_{0,T} = [B \ A]\mathcal{S}$ into the right-hand side of (28) yields

$$(X_{1,T} - \Xi_{0,T})(I_T - \mathcal{S}^\dagger\mathcal{S}) = 0. \quad (29)$$

Equation (29) is a *model-free* type of condition that must be satisfied for all noisy linear system data realizations, and hence is a consistency relation for any feasible set $\mathcal{S}$ of data.

Given the constraint (29), the MLE problem then becomes

$$\max_{\Xi_{0,T}, \Sigma_\xi} \mathcal{J}_{\text{ML}}(\Xi_{0,T}, \Sigma_\xi) = \sum_{k=0}^{T-1} \log \rho_\xi(\xi_k) \quad (30a)$$

$$(X_{1,T} - \Xi_{0,T})(I_T - \mathcal{S}^\dagger\mathcal{S}) = 0, \quad (30b)$$

where $\rho_\xi(\xi)$ is the probability density function (PDF) of the random vector ξ , given as

$$\rho_\xi(\xi) = \frac{1}{(2\pi)^{n/2}} (\det \Sigma_\xi)^{-1/2} \exp \left(-\frac{1}{2} \xi^\top \Sigma_\xi^{-1} \xi \right). \quad (31)$$

Remark 5 Since the dynamics (1) are uncertain, it follows that there may be multiple noise realizations $\Xi_{0,T}$ that satisfy the linear dynamics constraints (29). Consequently, the purpose of the constrained MLE problem (30) is to find the most likely noise realization sequence $\hat{\Xi}_{0,T}$ given that the noise is normally distributed according to (31).

The next theorem provides the optimal solution to the MLE problem (30).

Theorem 6 The solution to the MLE problem for the most probable noise realization $\Xi_{0,T}$ and disturbance co-

variance matrix Σ_ξ is given by

$$\Xi_{0,T}^* = X_{1,T}(I_T - \mathcal{S}^\dagger\mathcal{S}), \quad (32a)$$

$$\Sigma_\xi^* = \frac{1}{T} X_{1,T}(I_T - \mathcal{S}^\dagger\mathcal{S}) X_{1,T}^\top. \quad (32b)$$

Remark 7 The solution (32) of the MLE program in (30) is contingent on $\Sigma_\xi \succ 0$. In fact, the optimal covariance estimate is simply the sample covariance of the dataset with respect to the estimated noise realizations, i.e., $\Sigma_\xi^* = \frac{1}{T} \Xi_{0,T}^* (\Xi_{0,T}^*)^\top$. If Σ_ξ is singular, however, then $\log \det \Sigma_\xi$ is undefined, hence the problem is infeasible. As a result, regularization techniques (e.g., GLASSO [25], distributionally robust estimation [26]) should be used in these cases instead. It should also be noted that such degenerate cases arise when the number of disturbance channels is less than the number of states channels, i.e., $D \in \mathbb{R}^{n \times d}$, with $d < n$.

4.2 Indirect Design Estimation

We conclude this section by observing that an alternative way to extract disturbance information from noisy data is by examining the difference between the observed state and the state prediction from the dynamics model. Referring back to (2), and assuming knowledge of the system matrices A and B , the disturbance would be given by

$$\xi_k^{(d)} = x_{k+1}^{(d)} - Ax_k^{(d)} - Bu_k^{(d)}, \quad k \in \llbracket T-1 \rrbracket_0. \quad (33)$$

Given the collected data, and under the rank condition (15), an estimate $(\hat{B}, \hat{A})$ of the system matrices can be obtained as the unique solution to the (noiseless) least-squares problem

$$[\hat{B} \ \hat{A}] = \operatorname{argmin}_{B,A} \|X_{1,T} - [B \ A]\mathcal{S}\|_F = X_{1,T}\mathcal{S}^\dagger. \quad (34)$$

Concatenating (33) over the entire sampling horizon yields the equality $\Xi_{0,T} = X_{1,T} - [B \ A]\mathcal{S}$ (see also (27)). We may then estimate the noise sequence from the estimated model parameters in (34) as $\hat{\Xi}_{0,T} = X_{1,T} - [\hat{B} \ \hat{A}]\mathcal{S}$, or

$$\hat{\Xi}_{0,T} = X_{1,T}(I_T - \mathcal{S}^\dagger\mathcal{S}), \quad (35)$$

from which we may compute the disturbance covariance as in (32b). The procedure, then, is to first estimate the nominal model, then estimate the disturbance structure, and lastly solve the associated CS problem using these estimated model parameters.

Notice that the optimal noise realization under the assumption of a *noiseless* system, (35), is equivalent to that of the optimal ML noise realization, (32a). Thus, the CE estimation solution (35) is equivalent to the MLE solution (32a) under a known disturbance structure, implying that there may be deeper parallels between indirect

and direct design methods in the context of noisy data. For an overview of this observation, please see [27].

The estimated value of $\hat{\Xi}_{0,T}$ (and, perhaps, $\hat{\Sigma}_{\xi}$) will differ, in general, from its true value $\Xi_{0,T}$. For control design, we need to quantify the estimation error $\Delta\Xi_{0,T} \triangleq \Xi_{0,T} - \hat{\Xi}_{0,T}$. In the next section, we therefore derive high-probability bounds on the estimation error $\Delta\Xi_{0,T}$ which is used in Section 6 in a robust data-driven design by enforcing the running and terminal state constraints to hold within the given uncertainty sets, defined by these bounds, with high probability.

5 Uncertainty Set Synthesis

Given the noise estimation techniques outlined in Section 4, we now derive high-confidence bounds for the noise realization estimation errors that will be used later in the robust control design. As mentioned earlier, this is a necessary step to account for the model mismatch due to noisy data. In particular, we leverage the statistical properties of the MLE-based noise realization estimate to construct tractable uncertainty sets for the estimation error.

5.1 Moment-Based Ambiguity Sets

In light of the preceding discussion, we are interested in practical bounds we can implement to ensure robust satisfaction of the constraints for the DD-MS (e.g., (10a)) and DD-CS (e.g., (24c)) problems. To do so, and equipped with the MLE noise realization estimate (32a), we will use the statistical properties of $\hat{\Xi}_{0,T}$ to generate an ellipsoidal uncertainty set based on some degree of confidence $1 - \delta \in [0.5, 1)$. In the context of the control design problem, this will imply that the resulting controller will steer the system to the desired final distribution *for all* uncertainty estimates $\Delta\Xi_{0,T} \in \Delta$, in some compact set Δ , with a prescribed degree of confidence level $1 - \delta$.

For simplicity, assume $\Sigma_{\xi} \succ 0$ is known. First, we re-write the MLE problem (30) in terms of the vectorized parameters to be estimated $\xi \triangleq \text{vec}(\Xi_{0,T}) = [\xi_0^T, \dots, \xi_{T-1}^T]^T \in \mathbb{R}^{nT}$ as the optimization problem below, where for notational simplicity, we have dropped the superscript (d) from $\xi^{(d)}$:

$$\min_{\xi} \mathcal{J}_{\text{ML}}(\xi \mid \mathbb{D}) = \frac{1}{2} \xi^T (I_T \otimes \Sigma_{\xi}^{-1}) \xi \quad (36a)$$

$$C(\xi) \triangleq (\Gamma \otimes I_n) \xi - \lambda = 0, \quad (36b)$$

where $\Gamma \triangleq I_T - \mathcal{S}^\dagger \mathcal{S} \in \mathbb{R}^{T \times T}$, and $\lambda \triangleq \text{vec}(X_{1,T} \Gamma)$. It can be shown [28] that, as the number of samples grows, the unconstrained ML noise estimate $\hat{\xi}$ is asymptotically efficient and converges to a normal distribution

as $\sqrt{T} \Delta \xi \xrightarrow{d} \mathcal{N}(0, \mathcal{I}^{-1})$, where $\Delta \xi \triangleq \xi - \hat{\xi}$ and where $\mathcal{I} = \mathbb{E}_{\xi} \left[\frac{\partial^2}{\partial \xi^2} \mathcal{J}_{\text{ML}}(\xi \mid \mathbb{D}) \right]$ is the *Fisher Information Matrix* (FIM), which is given by $\mathcal{I} = I_T \otimes \Sigma_{\xi}^{-1}$ in the unconstrained case.

