Gaussian behaviors: representations and data-driven control

András Sasfi, Ivan Markovsky, Alberto Padoan, Florian Dörfler

Abstract—We propose a modeling framework for stochastic systems based on Gaussian processes. Finite-length trajectories of the system are modeled as random vectors from a Gaussian distribution, which we call a Gaussian behavior. The proposed model naturally quantifies the uncertainty in the trajectories, yet it is simple enough to allow for tractable formulations. We relate the proposed model to existing descriptions of dynamical systems including deterministic and stochastic behaviors, and linear time-invariant (LTI) state-space models with Gaussian process and measurement noise. Gaussian behaviors can be estimated directly from observed data as the empirical sample covariance under the assumption that the measured trajectories are from independent experiments. The distribution of future outputs conditioned on inputs and past outputs provides a predictive model that can be incorporated in predictive control frameworks. We show that subspace predictive control (SPC) is a certainty-equivalence control formulation with the estimated Gaussian behavior. Furthermore, the regularized data-enabled predictive control (DeePC) method is shown to be a distributionally optimistic formulation that optimistically accounts for uncertainty in the Gaussian behavior. To mitigate the excessive optimism of DeePC, we propose a novel distributionally robust control formulation, and provide a convex reformulation allowing for efficient implementation.

I. INTRODUCTION

Recent data-driven control methods based on the behavioral approach [1]–[7] have gained significant attention. These formulations rely on behavioral systems theory that treats systems as sets of trajectories, and allows to represent linear time-invariant (LTI) systems directly with data [8], [9]. The methods exploit this data representation and typically add regularization to the problem [10], [11] for a posteriori robustification. The resulting formulations can achieve comparable or even superior performance compared to classical methods consisting of an identification and a control step, especially in challenging scenarios involving nonlinear effects and large amount of uncertainty [10], [12]. Existing works explain this observation by the robustifing effect of regularization, yet they rely on the inherently deterministic behavioral description [2], [3], [13]. On the other hand, only a few works propose frameworks that start directly from a stochastic system model. In fact, the survey [12] points out that a major open problem in this field is precisely the extension of data representations to the stochastic domain.

I. Markovsky is with the Catalan Institution for Research and Advanced Studies, 08010 Barcelona, Spain, and also with the International Centre for Numerical Methods in Engineering, 08034 Barcelona, Spain (e-mail: ivan.markovsky@cimne.upc.edu).

Stochastic extensions to behavioral systems theory have been defined bottom up in the literature [14], [15]. However, the existing works provide general and abstract (and thus also blunt) perspectives that hinder the practical applicability of the frameworks. Stochastic behaviors have also been modeled in the literature using polynomial chaos expansions [16], [17]. However, the complexity of the resulting methods grow with the order of the expansion, limiting scalability. Recently, a stochastic interpretation of data-driven control formulations was proposed in [18], highlighting that regularization accounts for the uncertainty in a linear model estimated from data. In this work, we take a different approach, and propose a new bottom-up stochastic modeling framework that admits a data representation and leads to tractable control formulations which can be solved efficiently.

The contributions of this work are the following. First, in Section III, we propose a stochastic modeling framework based on Gaussian processes that enables us to model trajectories as normally distributed random vectors. We define the distribution of trajectories to be a *Gaussian behavior*, characterized by its mean and covariance, which inherently captures uncertainty. Gaussian behaviors give rise to predictive models given by the conditional distribution, which can be calculated readily in the proposed framework. Furthermore, in a data-driven context, Gaussian behaviors can be estimated directly as the empirical sample covariance, which is also used as a system representation in [19]. We show that Gaussian behaviors are a special class of stochastic behaviors. Furthermore, the deterministic behavioral description of LTI systems as subspaces is captured by the proposed definition with zero mean and singular covariance matrix. Finally, stochastic LTI systems given by state-space equations with normally distributed initial state, process and measurement noise are also special cases of the proposed definition by imposing additional structure on the covariance. Additionally, Gaussian behaviors are closely connected to covariance steering problems [20]-[23] that usually consider linear stochastic systems with a normally distributed initial state and aim at finding a feedback controller that achieves a desired distribution for the closed-loop output.

Second, we interpret existing data-driven control formulations in the proposed framework and derive a new robust method in Section IV. In particular, subspace predictive control (SPC) is a certainty-equivalence formulation in this framework that uses the sample covariance to predict the future outputs. To account for inaccuracies in the predictive model, we optimize the predictive distribution within an uncertainty set centered around the estimate. This results in a *distributionally optimistic* formulation, which is shown

A. Sasfi and F. Dörfler are with the Department of Information Technology and Electrical Engineering, ETH Zürich, 8092 Zürich, Switzerland (e-mail: {asasfi, doerfler}@control.ee.ethz.ch).

A. Padoan is with the Department of Electrical & Computer Engineering, University of British Columbia. (e-mail: apadoan@ece.ubc.ca).

to be equivalent to the regularized data-enabled predictive control (DeePC) approach. In this method, the distribution that minimizes the control cost is selected from the uncertainty set. In contrast, we propose a distributionally robust formulation that optimizes the control input under the worstcase mismatch between the true and estimated distributions. The resulting min-max problem can be reformulated as a convex problem, enabling efficient computation.

II. PRELIMINARIES

A. Notation

The image and kernel of a matrix A are denoted by $\operatorname{im}(A)$ and $\operatorname{ker}(A)$, respectively. Furthermore, A^{\dagger} denotes the Moore–Penrose inverse of A. The set of symmetric positive (semi) definite matrices is denoted by $\mathbb{S}_{>0}$ ($\mathbb{S}_{\geq 0}$). The identity matrix of dimension $n \times n$ is I_n . Furthermore, the (semi-) norm of x weighted by a matrix $M \in \mathbb{S}_{>0}$ ($\mathbb{S}_{\geq 0}$) is denoted by $||x||_M = x^{\top}Mx$.