In our setting, the constraint Jacobian $J(\xi) \triangleq \partial C(\xi) / \partial \xi = \Gamma \otimes I_n$ is generally *rank-deficient* because Γ is an orthogonal projector. The constraint manifold is nevertheless smooth, and asymptotic efficiency holds on the tangent space if we enforce only an *independent* set of equality constraints. Concretely, let $V_{\perp} \in \mathbb{R}^{T \times p}$ have orthonormal columns spanning $\text{range}(\Gamma)$ (so that $p = \text{rank}(\Gamma) = T - (n+m)$), and define the full row-rank Jacobian

$$J_r \triangleq (V_{\perp}^T \otimes I_n) \in \mathbb{R}^{(pn) \times (nT)},$$

which satisfies $\text{range}(J_r) = \text{range}(J)$ and encodes the same constraint manifold. Under standard regularity conditions, and with J_r used as the effective constraint Jacobian, the constrained MLE is $\sqrt{T}$ -consistent, asymptotically normal, and asymptotically efficient on the tangent space of the constraints [29,30].

It follows from [31] that the asymptotic covariance of the constrained MLE error is the projection of $\mathcal{I}^{-1}$ onto the tangent space. Equivalently, one may write

$$\Sigma_{\Delta} = \mathcal{I}^{-1} - \mathcal{I}^{-1} J_r^T (J_r \mathcal{I}^{-1} J_r^T)^{-1} J_r \mathcal{I}^{-1} \quad (37a)$$

$$= \mathcal{I}^{-1} - \mathcal{I}^{-1} J^T (J \mathcal{I}^{-1} J)^{\dagger} J \mathcal{I}^{-1}. \quad (37b)$$

Using the MLE error covariance matrix (37), we can construct high-confidence uncertainty sets for use later in a robust control design (Section 6). To this end, we first compute the analytical form of the error covariance for the constrained MLE problem (30).

Lemma 8 *The distribution of the error of the constrained ML estimator (30) for the unknown noise realization $\xi = \text{vec}(\Xi_{0,T})$ converges, as $T \rightarrow \infty$, to the normal distribution $\mathcal{N}(0, \Sigma_{\Delta})$, that is, $\Delta \xi \sim \mathcal{N}(0, \Sigma_{\Delta})$, where $\Sigma_{\Delta} = \mathcal{S}^\dagger \mathcal{S} \otimes \Sigma_{\xi}$.*

Given the asymptotic noise estimation error covariance Σ_{Δ} , one can build uncertainty sets for $\Delta\Xi_{0,T}$ by taking quantiles of the associated Gaussian distribution. A conservative option (used in our earlier work [20]) treats the joint error as full-dimensional and yields a radius with nT degrees of freedom; we omit that derivation here for brevity.

Instead, we directly construct the tight confidence set by accounting for the fact that Σ_{Δ} is singular (Lemma 8), which reduces the effective degrees of freedom. We first provide a formal definition of a normal distribution that takes into account singular covariance matrices.

Corollary 9 For the MLE problem (30), the associated $(1 - \delta)$ -quantile uncertainty set may be taken as

$$\Delta^* = \left\{ \|\Delta \Xi_{0,T}\| \leq \|\Sigma_\xi^{1/2}\| Q_{\chi^2_{n(n+m)}}^{1/2} (1 - \delta) \right\}, \quad (38)$$

that is, $\mathbb{P}\{\Delta \Xi_{0,T} \in \Delta^*\} \geq 1 - \delta$.

Proposition 10 ([26]) Let $\mathbb{P}$ be a normal distribution on $\mathbb{R}^p$ with mean $\mu \in \mathbb{R}^p$ and covariance matrix $\Sigma \succeq 0$, that is, $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$, and let $r = \text{rank}(\Sigma)$. Then, $\mathbb{P}$ is supported on $\text{supp}(\mathbb{P}) \triangleq \{\mu + Ev : v \in \mathbb{R}^r\}$, and its density with respect to the Lebesgue measure on $\text{supp}(\mathbb{P})$ is given by

$$\rho_{\mathbb{P}}(\xi) \triangleq \frac{1}{\sqrt{(2\pi)^r \det(D)}} e^{-\frac{1}{2}(\xi - \mu)^\top E D^{-1} E^\top (\xi - \mu)}, \quad (39)$$

where $D \succ 0$ is the diagonal matrix of the positive eigenvalues of Σ , and $E \in \mathbb{R}^{p \times r}$ is the matrix whose columns correspond to the orthonormal eigenvectors of the positive eigenvalues of Σ .

Next, we use this result to construct a confidence ellipsoid for a normal distribution with a singular covariance matrix by recognizing that $(\xi - \mu)^\top E D^{-1} E^\top (\xi - \mu)$ is a χ_r^2 random variable with r degrees of freedom [32].

Proposition 11 Given the normal distribution $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$ with mean $\mu \in \mathbb{R}^p$ and covariance matrix $\Sigma \succeq 0$, the associated uncertainty set

$$\mathcal{U}_{\mathcal{N}} \triangleq \{\xi \in \text{supp}(\mathbb{P}) : (\xi - \mu)^\top E D^{-1} E^\top (\xi - \mu) \leq Q_{\chi_r^2}(1 - \delta)\} \quad (40)$$

contains the $(1 - \delta)$ -quantile of the distribution $\mathbb{P}$, that is, $\mathbb{P}(\xi \in \mathcal{U}_{\mathcal{N}}) \geq 1 - \delta$.

The uncertainty set for the joint distribution $\Delta \Xi_{0,T}$ from MLE can now be constructed by recognizing that $\text{rank}(\Sigma_\Delta) = n(n + m)$ since¹ $\Sigma_\Delta = \mathcal{S}^\dagger \mathcal{S} \otimes \Sigma_\xi$ from Lemma 8, and $\text{rank}(\mathcal{S}^\dagger \mathcal{S}) = n + m$. We arrive at the set

$$\Delta^* = \left\{ \|\Delta \Xi_{0,T}\| \leq \|\Sigma_\xi^{1/2}\| Q_{\chi^2_{n(n+m)}}^{1/2} (1 - \delta) \right\}, \quad (41)$$

which guarantees that $\mathbb{P}(\Delta \Xi_{0,T} \in \Delta^*) \geq 1 - \delta$. Since, in general, $r = n(n + m) \ll nT$, it follows [33] that $Q_{\chi_r^2}(1 - \delta) < Q_{\chi_{Tn}^2}(1 - \delta)$, and hence $\Delta^* \subseteq \Delta$.

Remark 12 The above analysis relied on the asymptotic properties of the MLE estimator to obtain the error covariance of the noise realization estimates through Lemma 8. One natural question is whether the derived bound still holds for the case of a finite T -long input/state

¹ In general, Σ_Δ becomes a very low rank matrix, since when T is large, $Tn \gg n(n + m)$.

dataset. $\mathbb{D}$. In the full-length version [?] we show that, in fact, the finite-sample norm estimate coincides with the uncertainty set radius $\rho = \|\Sigma_\xi^{1/2}\| Q_{\chi^2_{n(n+m)}}^{1/2} (1 - \delta)$ in (41) as long as the data is persistently exciting.

6 Robust DD-DS

Equipped with the machinery to efficiently estimate and bound the uncertainty due to the noise, we are now in a position to tackle the uncertain convex programs in (18) and (24). The analysis in this section will lead to robust versions of the DD-MS and DD-CS problems. To this end, notice that the DD-MS problem essentially becomes a robust control problem with unstructured model uncertainty, albeit only with high probability guarantees. Hence, in Section 6.1, we use techniques from system-level synthesis (SLS) [34] to tractably enforce terminal constraint satisfaction (with high probability) for all bounded model uncertainties arising from the methods in Section 5. We note that this procedure is similar to the *Coarse-ID* proposed in [35], which used SLS to solve a robust control problem using uncertainty bounds constructed during the model identification step. Therein the authors employed techniques from RMT to construct tight uncertainty sets, which provide a promising framework for analyzing the estimation errors arising from both noise and model estimation. The main difference between our work and [35] is that we employ an indirect design technique *solely* for the mean steering problem, while other robust optimization methods (i.e., the robust counterpart of uncertain LMIs), based on a direct design, are utilized to address the covariance steering problem. In this regard, the DD-CS problem requires robust satisfaction of LMI constraints over the planning horizon that encode the covariance propagation constraints under the chosen feedback control strategy. In Section 6.3, we will form the robust counterpart of these semi-infinite constraints, and use techniques from robust optimization to tractably enforce these as equivalent, deterministic LMIs. First, we address the robust DD-MS control design problem.

6.1 Robust Problem Formulation for DD-MS

Solving the DD-MS program (18) by replacing the true $\Xi_{0,T}$ with its estimate $\hat{\Xi}_{0,T}$ will result in optimal controllers that do *not* satisfy the terminal constraint $\mu_N = \mu_f$ due to the inaccuracy in the estimated model from the indirect design step (34). The true mean dynamics are

$$\mu_{k+1} = (\hat{A} + \Delta A)\mu_k + (\hat{B} + \Delta B)v_k, \quad (42)$$

where the nominal matrices $[\hat{B} \ \hat{A}] = X_{1,T} \mathcal{S}^\dagger$ are computed from CE estimation, and where the model deviations $\Delta A, \Delta B$ are bounded as

$$\|[\Delta B \ \Delta A]\| \leq \rho(\delta)/\sigma_{\min}(\mathcal{S}) \triangleq \alpha(\delta), \quad (43)$$

where we have used the fact

$$\|\Delta \Xi_{0,T}\| = \|[\Delta B \ \Delta A] \mathcal{S}\| \geq \sigma_{\min}(\mathcal{S}) \|[\Delta B \ \Delta A]\|, \quad (44)$$

where $\|\Delta \Xi_{0,T}\| \leq \rho(\delta)$ from Corollary 9.

Note, however, that enforcing the terminal constraint $\mu_N = \mu_f$ for all uncertainties $\|[\Delta B \ \Delta A]\| \leq \alpha(\delta)$ is intractable, in general. Instead, we relax the pointwise terminal constraint to a terminal set given by a polytope such that $\mu_N \in \mathcal{X}_f \triangleq \{x : F_{x_N} x \leq b_{x_N}\}$, and require robust satisfaction of the constraint $\mu_N \in \mathcal{X}_f$, for all $\|\Delta A\| \leq \varepsilon_A$ and for all $\|\Delta B\| \leq \varepsilon_B$, for some $\varepsilon_A, \varepsilon_B > 0$. Along these lines, and in order to enhance tractability, we also impose polyhedral constraints on the transient motion of the mean state and the feed-forward input as $\mu_k \in \mathcal{X}_k \triangleq \{x : F_{x_k} x \leq b_{x_k}\}$ and $v_k \in \mathcal{U}_k \triangleq \{u : F_{u_k} u \leq b_{u_k}\}$. **Lastly, instead of the open-loop control v_k , we introduce a feedback mean control in terms of the mean state history as follows $v_k = \sum_{i=0}^k L_{k,i} \mu_i$.**

For notational convenience, let the *nominal* mean state be denoted by $\bar{\mu}_k$ which satisfies the error-free dynamics $\bar{\mu}_{k+1} = \hat{A} \bar{\mu}_k + \hat{B} \bar{v}_k$ and let $\bar{v}_k = \sum_{i=0}^k L_{k,i} \bar{\mu}_i$. In summary, the **robust DD-MS (R-DD-MS) problem is posed as follows.**

$$\begin{aligned} \min_{L_{k,i}} \tilde{J}_\mu &= \sum_{k=0}^N \left((\bar{\mu}_k - x_k^r)^\top Q_k (\bar{\mu}_k - x_k^r) + \bar{v}_k^\top R_k \bar{v}_k \right) \\ &\quad \text{such that, for all } k \in \llbracket N-1 \rrbracket \text{ and for all } \|\Delta A\| \leq \varepsilon_A \\ &\quad \text{and } \|\Delta B\| \leq \varepsilon_B, \\ \bar{\mu}_{k+1} &= \hat{A} \bar{\mu}_k + \hat{B} \bar{v}_k, \\ \mu_{k+1} &= (\hat{A} + \Delta A) \mu_k + (\hat{B} + \Delta B) v_k, \\ v_k &= \sum_{i=0}^k L_{k,i} \mu_i, \\ \mu_k &\in \mathcal{X}_k, \ v_k \in \mathcal{U}_k, \ \mu_N \in \mathcal{X}_f. \end{aligned} \quad \begin{aligned} (45a) \\ (45b) \\ (45c) \\ (45d) \\ (45e) \end{aligned}$$

There are **numerous methods in the robust control literature to tackle the R-DD-MS problem** in (45), ranging from the early works of tube-MPC [2] to disturbance-feedback with lumped uncertainty [36], to the more modern methods using system level synthesis (SLS) [34]. Below, we use SLS to solve the R-DD-MS problem by reformulating the semi-infinite program (45) as a tractable SDP.

6.2 Solution to the Robust DD-MS Problem via SLS

The SLS approach to robust control aims at transforming the optimization problem over feedback control laws to one over closed-loop system *responses*, i.e., linear

maps from the uncertainty process to the states and inputs in the closed loop. To this end, and for the nominal dynamics (45b), define the augmented state and control inputs as $\bar{\mu} = [\bar{\mu}_0^\top, \dots, \bar{\mu}_N^\top]^\top$, $\bar{v} = [\bar{v}_0^\top, \dots, \bar{v}_N^\top]^\top$, and the vector² $\tilde{w} \triangleq [\mu_0; w] = [\mu_0; 0_{Nn}]$. Let the control input $\bar{v} = L \bar{\mu}$, where

$$L = \begin{bmatrix} L_{0,0} & & & \\ L_{1,0} & L_{1,1} & & \\ \vdots & \ddots & \ddots & \\ L_{N,1} & \cdots & L_{N,N-1} & L_{N,N} \end{bmatrix}, \quad (46)$$

and concatenate the dynamics matrices as $\hat{\mathcal{A}} \triangleq \text{blkdiag}(I_N, 0) \otimes \hat{A}$ and $\hat{\mathcal{B}} \triangleq \text{blkdiag}(I_N, 0) \otimes \hat{B}$. Let $Z \in \mathbb{R}^{n(N+1) \times n(N+1)}$ be the block-downshift operator, that is, a matrix with the identity matrix on the first block sub-diagonal and zeros elsewhere. Under the feedback controller L , the closed-loop behavior of the nominal system (45b) can be represented compactly as

$$\bar{\mu} = Z(\hat{\mathcal{A}} + \hat{\mathcal{B}}L)\bar{\mu} + \tilde{w}. \quad (47)$$

The closed-loop map from $\tilde{w} \mapsto (\bar{\mu}, \bar{v})$ is given by

$$\begin{bmatrix} \bar{\mu} \\ \bar{v} \end{bmatrix} = \begin{bmatrix} (I - Z(\hat{\mathcal{A}} + \hat{\mathcal{B}}L))^{-1} \\ L(I - Z(\hat{\mathcal{A}} + \hat{\mathcal{B}}L))^{-1} \end{bmatrix} \tilde{w} = \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} \tilde{w}, \quad (48)$$

where the matrices $\bar{\Phi}_x$ and $\bar{\Phi}_u$ are the *nominal* system responses under the action of the feedback controller L in (46) on the LTI system (45b). The essence of the SLS approach is to treat these closed-loop system maps as the decision variables in the resulting optimization problem. In order to satisfy the closed-loop dynamics in (48), the matrices $\bar{\Phi}_x$ and $\bar{\Phi}_u$ must be constrained to an affine subspace parameterizing all system responses, similar to the subspace relations in (29) that encode the LTI dynamics of the realization data. The next theorem formalizes this intuition and provides the corresponding controller.

Theorem 13 ([34]) *Consider the nominal system dynamics (45b) with state feedback law $\bar{v} = L \bar{\mu}$, where L is a block-lower triangular matrix. Then, the following are true:*

i) *The affine subspace defined by*

$$\begin{bmatrix} I - Z\hat{\mathcal{A}} - Z\hat{\mathcal{B}} \end{bmatrix} \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} = I, \quad (49)$$

² This is a special case of the more general expression $w = [\mu_0^\top, w_0^\top, \dots, w_{N-1}^\top]^\top$, where $w_k \in \mathcal{W}$ are additive bounded uncertainties to the dynamics (45b).

parameterizes all possible system responses (48).

ii) For any block-lower triangular matrices $\bar{\Phi}_x$ and $\bar{\Phi}_u$ satisfying (49), the controller $L = \bar{\Phi}_u \bar{\Phi}_x^{-1}$ achieves the desired response.