B. Stochastic processes

Classically, a stochastic process (or, briefly, process) is defined as a family of random variables $\{w_t\}_{t\in\mathbb{T}}$, where w_t is a random variable for each time t in the time range \mathbb{T} [24]. Given their finite-dimensional distributions, the entire process is uniquely characterized via Kolmogorov's extension theorem [24].

A stochastic process $\{w_t\}_{t\in\mathbb{T}}$ is wide-sense stationary [25, Sec. 10.1] if its mean is constant, i.e.,

$$\mathbb{E}[w_t] = \mu$$
, for all $t \in \mathbb{T}$,

for some vector μ , and the covariance between w_{t_i} and w_{t_j} depends only on the time shift $t_i - t_j$, i.e.,

$$\operatorname{Cov}(w_{t_i}, w_{t_j}) = K(t_i - t_j), \text{ for all } t_i, t_j \in \mathbb{T},$$

for some matrix-valued covariance function $K(\cdot)$.

A stochastic process is said to be a *Gaussian process* if, for any finite set of time instants

$$T = \{t_1, t_2, \dots, t_k\} \subset \mathbb{T},$$

the joint probability distribution of the corresponding random variables $\{w_t\}_{t \in T}$ is Gaussian, that is,

$$\{w_t\}_{t\in T} \sim \mathcal{N}(\mu_T, K_T),$$

where $\mu_T = [\mu(t_1) \dots \mu(t_k)]^\top$ is the mean vector, and K_T is the $k \times k$ covariance matrix with entries

$$K_T(i,j) = \mathbb{E}[(w_{t_i} - \mu(t_i))(w_{t_j} - \mu(t_j))^{\top}]$$
 for $t_i, t_j \in T$.

C. Behavioral systems theory for discrete-time LTI systems

Behavioral systems theory abstractly defines systems as a set of trajectories, called the *behavior*, regardless of its representation in terms of particular equations. For LTI systems, the behavior is a shift-invariant subspace of the set of all possible trajectories [11]. In this deterministic context, with a slight abuse of notation, we denote the realization of the random vector by the same symbol, i.e., $w_t \in \mathbb{R}^q$. We focus on system trajectories of finite length L > 0, which we denote by $w = [w_t^{\top} \ w_{t+1}^{\top} \ \dots \ w_{t+L-1}^{\top}]^{\top}$. The restricted behavior, denoted by \mathcal{B}_L , is the set of all system trajectories of length L > 0, and it is defined as

$$\mathcal{B}_L = \{ w \in \mathbb{R}^{qL} \mid w \text{ is a length-} L \text{ trajectory of the system} \}.$$

For LTI systems and for large enough L, $\mathcal{B}_L \subseteq \mathbb{R}^{qL}$ is a subspace of dimension d = mL + n, where n is the order of a minimal realization of the system and m is the number of inputs. The behavior has multiple representations, two of which are reviewed below.

1) State-space representation: Without loss of generality, each trajectory $w_t \in \mathbb{R}^q$ of an LTI system can be partitioned into inputs $u_t \in \mathbb{R}^m$ and outputs $y_t \in \mathbb{R}^{q-m}$. Furthermore, every discrete-time LTI system is described by the statespace representation

$$x_{t+1} = Ax_t + Bu_t,$$

$$y_t = Cx_t + Du_t,$$

with state $x_t \in \mathbb{R}^n$. By convention, we denote the input and output trajectories of length L by u and y. By linearity, the output trajectory can be expressed as a linear combination of the initial state and the inputs

$$y = \mathcal{O}_L x_{\text{ini}} + \mathcal{T}_L u,$$

where x_{ini} is the initial state, while \mathcal{O}_L and \mathcal{T}_L are the extended observability and convolution matrices, defined as

$$\mathcal{O}_{L} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{L-1} \end{bmatrix}, \ \mathcal{T}_{L} = \begin{bmatrix} D & 0 & 0 & \dots & 0 \\ CB & D & 0 & \dots & 0 \\ CAB & CB & D & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ CA^{L-2}B & \dots & \dots & CB & D \end{bmatrix}.$$

A basis for the subspace \mathcal{B}_L can be readily constructed from \mathcal{O}_L and \mathcal{T}_L , giving

$$\mathcal{B}_L = \operatorname{im} \left(P \cdot \begin{bmatrix} I_{mL} & 0 \\ \mathcal{T}_L & \mathcal{O}_L \end{bmatrix} \right)$$

where P is a suitable permutation matrix.

2) Data representation: Assume that measurements of D > qL trajectories w^i , $i \in \{1, \ldots, D\}$, from the system are available, and define the data matrix $W = [w^1 \ldots w^D]$. According to [26], if the inputs are collectively persistently exciting, W has rank mL + n, and it spans the restricted behavior, i.e.,

$$\mathcal{B}_L = \operatorname{im}(W).$$

Therefore, sufficiently rich data directly provide representations of deterministic LTI behaviors.

D. Stochastic LTI systems

Consider a stochastic LTI system described by the statespace equations

$$x_{t+1} = Ax_t + Bu_t + \xi_t,$$

$$y_t = Cx_t + Du_t + \eta_t,$$
(1)

where $\xi_t \in \mathbb{R}^n$ and $\eta_t \in \mathbb{R}^p$ are the process and measurement noise, respectively. Assume that both ξ_t and η_t

are independent and identically distributed with zero mean and covariance Σ_{ξ} and Σ_{η} , respectively. Then the output trajectory be expressed as a linear combination of the initial state, the input trajectories, and noise realizations

$$y = \mathcal{O}_L x_{\rm ini} + \mathcal{T}_L u + \mathcal{T}_L^{\varsigma} \xi + \eta,$$

where ξ and η are the length-L trajectories of process and measurement noise, and \mathcal{T}_L^{ξ} is a convolution matrix defined as \mathcal{T}_L with B = I, D = 0.