Thus, instead of requiring satisfaction of the dynamic constraints (45b), we can, equivalently, require satisfaction of the system map constraints (49) that achieve the desired system response. Moreover, this framework can be extended to handle model uncertainty as well, through the following theorem.

Theorem 14 ([34]) *Let $\bar{\Lambda}$ be an arbitrary block-lower triangular matrix, and suppose that $\bar{\Phi}_x$ and $\bar{\Phi}_u$ are block-lower triangular matrices that satisfy*

$$\begin{bmatrix} I_{(N+1)n} - Z\hat{A} - Z\hat{B} \end{bmatrix} \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} = I_{(N+1)n} - \bar{\Lambda}. \quad (50)$$

If $(I - \bar{\Lambda}_{i,i})^{-1}$ exists for all $i = 0, \dots, N$, then the controller $L = \bar{\Phi}_u \bar{\Phi}_x^{-1}$ achieves the system response

$$\begin{bmatrix} \mu \\ v \end{bmatrix} = \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} (I - \bar{\Lambda})^{-1} \tilde{w}. \quad (51)$$

To see how Theorem 14 allows us to encode model uncertainty into the SLS framework, note that the nominal responses also approximately satisfy (49) with respect to the true model (45c) with an extra perturbation term given by

$$\begin{aligned} \begin{bmatrix} I - Z\hat{A} - Z\hat{B} \end{bmatrix} \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} &= I - Z \begin{bmatrix} \Delta A & \Delta B \end{bmatrix} \begin{bmatrix} \bar{\Phi}_x \\ \bar{\Phi}_u \end{bmatrix} \\ &\triangleq I - \Lambda \bar{\Phi}, \end{aligned} \quad (52)$$

where in the first equality we use the fact that $\bar{\Phi}_x$ and $\bar{\Phi}_u$ satisfy (49), and in the second equality, we define $\bar{\Phi} \triangleq [\bar{\Phi}_x; \bar{\Phi}_u]$, and $\Lambda \triangleq Z[\Delta A \ \Delta B]$. As a result, invoking Theorem 14 we conclude that the controller $L = \bar{\Phi}_u \bar{\Phi}_x^{-1}$, computed using the system estimates $\{\hat{A}, \hat{B}\}$, achieves the response (51) on the actual system $\{\mathcal{A}, \mathcal{B}\}$, with $\bar{\Lambda} = \Lambda \bar{\Phi}$.

Next, following standard SLS lifting [34], define $\mathcal{Q} \triangleq \text{blkdiag}(Q_0, \dots, Q_N)$, $\mathcal{R} \triangleq \text{blkdiag}(R_0, \dots, R_N)$, and $\mathcal{M} \triangleq \text{blkdiag}(\mathcal{Q}, \mathcal{R})$. Since $\tilde{w}$ has only the initial-state block nonzero, the cost depends only on the first block-column of the response. Partition $\bar{\Phi} = [\bar{\Phi}^0 \ \bar{\Phi}^{\tilde{w}}]$, where $\bar{\Phi}^0 \in \mathbb{R}^{(N+1)(n+m) \times n}$ and $\bar{\Phi}^{\tilde{w}} \in \mathbb{R}^{(N+1)(n+m) \times Nn}$. We

then minimize

$$\bar{\mathcal{J}}_\mu = \left\| \mathcal{M}^{1/2} \bar{\Phi}^0 \mu_0 \right\|_2^2 - 2 \begin{bmatrix} \mathcal{Q} x^r \\ 0_{(N+1)m} \end{bmatrix}^\top \bar{\Phi}^0 \mu_0 + \left\| \mathcal{Q}^{1/2} x^r \right\|_2^2. \quad (53)$$

We also stack all constraints $x_k \in \mathcal{X}_k$, $u_k \in \mathcal{U}_k$, and $x_N \in \mathcal{X}_f$ as $F[\mu; v] \leq b$ (see [34,37] for details). In summary, the R-DD-MS problem (45) can be written as the following program

$$\min_{\bar{\Phi}_x, \bar{\Phi}_u} \left\| \mathcal{M}^{1/2} \bar{\Phi}^0 \mu_0 \right\|_2^2 - 2[\mathcal{Q} x^r; 0_{(N+1)m}]^\top \bar{\Phi}^0 \mu_0 + \left\| \mathcal{Q}^{1/2} x^r \right\|_2^2, \quad (54a)$$

$$\begin{bmatrix} I - Z\hat{A} - Z\hat{B} \end{bmatrix} \bar{\Phi} = I, \quad (54b)$$

$$F(I + \bar{\Phi} \Lambda (I - \bar{\Phi} \Lambda)^{-1}) \bar{\Phi}^0 \mu_0 \leq b. \quad (54c)$$

The main difficulty in (54) is in the robust constraints (54c), which are nonlinear in $\bar{\Phi}$. Following the work of [37], we can upper bound the LHS of the constraints and formulate sufficient conditions such that (54c) holds, for all $\|\Delta A\| \leq \varepsilon_A$ and $\|\Delta B\| \leq \varepsilon_B$. To this end, define the partitions $\bar{\Phi}^0 = [\bar{\Phi}_x^0; \bar{\Phi}_u^0]$ and $\bar{\Phi}^{\tilde{w}} = [\bar{\Phi}_x^{\tilde{w}}; \bar{\Phi}_u^{\tilde{w}}]$, where $\bar{\Phi}_x^0 \in \mathbb{R}^{(N+1)n \times n}$, $\bar{\Phi}_u^0 \in \mathbb{R}^{(N+1)m \times n}$, $\bar{\Phi}_x^{\tilde{w}} \in \mathbb{R}^{(N+1)n \times Nn}$, $\bar{\Phi}_u^{\tilde{w}} \in \mathbb{R}^{(N+1)m \times Nn}$. The convex approximation is stated in the following theorem.

Theorem 15 *Let $\tau, \gamma > 0$ and $\theta \in (0, 1)$. Consider the following convex optimization problem*

$$\min_{\bar{\Phi}_x, \bar{\Phi}_u} \left\| \mathcal{M}^{1/2} \bar{\Phi}^0 \mu_0 \right\|_2^2 - 2[\mathcal{Q} x^r; 0_{(N+1)m}]^\top \bar{\Phi}^0 \mu_0 \quad (55a)$$

$$\begin{bmatrix} I - Z\hat{A} - Z\hat{B} \end{bmatrix} \bar{\Phi} = I, \quad (55b)$$

$$F_j^\top \bar{\Phi}^0 \mu_0 + \|F_j^\top \bar{\Phi}^{\tilde{w}}\| \frac{1 - \tau^N}{1 - \tau} \gamma \leq b_j, \quad \forall j \in [J], \quad (55c)$$

$$\left\| \begin{bmatrix} \frac{\varepsilon_A}{\theta} \bar{\Phi}_x^{\tilde{w}} \\ \frac{\varepsilon_B}{1-\theta} \bar{\Phi}_u^{\tilde{w}} \end{bmatrix} \right\| \leq \tau, \quad \left\| \begin{bmatrix} \frac{\varepsilon_A}{\theta} \bar{\Phi}_x^0 \\ \frac{\varepsilon_B}{1-\theta} \bar{\Phi}_u^0 \end{bmatrix} \mu_0 \right\| \leq \gamma, \quad (55d)$$

where $F_j \in \mathbb{R}^{(N+1)(n+m)}$, $b_j \in \mathbb{R}$ denote the j th row and element of F and b , respectively. Let the solution to this optimization problem be $\bar{\Phi} = [\bar{\Phi}_x; \bar{\Phi}_u] = [\bar{\Phi}^0 \ \bar{\Phi}^{\tilde{w}}]$ where $\bar{\Phi}_x$ and $\bar{\Phi}_u$ are lower-block triangular matrices. Then, the controller $L = \bar{\Phi}_u \bar{\Phi}_x^{-1}$ guarantees the constraint satisfaction in (54c) for all possible model uncertainties $\|\Delta A\| \leq \varepsilon_A$ and $\|\Delta B\| \leq \varepsilon_B$ in (45), that is, the solution to (55) is a conservative convex approximation of (54).

6.3 Solution of the Robust DD-CS Problem

In this section, we reformulate the uncertain CS program (24) so that it is amenable to a tractable convex matrix feasibility problem. The original constraints in

(24c), when reformulated as the LMI constraints $G_k^\Sigma \succeq 0$ in (25b), may be robustly satisfied with high-probability, using the decomposition $\Xi_{0,T} = \hat{\Xi}_{0,T} + \Delta\Xi_{0,T}$ along with the established estimation error bounds, as the semi-infinite uncertain LMIs

$$\hat{G}_k^\Sigma + \Delta G_k^\Sigma(\Delta\Xi_{0,T}) \succeq 0, \quad \forall \|\Delta\Xi_{0,T}\| \leq \rho(\delta), \quad (56)$$

where,

$$\hat{G}_k^\Sigma = \begin{bmatrix} \Sigma_{k+1} - \Sigma_\xi & (X_{1,T} - \hat{\Xi}_{0,T})S_k \\ S_k^\top (X_{1,T} - \hat{\Xi}_{0,T})^\top & \Sigma_k \end{bmatrix} \succeq 0, \quad (57)$$

is the nominal covariance LMI, and

$$\Delta G_k^\Sigma = \begin{bmatrix} 0_n & -\Delta\Xi_{0,T}S_k \\ -S_k^\top \Delta\Xi_{0,T}^\top & 0_n \end{bmatrix} \succeq 0 \quad (58)$$

is the perturbation to the covariance LMI. Next, we rewrite the perturbation matrix (58) as

$$\Delta G_k^\Sigma = \Theta^\top(S_k) \Delta\Xi_{0,T}^\top \Pi + \Pi^\top \Delta\Xi_{0,T} \Theta(S_k), \quad (59)$$

where $\Theta^\top(S_k) \triangleq [0_{n,T}; -S_k^\top]$ and $\Pi^\top = [I_n; 0_n]$. Finally, using [38], we may equivalently represent the uncertain LMI (56) as the following standard LMI

$$\begin{bmatrix} \lambda I_T & \rho(\delta) \Theta(S_k) \\ \rho(\delta) \Theta^\top(S_k) & \hat{G}_k^\Sigma(\Sigma_k, \Sigma_{k+1}, S_k) - \lambda \Pi^\top \Pi \end{bmatrix} \succeq 0, \quad (60)$$

in terms of the decision variables S_k, λ, Σ_k and Σ_{k+1} .

7 Parametric Uncertainty DD-CS

Rather than deriving bounds on $\Delta\Xi_{0,T}$ and designing for the worst-case disturbance (Sections 5 and 6), in this section we treat the parametric disturbances as stochastic ones obeying a known probability law $\xi_k^{(d)} \sim \mathcal{N}(0, \Sigma_\xi)$. For simplicity, in this section, we assume that D is known. We pose a data-driven density steering problem with probabilistic parametric uncertainty (PU-DD-DS). In other words, we treat the noise vectors $\xi_k^{(d)}$ as random variables with known probability distributions rather than attempting to estimate their specific realization. This approach has the advantage of allowing us to design a controller that is inherently robust to the entire distribution of possible noise realizations, rather than being robustly optimized for a single estimated noise instance.

7.1 Solution of the PU-DD-DS Problem

The dynamics of the actual system may be written as

$$\mathbf{x}_{k+1} = (\hat{A} + \Delta A(\Xi_{0,T}))\mathbf{x}_k + (\hat{B} + \Delta B(\Xi_{0,T}))\mathbf{u}_k + \xi_k, \quad (61)$$

where $[\hat{B} \ \hat{A}] = X_{1,T}S^\dagger$ is the CE estimate from (34), and $[\Delta B \ \Delta A] = -\Xi_{0,T}S^\dagger$ is the error from the CE design.

Partition the matrix $S^\dagger = [S_1, S_2] \in \mathbb{R}^{T \times (m+n)}$ in (26) into $S_1 \in \mathbb{R}^{T \times m}$ and $S_2 \in \mathbb{R}^{T \times n}$, and let $\tau_i \in \mathbb{R}^m$ denote the i th row of S_1 and $\sigma_i \in \mathbb{R}^n$ denote the i th row of S_2 . Since $\Xi_{0,T} = [\xi_0^{(d)}, \dots, \xi_{T-1}^{(d)}]$, the dynamics (61) can be written as

$$\mathbf{x}_{k+1} = \left(\hat{A} - \sum_{i=0}^{T-1} \xi_i^{(d)} \tau_i^\top \right) \mathbf{x}_k + \left(\hat{B} - \sum_{i=0}^{T-1} \xi_i^{(d)} \sigma_i^\top \right) \mathbf{u}_k + \xi_k. \quad (62)$$

Forming the outer product $\tilde{\mathbf{x}}_{k+1} \tilde{\mathbf{x}}_{k+1}^\top$, where $\tilde{\mathbf{x}}_k \triangleq \mathbf{x}_k - \mu_k$, and taking expectations yields the following expression for the covariance dynamics

$$\Sigma_{k+1} = \hat{\Sigma}_{x_{k+1}} + \Delta\Sigma_{x_{k+1}}^{(1)} + \Delta\Sigma_{x_{k+1}}^{(2)}, \quad (63)$$

where,

$$\begin{aligned} \hat{\Sigma}_{x_{k+1}} &= \hat{A}\Sigma_k\hat{A}^\top + \hat{A}\Sigma_{x_k, u_k}\hat{B}^\top \\ &\quad + \hat{B}\Sigma_{x_k, u_k}^\top\hat{A}^\top + \hat{B}\Sigma_{u_k}\hat{B}^\top + \Sigma_\xi, \end{aligned} \quad (64a)$$

$$\begin{aligned} \Delta\Sigma_{x_{k+1}}^{(1)} &= \sum_{i=0}^{T-1} \left(\sigma_i^\top \Sigma_k \sigma_i + \sigma_i^\top \Sigma_{x_k, u_k} \tau_i \right. \\ &\quad \left. + \tau_i^\top \Sigma_{x_k, u_k}^\top \sigma_i + \tau_i^\top \Sigma_{u_k} \tau_i \right) \Sigma_\xi, \end{aligned} \quad (64b)$$

$$\Delta\Sigma_{x_{k+1}}^{(2)} = \sum_{i=0}^{T-1} (\sigma_i^\top \mu_k + \tau_i^\top v_k)^2 \Sigma_\xi, \quad (64c)$$

where $\Sigma_{u_k} \triangleq \mathbb{E}[(\mathbf{u}_k - \mathbb{E}[\mathbf{u}_k])(\mathbf{u}_k - \mathbb{E}[\mathbf{u}_k])^\top]$ and $\Sigma_{x_k, u_k} \triangleq \mathbb{E}[(\mathbf{x}_k - \mathbb{E}[\mathbf{x}_k])(\mathbf{u}_k - \mathbb{E}[\mathbf{u}_k])^\top]$, and where we have used the fact that the noise follows an i.i.d normal distribution, and $\mathbb{E}[\xi_k \xi_j^\top] = \delta_{kj} I_n$ and $\mathbb{E}[\xi_i^{(d)} \xi_k^\top] = 0$. With the affine feedback controller (6) the covariance matrices are given by

$$\Sigma_{u_k} = K_k \Sigma_k K_k^\top, \quad \Sigma_{x_k, u_k} = \Sigma_k K_k^\top. \quad (65)$$

Note that with multiplicative uncertainties, the mean and covariance designs become coupled through the extra term (64c) entering equation (63); this is in contrast to the case of only additive disturbances, where the mean and covariance subproblems are decoupled.

For the uncertain mean dynamics, we have the following expression

$$\mu_{k+1} = (\hat{A} + \Delta A(\Xi_{0,T}))\mu_k + (\hat{B} + \Delta B(\Xi_{0,T}))v_k.$$

To proceed, we follow an CE approach, by neglecting the model error matrices, which yields the approximate mean dynamics

$$\mu_{k+1} \approx \hat{A}\mu_k + \hat{B}v_k, \quad (66)$$

with the additional caveat that the true terminal mean constraints may not be satisfied in practice, due to the model mismatch. Future work will investigate ways to incorporate robust design techniques, such as those in Section 6.1, in the context of PU-DD-DS. In summary, the PU-DD-DS problem is given as follows.