E. Data-driven predictive control formulations

Various data-driven predictive control formulations exploit the data representation of the behavior [1], [4], [7], [27]. To ensure compatibility with the most recent measurements, the initial part of the trajectory $w_{\text{ini}} = [u_{\text{ini}}^{\top} y_{\text{ini}}^{\top}]^{\top}$ of length $(L - L_{\text{f}})$ is considered to be given, and the future part $w_{\text{f}} = \sum_{i=1}^{T} |u_{\text{ini}}|^{\top}$ $[u_{\rm f}^{\top} y_{\rm f}^{\top}]^{\top}$ of length $L_{\rm f}$ is a free variable to be optimized. The whole trajectory is $w = [w_{\text{ini}}^{\top} \ w_{\text{f}}^{\top}]^{\top}$. The control objective is often posed as minimizing the quadratic cost

$$c_{\rm ctrl}(w_{\rm f}) = (w_{\rm f} - w_{\rm ref})^{\top} \begin{bmatrix} R & 0\\ 0 & Q \end{bmatrix} (w_{\rm f} - w_{\rm ref}),$$

where $w_{\rm ref}$ is the reference trajectory, and $R \in \mathbb{S}_{>0}$, $Q \in \mathbb{S}_{\geq 0}$ are weight matrices for the entire future input and output trajectories, respectively. The future trajectory $w_{\rm f} = [u_{\rm f}^{\top} \ y_{\rm f}^{\top}]^{\top}$ is often restricted to a non-empty, closed and convex constraint set $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$.

We review two data-driven predictive control formulations below. Let us partition the rows of the data matrix W as $W = [W_{\rm p}^{\top}, U_{\rm f}^{\top}, Y_{\rm f}^{\top}]^{\top}$, where the block rows correspond to $u_{\rm ini}, y_{\rm ini}, u_{\rm f}$ and $y_{\rm f}$, respectively. The SPC formulation from [11], [27] is then given as

$$\min_{w_{\rm f} \in \mathcal{W}} c_{\rm ctrl}(w_{\rm f})$$
s.t. $y_{\rm f} = Y_{\rm f} \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix}^{\dagger} \begin{bmatrix} w_{\rm ini} \\ u_{\rm f} \end{bmatrix}.$

$$(2)$$

Furthermore, the DeePC method from [10] is based on

$$\min_{w_{f} \in \mathcal{W}, g \in \mathbb{R}^{D}} c_{ctrl}(w_{f}) + \lambda_{g} \cdot h(g)$$
s.t.
$$\begin{bmatrix}
w_{ini} \\
u_{f} \\
y_{f}
\end{bmatrix} = \begin{bmatrix}
W_{p} \\
U_{f} \\
Y_{f}
\end{bmatrix} g,$$
(3)

where h(g) is a regularizer and $\lambda_g \ge 0$ is the regularization weight. Typical choices for the function h(g) are the (squared) 1-norm, 2-norm, and suitably projected 2-norm of g, see [10], [11].

III. GAUSSIAN BEHAVIORS

We propose a pragmatic yet foundational modeling framework for stochastic dynamical systems based on Gaussian processes enabling tractable control formulations. In particular, we model trajectories of (discrete-time) dynamical systems as stochastic processes $\{w_t\}_{t\in\mathbb{Z}}$. Similarly to behavioral systems theory, we consider models that describe entire trajectories of the system. Furthermore, to ensure that the proposed model is amenable to computation, we focus on finite-length trajectories. For the rest of the paper, we restrict our attention to a special class of systems, whose trajectories are wide-sense stationary Gaussian processes of q-variate random vectors w_t . Note that wide-sense stationarity is closely connected to time-invariance of the dynamical system [28].

We consider each trajectory of length L > 0 to be a random vector containing L consecutive time instants of the random process $\{w_t\}_{t\in\mathbb{Z}}$ denoted by

$$w = [w_t^\top \ w_{t+1}^\top \ \dots \ w_{t+L-1}^\top]^\top.$$

Due to the stationarity of the process, all trajectories of length L follow the same Gaussian distribution

Definition 1: A (finite-length) Gaussian behavior is a random length-L trajectory $w \in \mathbb{R}^{qL}$ distributed as a multivariate Gaussian, that is,

$$w \sim \mathcal{N}(\mu_w, \Sigma_w),$$

with mean $\mu_w \in \mathbb{R}^{qL}$ and covariance matrix $\Sigma_w \in \mathbb{S}_{\geq 0}$. The covariance Σ_w captures the correlation between the components of w. This model naturally quantifies uncertainty in the trajectories and enables tractable methods to perform downstream tasks, such as prediction or control. Furthermore, the mean and covariance can be readily identified from data, leading to a data representation of the system based on empirical sample covariances (see Lemma 1), which was recently proposed also in [19]. Finally, this model is consistent with existing modeling approaches for stochastic LTI systems, as discussed next.

Remark 1: Gaussian processes are widely used in machine learning to model the relationship between a regressor vector and a dependent variable [29]. Gaussian processes in this context are collections of (random) dependent variables indexed by the regressor vector. The covariance (or kernel) function defines the covariance between two dependent variables as a function of the corresponding regressors. In this work, on the other hand, the random variables are the inputs and outputs of a dynamical system indexed by time. Since we work with stationary processes, the covariance function $K(\cdot)$ only depends on the time shift $t_i - t_j$. Furthermore, as we consider length L trajectories, we are only interested in the first L values of the covariance function $K(0), \ldots, K(L-1)$, which constitute the covariance matrix Σ_w in Definition 1. These function values can be identified from measured trajectories of length L as described later.