Problem 4 (PU-DD-DS) *Given the dataset $\mathbb{D}$ corresponding to the unknown linear system (1) with the nominal model $[\hat{B} \ \hat{A}] = X_{1,T}S^\dagger$, find the optimal control sequence $\{\mathbf{u}_k\}_{k=0}^{N-1}$ that minimizes the cost (5), subject to the (approximate) mean and (exact) covariance dynamics (66), (63), respectively, initial state (3) and terminal boundary conditions $\mu_N = \mu_f$, $\Sigma_N \preceq \Sigma_f$.*

Next, we provide a convex reformulation of Problem 4, which is similar to the derivation in [39]. To this end, first notice that the third term (64c) is quadratic in the decision variables $\{\mu_k, v_k\}$. Additionally, the control parameterization (65) results in a nonlinear program in the decision variables $\{\Sigma_{u_k}, K_k, \Sigma_{x_k}\}$. To remedy the former issue, we relax the equality constraint (64c) by introducing the new decision variables Σ_{ik}^Δ for each $i \in \llbracket T-1 \rrbracket_0$ and $k \in \llbracket N-1 \rrbracket_0$ such that

$$\Sigma_{ik}^\Delta \succeq (\sigma_i^\top \mu_k + \tau_i^\top v_k) \Sigma_\xi (\sigma_i^\top \mu_k + \tau_i^\top v_k), \quad (67)$$

which, using the Schur complement, can be recast as the LMI

$$\begin{bmatrix} \Sigma_{ik}^\Delta & (\sigma_i^\top \mu_k + \tau_i^\top v_k) I_n \\ (\sigma_i^\top \mu_k + \tau_i^\top v_k) I_n & \Sigma_\xi^{-1} \end{bmatrix} \succeq 0. \quad (68)$$

To address the nonlinear dependence in the control parameterization, and in a similar manner to the theory developed in Section 2.1, define the new decision variables $U_k \triangleq K_k \Sigma_{x_k} = \Sigma_{x_k, u_k}^\top$. The control covariance thus becomes $\Sigma_{u_k} = U_k \Sigma_{x_k}^{-1} U_k^\top$, which is still nonlinear in the decision variables. To this end, we relax this equality constraint by introducing yet another new decision variable Y_k such that

$$Y_k \succeq U_k \Sigma_{x_k}^{-1} U_k^\top \implies \begin{bmatrix} Y_k & U_k \\ U_k^\top & \Sigma_{x_k} \end{bmatrix} \succeq 0. \quad (69)$$

In summary, the relaxed PU-DD-DS problem is given by the SDP in (70)

$$\min_{\Sigma_k, U_k, Y_k, \Sigma_{ik}^\Delta} \sum_{k=0}^{N-1} (\mu_k^\top Q_k \mu_k + v_k^\top R_k v_k + \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k Y_k)), \quad (70a)$$

such that, for all $k \in \llbracket N-1 \rrbracket_0$, $i \in \llbracket T-1 \rrbracket_0$,

$$\mu_{k+1} = \hat{A}\mu_k + \hat{B}v_k \quad (70b)$$

$$\Sigma_{x_{k+1}} = \hat{\Sigma}_{x_{k+1}} + \sum_{i=0}^{T-1} \Sigma_{ik}^\Delta + \sum_{i=0}^{T-1} \left(\sigma_i^\top \Sigma_{x_k} \sigma_i + \sigma_i^\top U_k^\top \tau_i + \tau_i^\top U_k \sigma_i + \tau_i^\top Y_k \tau_i \right) \Sigma_\xi, \quad (70c)$$

$$\hat{\Sigma}_{x_{k+1}} = \hat{A} \Sigma_{x_k} \hat{A}^\top + \hat{A} U_k^\top \hat{B}^\top + \hat{B} U_k \hat{A}^\top + \hat{B} Y_k \hat{B}^\top + \Sigma_\xi, \quad (70d)$$

$$Y_k \succeq U_k \Sigma_{x_k}^{-1} U_k^\top, \quad (70e)$$

$$\Sigma_{ik}^\Delta \succeq (\sigma_i^\top \mu_k + \tau_i^\top v_k) \Sigma_\xi (\sigma_i^\top \mu_k + \tau_i^\top v_k)^\top, \quad (70f)$$

$$\Sigma_f - \Sigma_N \succeq 0, \quad (70g)$$

$$\mu_f - \mu_N = 0. \quad (70h)$$

8 Numerical Example

In this section, we compare all previous methods developed in the previous sections, using the following linear system

$$A = \frac{1}{2} \begin{bmatrix} 1 & -1 \\ 2 & 1 \end{bmatrix}, \quad B = I_2, \quad D = 0.1 I_2.$$

The initial state is normally distributed with mean $\mu_0 = [2, 10]^\top$ and covariance $\Sigma_0 = (1/3)^2 I_2$. The target terminal distribution has mean $\mu_f = [-1; -1]^\top$ and covariance $\Sigma_f = 0.25 \Sigma_0$. For the objective function, we set $Q = I_2$ and $R = 10 I_2$. For the robust mean design (Section 6.1) we set the terminal constraint set as $\mathcal{X}_f = \{x : |\mu_{f,i} - x_i| \leq 0.5, i = 1, 2\}$, where $x_i \in \mathbb{R}$ denotes the i th element of the state.

8.1 Noise Realization Estimation Analysis

We first start with an analysis of the estimation errors $\Delta \Xi_{0,T}$ resulting from the MLE problem (30). To this end, we ran the noise estimation procedure for 10^6 random trials, where the data was randomly generated for each trial from inputs $\mathbf{u}_k^{(d)} \sim U[-1, 1]$, initial state $\mathbf{x}_0^{(d)} \sim \mathcal{N}(0, I_2)$, and disturbances $\mathbf{w}_k^{(d)} \sim \mathcal{N}(0, DD^\top)$. In the full-length version [?] we provide additional plots of the individual noise estimation errors $\Delta \xi_k$ for varying sampling horizon lengths T . **Notably, we see that, indeed, as we gather more noisy data, the covariance of the estimation errors of each individual noise term converges to its exact value. The total magnitude of the**

joint estimation error, $\|\Delta\Xi_{0,T}\|$, as shown in Figure 2, however, does not converge to zero. This is explained by noting that while the individual estimation errors converge with more samples, the compounded error remains fixed due to the increasing number of elements to be estimated (see the full-length version [?] for a more formal derivation of the normed error distribution). Additionally, Table 1 shows the efficacy of the various upper bounds constructed in Section 5.

Table 1

Upper bounds for noise uncertainty set constructions compared to true error quantile ρ^* for various confidence levels δ .

δ	0.1	0.2	0.3	0.4	0.5
ρ^*	0.326	0.292	0.269	0.249	0.231
Tight	0.326	0.293	0.269	0.249	0.231
Loose ($T = 10$)	0.533	0.500	0.477	0.458	0.440
Loose ($T = 100$)	1.503	1.472	1.449	1.430	1.412
Loose ($T = 1,000$)	4.562	4.531	4.509	4.489	4.471

As mentioned in Section 5.1, the uncertainty set (41) based on the subspace decomposition of the singular joint Gaussian density (*Tight* row) has the *smallest* conservativeness of all the three alternatives, as it is almost exactly equal to the true quantile. Additionally, the original uncertainty set construction (*Loose* row) from our previous work in [20] (see Corollary 9) provides unnecessarily too loose bounds, and degrades rather quickly for large sampling horizons. We thus confirm that the uncertainty set constructed from the subspace density (39) and associated confidence set in Proposition 11 has the tightest overapproximation to the true quantile, and the original construction in [20] is the most conservative due to the spurious extra degrees of freedom. For the rest of the analysis, in this section, we choose to use the least conservative upper bound $\rho(\delta) = \|\Sigma_{\xi}^{1/2}\| Q_{\chi_{n(n+m)}^2}^{1/2} (1-\delta)$ from (41).




Fig. 2. Empirical distribution of norm of joint estimation errors $\Delta\Xi_{0,T}$, for varying sampling horizons T .