A. Relation to existing system descriptions

1) Stochastic behaviors: The notion of stochastic behaviors is formalized in [15] as a extension of the deterministic notion of dynamical systems. It incorporates a probability space structure into the set of admissible system trajectories and it includes as special cases both the deterministic notion of dynamical system-viewed as a set of trajectories-and the traditional notion of stochastic process [24]. In this context, a stochastic dynamical system is defined as a quadruple

$$(\mathbb{T}, \mathbb{W}, \mathcal{F}, P),$$

where \mathbb{T} is the *time axis* (e.g., $\mathbb{T} = \mathbb{Z}$, $\mathbb{T} = \mathbb{R}$), \mathbb{W} is the *signal set* (e.g., $\mathbb{W} = \mathbb{R}^q$ or $\mathbb{W} = \mathbb{C}^q$), \mathcal{F} is a σ -algebra of events, that is, measurable subsets of the space of trajectories \mathbb{W}^T , and P is a *probability measure* on $(\mathbb{W}^T, \mathcal{F})$, assigning probabilities to events. Opposed to the traditional notion of stochastic processes, the behavioral approach [8] shifts the focus from individual random variables to the structure of admissible trajectories in \mathbb{W}^T , where P assigns a probability to entire trajectories. This interpretation, further developed in [15], highlights that stochastic dynamical systems generalize classical stochastic processes by emphasizing event spaces rather than generating random variables.

Our approach of modeling stochastic systems as Gaussian processes fits in this framework as a special case with $\mathbb{T} = \mathbb{Z}$ and $\mathbb{W} = \mathbb{R}^q$. Furthermore, the σ -algebra of events \mathcal{F} is the Borel σ -algebra on $(\mathbb{R}^q)^{\mathbb{Z}}$, generated by the open sets under the product topology (i.e., the topology of pointwise convergence). We use the standard Borel σ -algebra in this work, since it suffices for our purposes, while we leave the investigation of more general structures to future work. The probability measure P is the Gaussian process in the generality of the definitions given in [14], [15], our definition of Gaussian behaviors in Definition 1 allows us to propose actionable, i.e., simple and tractable, identification and control techniques.

2) *Deterministic LTI behaviors:* Deterministic restricted LTI behaviors can be regarded as special cases of the Gaussian behavior in Definition 1.

Proposition 1: Consider a Gaussian random vector $w \in \mathbb{R}^{qL}$ with mean $\mu_w = 0$ and covariance matrix $\Sigma_w \in \mathbb{S}_{\geq 0}$. Assume $\Sigma_w = BB^{\top}$ for some full column rank matrix $B \in \mathbb{R}^{qL \times d}$. Then, $\operatorname{rank}(\Sigma_w) = d$ and $w \in \operatorname{im}(B)$ with probability one.

Proof: Note that $w \sim \mathcal{N}(0, BB^{\top})$ can be written as w = Bx, with some $x \sim \mathcal{N}(0, I_d)$. Thus, w belongs to the image of B with probability 1 (cf., e.g., [30, Sec. 2.5.4]).

Proposition 1 shows that the Gaussian behavior with a suitably chosen mean and singular covariance matrix restricts the trajectories to $\mathcal{B}_L = \operatorname{im}(B)$. Nevertheless, a Gaussian behavior contains more information about the finite-length trajectories, namely, their probability distribution on \mathcal{B}_L . Note, however, that if we fix d "free" variables (the inputs and n output components) in w, the trajectory becomes deterministic. Therefore, by fixing some of the variables, we recover the deterministic relation between the signal components that is given by \mathcal{B}_L . Note that for a general $\mu_w \neq 0$, w belongs to a hyperplane instead of a subspace, and hence, the Gaussian behavior would describe an affine system [31].

3) Stochastic LTI systems: Now we show that the statespace representation of stochastic LTI systems with normally distributed initial state can also be recovered with a specific choice of the covariance matrix.

Proposition 2: Consider the stochastic LTI system (1). Assume that $w = [u^{\top} y^{\top}]^{\top}$ satisfies (1), with $x_{\text{ini}} \sim \mathcal{N}(0, \Sigma_{\text{ini}}), u \sim \mathcal{N}(0, \Sigma_u), \xi \sim \mathcal{N}(0, \Sigma_{\xi}), \text{ and } \eta \sim$ $\mathcal{N}(0, \Sigma_{\eta})$. Assume further that x_{ini} and u are independent. Then $w \sim \mathcal{N}(0, \Sigma_w)$ with covariance

$$\Sigma_w = \begin{bmatrix} \Sigma_u & \Sigma_u \mathcal{T}_L^\top \\ \mathcal{T}_L \Sigma_u & \mathcal{O}_L \Sigma_{\rm ini} \mathcal{O}_L^\top + \mathcal{T}_L \Sigma_u \mathcal{T}_L^\top + \mathcal{T}_L^\xi \Sigma_\xi (\mathcal{T}_L^\xi)^\top + \Sigma_\eta \end{bmatrix}.$$

Proof: The claim follows by writing the trajectory as

$$w = \begin{bmatrix} u \\ y \end{bmatrix} = \begin{bmatrix} 0 & I_{mL} & 0 & 0 \\ \mathcal{O}_L & \mathcal{T}_L & \mathcal{T}_L^{\xi} & I \end{bmatrix} \begin{bmatrix} x_{\text{ini}} \\ u \\ \xi \\ \eta \end{bmatrix}$$

where $[x_{\text{ini}}^{\top} \ u^{\top} \ \xi^{\top} \ \eta^{\top}]^{\top} \sim \mathcal{N}(0, \text{diag}(\Sigma_{\text{ini}}, \Sigma_u, \Sigma_{\xi}, \Sigma_{\xi})),$ with $\text{diag}(\Sigma_{\text{ini}}, \Sigma_u, \Sigma_{\xi}, \Sigma_{\eta})$ denoting a block diagonal matrix with blocks $\Sigma_{\text{ini}}, \Sigma_u, \Sigma_{\xi}$, and Σ_{η} .

Proposition 2 shows that the covariance Σ_w in Definition 1 can be constructed from the stochastic state-space representation (1) and the covariances of u, x_{ini} , ξ , and η . Note that the initial state x_{ini} is assumed to be normally distributed, which is a common assumption in the covariance steering literature [20]–[23] as well. The mean μ_w is set to zero, implying that the output must be zero mean, if x_{ini}, u, ξ, η are also zero mean. In contrast to the deterministic behaviors in Proposition 1, the covariance is always non-singular if Σ_u and Σ_η are non-singular. Thus, we assume the following throughout the rest of the paper.

Assumption 1: Assume $\mu_w = 0$ and Σ_w is non-singular.

B. Prediction using Gaussian behaviors

The Gaussian behavior in Definition 1 describes the distribution of the whole trajectory w. In many downstream tasks, some parts of the trajectory are already observed, or are fixed priori. For example, in predictive control, the initial part w_{ini} of the trajectory is observed, and the future inputs u_f are decision variables that we can fix. The probability of the future outputs y_f given w_{ini} and u_f is expressed by the conditional distribution and can be easily computed for Gaussians. Recall Assumption 1 and partition Σ_w into block corresponding to $[w_{ini}^{\top}, u_f^{\top}]^{\top}$ and y_f

$$\Sigma_w = \begin{bmatrix} \Sigma_w^{11} & \Sigma_w^{12} \\ \Sigma_w^{21} & \Sigma_w^{22} \end{bmatrix}$$

The distribution of $y_{\rm f}$ conditioned on $w_{\rm ini}$ and $u_{\rm f}$ is given as $y_{\rm f}|u_{\rm f}, w_{\rm ini} \sim \mathcal{N}(\mu_{\rm pred}, \Sigma_{\rm pred})$, with

$$\mu_{\rm pred} = \Sigma_w^{21} (\Sigma_w^{11})^{-1} \begin{bmatrix} w_{\rm ini} \\ u_{\rm f} \end{bmatrix}, \qquad (4)$$

$$\Sigma_{\text{pred}} = \Sigma_w^{22} - \Sigma_w^{21} (\Sigma_w^{11})^{-1} \Sigma_w^{12}.$$
 (5)

We call $\mathcal{N}(\mu_{\text{pred}}, \Sigma_{\text{pred}})$ the *predictive distribution*. Note that the mean of the predictive distribution is linear in w_{ini} and u_{f} , while the covariance is independent of them.

C. Identifying Gaussian behaviors

An important advantage of Gaussian behaviors is that they can be readily identified using the sample covariance of data $\hat{\Sigma}_w = (1/D) \cdot WW^{\top}$. It is worthwhile noting that in the deterministic LTI setting, the sample covariance serves as the data representation of the behavior, since $\operatorname{im}(\hat{\Sigma}_w) =$ $\operatorname{im}(W) = \mathcal{B}_L$. Consequently, existing data-driven methods that exploit the image of W can be directly interpreted in the proposed framework using the image of $\hat{\Sigma}_w$. Analogously to deterministic behaviors in Section II-C, one might say that the sample covariance is the data representation of Gaussian behaviors. In practice, however, only finite data is available from the system, and therefore the covariance estimate $\hat{\Sigma}_w$ of the true covariance Σ_w is usually inaccurate.

Let us pose the following assumption on the available data. Assumption 2: Assume w^i , $i \in \{1, ..., D\}$ are independent samples from $\mathcal{N}(0, \Sigma_w)$, and W is full row rank. The rank condition on W can be interpreted as a persistence of excitation condition analogously to the generalized persistency of excitation condition for deterministic LTI behaviors [11], which is equivalent to identifiability. Furthermore, the samples are assumed to be independent to ensure that the sample covariance can be used to estimate Σ_w .

Under Assumption 2, the estimate $\hat{\Sigma}_w$ of Σ_w that maximizes the likelihood of observing the data is given by the sample covariance, as formalized below.

Lemma 1: Let Assumptions 1 and 2 hold. Then

$$\hat{\Sigma}_w = \frac{1}{D} W W^\top = \arg \max_{\Sigma_w \in \mathbb{S}_{>0}} \prod_{i=1}^D p_w(w^i)$$

where $p_w(w^i)$ denotes the probability density function of $\mathcal{N}(0, \Sigma_w)$ evaluated at w^i .

Proof: The optimizer $\hat{\Sigma}_w$ can be expressed by taking the logarithm of the cost and using the first order optimality condition. See, e.g., [32, Thm. 3.1.5] for the derivation.

Note that it is possible to relax Assumption 2 and allow linear correlation between the samples w^i . Then, the maximum likelihood estimator of Σ_w is given as a weighted sample covariance [33], [34]. Furthermore, finite-sample bounds on the estimation error are available in the literature [35]. From $\hat{\Sigma}_w$, the estimate of the conditional distribution $y_f|u_f, w_{ini}$ can be derived as follows.