8.2 DD-MS Robustness Analysis

Next, we turn our attention to analyzing the effect of an unknown model on the resulting optimal mean trajectories. To this end, we compare the CE (i.e., using the




Fig. 3. Comparison of (left) DD-MS and (right) R-DD-MS terminal splashpoints for 500 randomized trials of data collected from the true model over a horizon $T = 30$, with empirical error distribution (bottom).

subspace predictor $[\hat{B} \ \hat{A}] = X_{1,T} \mathcal{S}^\dagger$) design with the robust design outlined in Section 6.1. We remark that the **MLE noise estimation scheme in Section 4.1 does not influence the nominal model estimate since $\hat{\Xi}_{0,T} \mathcal{S}^\dagger = X_{1,T} (I_T - \mathcal{S}^\dagger \mathcal{S}) \mathcal{S}^\dagger = 0$** . To compare the two designs, we ran a set of 500 trials with data generated randomly for each trial in a similar vein to the noise estimation study. We use the upper bound $\alpha = \rho(\delta)/\sigma_{\min}(\mathcal{S})$, where $\rho(\delta)$ as in (41), for the model estimation error, with $\delta = 0.1$ for the robust DD-MS design, and a data collection horizon $T = 30$.

Figure 3 shows the difference between the optimal mean trajectories for the two designs with respect to the *true* model $\{A, B\}$. Clearly, the robust design has better control over the dispersion of the terminal mean trajectories than that of the CE design (as done in [20]), as intended.

We conclude the discussion on the DD-MS design by focusing on the specific parameters involved, namely, the robustness level α and the terminal constraint box $\mathcal{X}_f$. **One common issue with robust MPC frameworks is the design of the terminal set in order to ensure recursive feasibility [40]**. In this work, however, we simply want a small enough terminal set, centered around the desired terminal mean that we can robustly steer the system trajectories to. As such, it is not guaranteed that the robust control problem will be feasible with a given terminal constraint set and model uncertainty $\alpha(\delta|\mathbb{D})$ under the dataset $\mathbb{D}$ and confidence level δ . Table 2 shows a quantitative comparison of the percentage of feasible solutions for varying levels of robustness and terminal set sizes. We see that, when the terminal box is large and

Table 2

Percentage of feasible solutions to R-DD-MC for varying model uncertainty δ and terminal constraint space size $|\mathcal{X}_f|$.

α	Box Width					
	0.5	0.4	0.3	0.2	0.1	0.05
0.05	100	100	100	100	100	99.4
0.1	100	100	100	100	96.2	50.6
0.15	100	100	100	100	43.8	17.4
0.2	100	100	100	71.0	13.4	8.00
0.25	100	99.4	42.4	8.40	6.20	4.60

when there is not much robustness, all problems are feasible. However, as we increase the level of robustness and decrease the size of the terminal box, many more problems become infeasible. Thus, the control designer has a trade-off between the desired accuracy of the nominal model with the level of precision in the terminal state.

8.3 DD-CS Robustness Analysis

Next, we proceed with the DD-CS analysis in the presence of noisy data, and subsequently with the synthesis of both the mean and covariance control designs. To this end, and similarly to the DD-MS analysis, we first begin with a study on the effect of robustness level ρ on the resulting terminal covariances. For reference, given the disturbance matrix $D = 0.1I_2$, confidence level $\delta = 0.1$, and sampling horizon $T = 30$, the associated MLE noise error uncertainty set is $\Delta^* = \{\|\Delta\Xi_{0,T}\| \leq \rho\}$, with $\rho = 0.3263$. Figure 5 shows the terminal covariances from the R-DD-CS design for varying levels of $\rho \in \{0, 0.3, 0.6, 0.9, 1.2, 1.5\}$, as well as for varying noise-to-precision ratio (NPR) $\sigma_\xi/\sigma_f \in \{0.6, 0.72, 0.84, 0.96\}$, evaluated on the true dynamics model $\{A, B, D\}$, where we assume $\Sigma_f = \sigma_f^2 I_2$ and $\Sigma_\xi = \sigma_\xi^2 I_2$. As the level of robustness to noise estimation errors increases, the terminal covariances become smaller when simulated on the true model because the optimal feedback gains anticipate more uncertainty than there is in reality. Note also that with no robustness (i.e., $\rho = 0$), the terminal covariances (black) do *not*, in general, satisfy the constraints (green) due to the non-zero estimation errors. Hence, robustness against estimation errors is essential for feasible covariance designs.

It is interesting to note that as the NPR increases, not only does the terminal covariance becomes larger, as expected, but also the R-DD-CS is more likely to become infeasible. According to [?], there is a theoretical lower bound on the achievable terminal covariance, given by $\Sigma_N \succeq D_{N-1}D_{N-1}^\top$. As the noise covariance σ_ξ increases and approaches σ_f , this lower bound becomes more constraining. At the same time, increasing the robustness

Table 3

Percentage of feasible solutions to R-DD-CS for varying noise estimation error bounds δ and disturbance variance σ_ξ^2 .

ρ	Noise variance σ_ξ^2			
	0.1 ²	0.12 ²	0.14 ²	0.16 ²
0.0	100	100	100	100
0.3	100	100	100	100
0.6	100	100	100	4.00
0.9	100	99.0	70.8	0.00
1.2	87.0	54.2	1.00	0.00
1.5	22.8	1.2	0.00	0.00


level ρ requires the covariance steering algorithm to aim for smaller values of Σ_N to ensure $\Sigma_N \preceq \Sigma_f$ holds for all bounded errors within ρ . However, when ρ becomes too large relative to the gap between σ_ξ and σ_f , the convex program (24) with robust constraints (60) becomes infeasible as it cannot reduce Σ_N below the theoretical lower bound while meeting the robustness constraints. Table 3 empirically verifies this relationship, showing the percentage of feasible solutions for each NPR across 500 random trials.

8.4 Closed-Loop Density Steering

Lastly, we combine the DD-MS and DD-CS designs and look at the resulting optimal trajectories. We choose a planning horizon of $N = 30$ and ran a set of 50 random trials for data collection, where, for each trial, we plot 10 Monte Carlo trajectories from randomly sampling the additive noise. Of all the designs, the R-DD-DS design performs the best in terms of achieving the closest terminal distribution to the desired one in the presence of noisy data. The PU-DD-DS design, which does *not* estimate the noise realizations but instead incorporates the distributional knowledge as multiplicative uncertainty, satisfies the terminal covariances for each trial, but is not robust against mean estimation errors, similar to vanilla DD-MS. Figuring out a way to robustly satisfy the terminal mean constraints in this parametric uncertainty framework is an interesting problem for future work.

9 Conclusion

We have presented a novel framework for data-driven stochastic optimal control for unknown linear systems with distributional boundary conditions, referred to as data-driven density steering (DD-DS). The proposed framework provides a comprehensive approach to the design of optimal controllers that steer the state distribution of an uncertain linear system to a desired termi-



(a) Non-robust design.

(b) Robust covariance design.

(c) Robust mean/covariance design.

Fig. 4. Terminal state mean and covariances of (a) certainty-equivalence design, (b) CE mean design + robust covariance design, and (c) robust mean + covariance designs.





Fig. 5. Terminal 3σ covariance ellipses for varying levels of robustness ρ and desired precision to noise ratio σ_f/σ_ξ .



(a) Model-based DS.

(b) Data-driven DS.

(c) Robust DD-DS.

(d) Parametric uncertainty DD-DS.

Fig. 6. Comparison of optimal closed-loop trajectories between (a) model-based, and (b-d) data-driven control design. The solid lines represent the mean trajectories and the ellipses represent the 3σ covariance ellipses for each randomized trial.

nal Gaussian distribution, using only input-state data collected from the actual system. By parameterizing the feedback gains directly in terms of the collected data, we have reformulated the data-driven distribution steering problem as an uncertain convex problem in terms of the unknown noise realizations. Using techniques from behavioral systems theory and statistical learning, we were able to develop tight, tractable uncertainty sets for the errors of these estimated noise realizations, which were subsequently used to formulate and solve robust data-driven extensions for the mean steering (DD-MS) and covariance steering (DD-CS) of the state, which guarantee high-probability constraint satisfaction under bounded estimation errors. Additionally, an alternative parametric uncertainty formulation (PU-DD-DS) was developed that treats model uncertainties probabilistically rather than deterministically. Extensive numerical studies demonstrated the efficacy of the proposed methods compared to certainty-equivalence and model-based approaches.