Lemma 2: For $\hat{\Sigma}_w = (1/D) \cdot WW^{\top}$ as in Lemma 1, the mean and covariance of the predictive distribution $\mathcal{N}(\mu_{\text{pred}}, \Sigma_{\text{pred}})$ can be expressed as

$$\hat{\mu}_{\text{pred}} = Y_{\text{f}} \begin{bmatrix} W_{\text{p}} \\ U_{\text{f}} \end{bmatrix}^{\dagger} \begin{bmatrix} w_{\text{ini}} \\ u_{\text{f}} \end{bmatrix}, \qquad (6)$$

$$\hat{\Sigma}_{\text{pred}} = \frac{1}{D} Y_{\text{f}} \left(I - \Pi \right) Y_{\text{f}}^{\top}, \tag{7}$$

where $\Pi := \begin{bmatrix} W_{\mathrm{p}} \\ U_{\mathrm{f}} \end{bmatrix}^{\dagger} \begin{bmatrix} W_{\mathrm{p}} \\ U_{\mathrm{f}} \end{bmatrix}$.

Proof: The formulas follow from (4) and (5) with

$$\Sigma_w^{11} = \frac{1}{D} \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix} \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix}^{\top}, \ \Sigma_w^{12} = \frac{1}{D} \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix} Y_{\rm f}^{\top},$$

and $\Sigma_w^{21} = (\Sigma_w^{12})^\top$, $\Sigma_w^{22} = \frac{1}{D} Y_{\mathrm{f}} Y_{\mathrm{f}}^\top$.

IV. DATA-DRIVEN CONTROL WITH GAUSSIAN BEHAVIORS

We propose stochastic predictive control formulations relying on our definition of Gaussian behaviors, and show that they are closely related to existing control formulations described in Section II-E. In the proposed control formulations, the future inputs u_f are considered to be deterministic decision variables, i.e., the trajectory is partially realized. The future outputs are random variables that depend on u_f and w_{ini} , and they follow the predictive probability distribution from Section III-B. Therefore, we define the control cost as

$$\tilde{c}_{\mathrm{ctrl}}(u_{\mathrm{f}}) := c_{\mathrm{ctrl}}([u_{\mathrm{f}}^{\top} (y_{\mathrm{f}}|u_{\mathrm{f}}, w_{\mathrm{ini}})^{\top}]^{\top}),$$

which is a family of random variables parametrized by $u_{\rm f}$. The proposed control problems aim at finding $u_{\rm f}$ that minimizes the expected control cost given $w_{\rm ini}$. As the $y_{\rm f}$ is a random variable in this setting, it is natural to include chance constraints on it in the following control formulations. However, we do not impose such constraints to simplify the presentation. The following formulations differ in how the expectation of the cost is calculated.

Remark 2: Aside from predictive control formulations, our definition of Gaussian dynamical systems fits naturally with covariance steering problems [20]–[23] and the covariance parametrization in [19]. In covariance steering, the feedback controller that achieves a desired distribution for the closed-loop output usually consists of a linear state feedback gain and a feedforward term. In our setup, w_{ini} can be considered as a (non-minimal) initial state. By tuning the feedback gain, the cross-correlation between u_f and w_{ini} in the covariance Σ_w can be shaped. Therefore, covariance of the predictive distribution $\mathcal{N}(\mu_{pred}, \Sigma_{pred})$ can be influenced by the controller design. Furthermore, the mean μ_{pred} can be assigned directly by a feedforward term.

A. Certainty-equivalence control

We propose a control formulation that calculates the input by minimizing the expected control cost. The expectation is calculated using the estimated predictive distribution from Section III-C. This is a certainty-equivalence approach, as we do not take into account the inaccuracy of the estimated model. Let $\hat{\mu}_{pred}$ and $\hat{\Sigma}_{pred}$ be defined as in (6) and (7). Then, the control problem can be formulated as follows

$$\min_{u_f \in \mathcal{U}} \quad \mathbb{E}_{\mathcal{N}\left(\hat{\mu}_{\text{pred}}, \hat{\Sigma}_{\text{pred}}\right)} [\tilde{c}_{\text{ctrl}}(u_f)]. \tag{8}$$

Next, we show that (8) has the same minimizer as the SPC method in (2) without output constraints, i.e., $\mathcal{Y} = \mathbb{R}^{d_y}$, where $d_y = (q - m)L_f$ is the dimension of y_f .

Theorem 1: If $u_{\rm f}^{\star}$ is a minimizer of (2) with $\mathcal{Y} = \mathbb{R}^{d_{\rm y}}$, then it is also a minimizer of (8).

Proof: The expectation of the control objective is

$$\mathbb{E}_{\mathcal{N}\left(\hat{\mu}_{\text{pred}},\hat{\Sigma}_{\text{pred}}\right)}[\tilde{c}_{\text{ctrl}}(u_{\text{f}})] \\ = \|u_{\text{f}} - u_{\text{ref}}\|_{R}^{2} + \mathbb{E}\left[\|(y_{\text{f}}|u_{\text{f}}, w_{\text{ini}}) - y_{\text{ref}}\|_{Q}^{2}\right] \\ = c_{\text{ctrl}}\left(\begin{bmatrix}u_{\text{f}}\\\hat{\mu}_{\text{pred}}\end{bmatrix}\right) + \operatorname{tr}\left(Q\hat{\Sigma}_{\text{pred}}\right).$$

Furthermore, due to the definition of $\hat{\mu}_{\text{pred}}$ in (6), $y_{\text{f}} = \hat{\mu}_{\text{pred}}$ which coincides with the explicit constraint in (2).

B. Distributionally optimistic control

We now show that, contrary to the certainty-equivalence formulation in (8), the regularized DeePC scheme is a distributionally optimistic formulation. Since we use finite data to identify $\hat{\mu}_{pred}$ and $\hat{\Sigma}_{pred}$, the estimates may be inaccurate. To account for this uncertainty, we propose to optimize over possible Gaussian distributions that are close to the estimated predictive model and lead to the lowest expected control cost. We quantify the closeness of two distributions by assuming an upper bound ϵ on their Kullback–Leibler (KL) divergence, leading to the distributionally optimistic problem

$$\min_{u_{f} \in \mathcal{U}} \min_{\mu \in \mathbb{R}^{d_{\mathcal{Y}}}, \Sigma \in \mathbb{S}_{>0}} \mathbb{E}_{\mathcal{N}(\mu, \Sigma)} [\tilde{c}_{\mathrm{ctrl}}(u_{f})] \\
\text{s.t. } D_{KL} \left(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N} \left(\hat{\mu}_{\mathrm{pred}}, \hat{\Sigma}_{\mathrm{pred}} \right) \right) \leq \epsilon.$$
(9)

The KL divergence D_{KL} measures the relative entropy between two distributions and for Gaussian distributions it takes the form [29]

$$D_{KL}\left(\mathcal{N}(\mu,\Sigma) \parallel \mathcal{N}(\hat{\mu}_{\text{pred}},\hat{\Sigma}_{\text{pred}})\right) = \frac{1}{2} \left(\operatorname{tr}(\hat{\Sigma}_{\text{pred}}^{-1}\Sigma) - d_{y} + \ln\left(\frac{\operatorname{det}(\hat{\Sigma}_{\text{pred}})}{\operatorname{det}(\Sigma)}\right) + \|\mu - \hat{\mu}_{\text{pred}}\|_{\hat{\Sigma}_{\text{pred}}^{-1}}^{2} \right).$$

Remark 3: An important benefit of the KL divergence is that the true covariance Σ only appears in terms that are independent of μ and $\hat{\mu}_{pred}$. The size of the uncertainty set ϵ is often a tuning parameter in robust formulations. Furthermore, we can rewrite the constraint in (9) as

$$\begin{aligned} \|\mu - \hat{\mu}_{\text{pred}}\|_{\hat{\Sigma}_{\text{pred}}^{-1}}^2 &\leq 2\epsilon \\ &- \left(\operatorname{tr}(\hat{\Sigma}_{\text{pred}}^{-1}\Sigma) - d_{\text{y}} + \ln\left(\operatorname{det}(\hat{\Sigma}_{\text{pred}})/\operatorname{det}(\Sigma) \right) \right), \end{aligned}$$

and consider the entire right-hand side of the inequality as a tuning parameter. Thus, the exact knowledge of Σ is not required. We use the KL divergence in (9) to exploit this benefit and also to highlight the similarities between robust formulations in the proposed framework and regularized data-driven control schemes. However, one can also consider other metrics, such as the Wasserstein [23] or Bhattacharyya distance [36].

To simplify (9), we lift the constraint on the KL divergence to the cost with a coefficient $\lambda \ge 0$ resulting in

$$\min_{\substack{u_{\mathrm{f}} \in \mathcal{U}, \ \mu \in \mathbb{R}^{d_{\mathrm{y}}}}} \mathbb{E}_{\mathcal{N}(\mu, \Sigma)} [\tilde{c}_{\mathrm{ctrl}}(u_{\mathrm{f}})] \\
+ \lambda \cdot D_{KL} \left(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N} \left(\hat{\mu}_{\mathrm{pred}}, \hat{\Sigma}_{\mathrm{pred}} \right) \right).$$
(10)

Note that due to the formula for the KL divergence and since $\mathbb{E}_{\mathcal{N}(\mu,\Sigma)}[\tilde{c}_{ctrl}(u_f)] = c_{ctrl}([u_f^{\top} \ \mu^{\top}]^{\top}) + tr(Q\Sigma)$, the minimizers u_f^{\star}, μ^{\star} of (10) are independent of the covariance Σ . For these reasons and those in Remark 3, we can drop the optimization over Σ . For a fixed (finite) value of Σ , the solution u_f^{\star}, μ^{\star} of the two problems coincide, if λ in (10) takes the value of the optimal Lagrange multiplier of the constraint in (9). In problem (10), λ is a tuning parameter that captures the uncertainty in the estimated predictive distribution. Intuitively, larger values of λ correspond to smaller uncertainty ϵ in (9), and vice versa. Note that as $\epsilon \rightarrow 0$, problem (9) collapses to the certainty-equivalence formulation (8), and hence, we expect the same to hold for (10) as $\lambda \rightarrow \infty$.

Next, we show that the projected 2-norm regularizer in the DeePC method and the softened constraint on the KL divergence in (10) play the same role, and therefore the two problems have the same minimizer. We consider the squared projected 2-norm regularizer $h(g) = ||(I - \Pi)g||_2^2$ also used, e.g., in [10], with Π defined as in Lemma 2.

Theorem 2: If $(w_{\rm f}^{\star}, g^{\star})$ is a minimizer of (3) with $h(g) = ||(I - \Pi)g||_2^2$ and $\mathcal{Y} = \mathbb{R}^{d_{\rm y}}$, then $u_{\rm f}^{\star} = U_{\rm f}g^{\star}$ and $\mu^{\star} = y_{\rm f}^{\star} = Y_{\rm f}g^{\star}$ is a minimizer of (10) with $\lambda = \lambda_g \frac{2}{D}$.

Note that solving (3) is computationally more expensive than solving (10), as it contains the optimization variable g, whose dimension scales with the number of data points D. Problem (10) accounts for the uncertainty in the estimated distribution in an optimistic fashion. The mean of the true distribution is chosen such that it minimizes the control cost. This optimism in the face of uncertainty has been observed perform well by balancing exploitation and exploration in adaptive settings [37], [38]. However, in the considered setup, one might expect the realized (closed-loop) control performance to degrade as the model uncertainty captured by ϵ increases, or equivalently, λ decreases. This intuition is supported by the empirical observation made in [10], stating that the realized control performance monotonically increases as λ_q in (3) is increased.