Acknowledgements

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

References

- [1] P. Dorato, A historical review of robust control, *IEEE Control Systems Magazine* 7 (2) (1987) 44–47.
- [2] W. Langson, I. Chrysoschoos, S. Raković, D. Mayne, Robust model predictive control using tubes, *Automatica* 40 (1) (2004) 125–133.
- [3] A. Mesbah, Stochastic model predictive control: An overview and perspectives for future research, *IEEE Control Systems Magazine* 36 (6) (2016) 30–44.
- [4] J. C. Willems, P. Rapisarda, I. Markovsky, B. L. De Moor, A note on persistency of excitation, *Systems and Control Letters* 54 (4) (2005) 325–329.
- [5] J. Coulson, J. Lygeros, F. Dörfler, Data-enabled predictive control: In the shallows of the DeePC, in: 18th European Control Conference (ECC), Naples, Italy, June 25–June 28, 2019, pp. 307–312.
- [6] J. Berberich, J. Köhler, M. A. Müller, F. Allgöwer, Data-driven model predictive control: closed-loop guarantees and experimental results, *Automatisierungstechnik* 69 (2021) 608 – 618.
- [7] C. De Persis, P. Tesi, Low-complexity learning of linear quadratic regulators from noisy data, *Automatica* 128 (2021) 109548.
- [8] F. Dörfler, P. Tesi, C. De Persis, On the certainty-equivalence approach to direct data-driven LQR design, *IEEE Transactions on Automatic Control* 68 (12) (2023) 7989–7996.
- [9] S. Kerz, J. Teutsch, T. Brüdigam, M. Leibold, D. Wollherr, Data-driven tube-based stochastic predictive control, *IEEE Open Journal of Control Systems* 2 (2023) 185–199.
- [10] G. Pan, R. Ou, T. Faulwasser, On a stochastic fundamental lemma and its use for data-driven optimal control, *IEEE Transactions on Automatic Control* 68 (10) (2023) 5922–5937.
- [11] T. Faulwasser, R. Ou, G. Pan, P. Schmitz, K. Worthmann, Behavioral theory for stochastic systems? A data-driven journey from Willems to Wiener and back again, *Annual Reviews in Control* 55 (2023) 92–117.
- [12] M. Yin, A. Iannelli, R. S. Smith, Stochastic data-driven predictive control: Regularization, estimation, and constraint tightening, 20th IFAC Symposium on System Identification (2024).
- [13] K. Okamoto, M. Goldshtein, P. Tsiotras, Optimal covariance control for stochastic systems under chance constraints, *IEEE Control Systems Letters* 2 (2) (2018) 266–271.
- [14] E. Bakolas, Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints, *Automatica* 91 (2018) 61–68.
- [15] Y. Chen, T. T. Georgiou, M. Pavon, Optimal steering of a linear stochastic system to a final probability distribution, part I, *IEEE Transactions on Automatic Control* 61 (5) (2016) 1158–1169.
- [16] J. Pilipovsky, P. Tsiotras, Covariance steering with optimal risk allocation, *IEEE Transactions on Aerospace and Electronic Systems* 57 (6) (2021) 3719–3733.
- [17] J. Ridderhof, J. Pilipovsky, 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.
- [18] J. Yin, Z. Zhang, E. Theodorou, P. Tsiotras, Trajectory distribution control for model predictive path integral control using covariance steering, in: International Conference on Robotics and Automation (ICRA), Philadelphia, PA, May 23–27, 2022, pp. 1478–1484.
- [19] J. Pilipovsky, P. Tsiotras, Data-driven covariance steering control design, in: 62nd IEEE Conference on Decision and Control (CDC), Marina Bay Sands, Singapore, Dec. 13–15, 2023, pp. 2610–2615.
- [20] J. Pilipovsky, P. Tsiotras, Data-driven robust covariance control for uncertain linear systems, in: Learning for Dynamics and Control, Oxford, UK, July 16–18, 2024.
- [21] M. ApS, The MOSEK optimization toolbox for MATLAB manual. Version 9.0. (2019).
URL <http://docs.mosek.com/9.0/toolbox/index.html>
- [22] I. Markovsky, F. Dörfler, Behavioral systems theory in data-driven analysis, signal processing, and control, *Annual Reviews in Control* 52 (2021) 42–64.
- [23] M. Rotulo, C. D. Persis, P. Tesi, Data-driven linear quadratic regulation via semidefinite programming, *IFAC-PapersOnLine* 53 (2) (2020) 3995–4000, 21st IFAC World Congress.
- [24] C. De Persis, P. Tesi, Formulas for data-driven control: Stabilization, optimality, and robustness, *IEEE Transactions on Automatic Control* 65 (3) (2020) 909–924.
- [25] J. Friedman, T. Hastie, R. Tibshirani, Sparse inverse covariance estimation with the graphical LASSO, *Biostatistics* 9 (3) (2007) 432–441.
- [26] V. A. Nguyen, D. Kuhn, P. Mohajerin Esfahani, Distributionally robust inverse covariance estimation: The Wasserstein shrinkage estimator, *Operations Research* 70 (1) (2022) 490–515.
- [27] F. Dörfler, J. Coulson, I. Markovsky, Bridging direct and indirect data-driven control formulations via regularizations

- and relaxations, *IEEE Transactions on Automatic Control* 68 (2) (2023) 883–897.
- [28] W. K. Newey, D. McFadden, Chapter 36 large sample estimation and hypothesis testing, in: *Handbook of Econometrics*, Vol. 4, Elsevier, 1994, pp. 2111–2245.
 - [29] J. Aitchison, S. D. Silvey, Maximum-likelihood estimation of parameters subject to restraints, *Annals of Mathematical Statistics* 29 (3) (1958) 813–828.
 - [30] C. J. Geyer, On the asymptotics of constrained M-estimation, *Annals of Statistics* 22 (4) (1993) 1993–2010.
 - [31] P. Stoica, B. C. Ng, On the Cramer-Rao bound under parametric constraints, *IEEE Signal Processing Letters* 5 (7) (1998) 177–179.
 - [32] P. C. Davis, C. H. Kapadia, J. T. Webster, Distribution of quadratic forms in the multivariate singular normal distribution, Technical Report 23, Department of Statistics, Southern Methodist University, oNR Contract N00014-68-A-0515, Project NR 042-260 (January 1969).
 - [33] B. K. Ghosh, Some monotonicity theorems for χ^2 , F , and t distributions with applications, *Journal of the Royal Statistical Society. Series B (Methodological)* 35 (3) (1973) 480–492.
 - [34] J. Anderson, J. C. Doyle, S. H. Low, N. Matni, System level synthesis, *Annual Reviews in Control* 47 (2019) 364–393.
 - [35] S. Dean, H. Mania, N. Matni, B. Recht, S. Tu, On the sample complexity of the linear quadratic regulator, *Foundations of Computational Mathematics* 20 (2020) 633–679.
 - [36] M. Bujarbaruah, U. Rosolia, Y. R. Stürz, F. Borrelli, A simple robust MPC for linear systems with parametric and additive uncertainty, in: *American Control Conference (ACC)*, New Orleans, LA, May 25–28, 2021, pp. 2108–2113.
 - [37] S. Chen, H. Wang, M. Morari, V. M. Preciado, N. Matni, Robust closed-loop model predictive control via system level synthesis, in: *59th IEEE Conference on Decision and Control (CDC)*, Jeju Island, Republic of Korea, Dec. 14–18, 2020, pp. 2152–2159.
 - [38] A. Ben-Tal, L. E. Ghaoui, A. Nemirovski, *Robust Optimization*, Vol. 28 of Princeton Series in Applied Mathematics, Princeton University Press, 2009.
 - [39] J. W. Knaup, P. Tsiotras, Computationally efficient covariance steering for systems subject to parametric disturbances and chance constraints, in: *62nd IEEE Conference on Decision and Control (CDC)*, Marina Bay Sands, Singapore, Dec. 13–15, 2023, pp. 1796–1801.
 - [40] J. Rossiter, *Model-Based Predictive Control: A Practical Approach*, CRC Press LLC, 2004.