Remark 4: In [18], the authors argue that model predictive control (MPC) and direct data-driven control techniques such as regularized DeePC differ in the way the expectation of the control cost is calculated. In MPC, the future output is predicted by some "true" model \mathcal{M} , and the expected cost given the model $\mathbb{E}[\tilde{c}_{ctrl}(u_f)|\mathcal{M}]$ is minimized in the control problem. On the other hand, the model is not know in direct data-driven control methods. Instead, the expectation of the expected cost given the model is minimized given the data \mathcal{D} , i.e., the control input minimizes $\mathbb{E}[\mathbb{E}[\tilde{c}_{ctrl}(u_f)|\mathcal{M}]|\mathcal{D}]$.

In this work, we consider \mathcal{M} to be the predictive distribution $\mathcal{N}(\mu_{\mathrm{pred}}, \Sigma_{\mathrm{pred}})$ of a Gaussian dynamical system, and argue that SPC and regularized DeePC first estimate \mathcal{M} from data, and then perform control. SPC is an MPC formulation in this setup. On the other hand, the model \mathcal{M} is an optimization variable in regularized DeePC, and we pose the constraint that \mathcal{M} should remain close to its estimate.

Proof: To prove the claim, we show that the costs of (3) and (10) are equal up to a constant, and the feasible sets coincide. Any term that does not depend on $u_{\rm f}, y_{\rm f}, \mu$ or g will be denoted by c. First note that the expected control cost in (10) can be written as $\mathbb{E}_{\mathcal{N}(\mu,\Sigma)}[\tilde{c}_{\rm ctrl}(u_{\rm f})] = c_{\rm ctrl}([u_{\rm f}^{\top} \mu^{\top}]^{\top}) + c.$

Next, we show that h(g) and the KL divergence in (10) are related. Let us decompose g in (3) as

$$g = g_{\text{part}} + g_{\text{hom}}, \quad g_{\text{part}} \in \operatorname{im}(W^{\top}), \ g_{\text{hom}} \in \ker(W).$$

Since $(I - \Pi)^{\top}(I - \Pi) = I - \Pi, \ g_{\text{part}} \perp g_{\text{hom}}, \text{ and } (I - \Pi)^{\top}$

 Π) $g_{\text{hom}} = 0$, the regularizer becomes

$$h(g) = ||g_{\text{part}}||^2_{(I-\Pi)} + ||g_{\text{hom}}||^2_2.$$

Note that $g_{\text{hom}} \in \text{ker}(W)$ can be removed from the constraint and only appears in the regularization term h(g) in the optimization, therefore $g_{\text{hom}} = 0$ for any optimal solution. Consequently, $g = g_{\text{part}} = W^{\dagger} \begin{bmatrix} w_{\text{ini}} \\ w_{\text{f}} \end{bmatrix}$. Furthermore, since $\begin{bmatrix} W_{\text{p}} \\ U_{\text{f}} \end{bmatrix} g_{\text{part}} = \begin{bmatrix} w_{\text{ini}} \\ u_{\text{f}} \end{bmatrix}$, the regularization on g_{part} can be connected as

$$h(g_{\text{part}}) = \left\| \begin{bmatrix} w_{\text{ini}} \\ w_{\text{f}} \end{bmatrix} \right\|_{(WW^{\top})^{-1}}^{2} - \left\| \begin{bmatrix} w_{\text{ini}} \\ u_{\text{f}} \end{bmatrix} \right\|_{\left(\begin{bmatrix} W_{\text{p}} \\ U_{\text{f}} \end{bmatrix} \begin{bmatrix} W_{\text{p}} \\ U_{\text{f}} \end{bmatrix}^{\top} \right)^{-1}$$

Observe that for any x and $M \in \mathbb{S}_{>0}$, the weighted norm satisfies $||x||_M^2 = -2\log(p(x)) - \tilde{c}$, where p denotes the probability density function of $\mathcal{N}(0, M^{-1})$, and \tilde{c} is a constant independent of x. Then

$$\begin{split} h(g_{\text{part}}) &= -\frac{2}{D} \log \left(p_w([w_{\text{ini}}^\top \ w_{\text{f}}^\top]^\top) \right) \\ &+ \frac{2}{D} \log \left(p_{\text{m}}([w_{\text{ini}}^\top \ u_{\text{f}}^\top]^\top) \right) + c, \end{split}$$

where p_w is the density of $\mathcal{N}(0, \Sigma_w)$, and $p_{\rm m}$ is density of the estimated marginal distribution over $w_{\rm ini}$ and $u_{\rm f}$ given as $\mathcal{N}\left(0, \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix} \begin{bmatrix} W_{\rm p} \\ U_{\rm f} \end{bmatrix}^{\top}\right)$. We can then apply the law of conditional probability $p_{y_{\rm f}|w_{\rm ini},u_{\rm f}}(y_{\rm f}|w_{\rm ini},u_{\rm f}) = p_w(w)/p_{\rm m}([w_{\rm ini}^{\top}\ u_{\rm f}^{\top}]^{\top})$, where $p_{y_{\rm f}|w_{\rm ini},u_{\rm f}}$ is the density of the predictive distribution $\mathcal{N}(\hat{\mu}_{\rm pred}, \hat{\Sigma}_{\rm pred})$, leading to

$$h(g_{\text{part}}) = -\frac{2}{D} \log(p_{y_{\text{f}}|w_{\text{ini}},u_{\text{f}}}(y_{\text{f}})) + c$$

$$= \frac{1}{D} (y_{\text{f}} - \hat{\mu}_{\text{pred}})^{\top} \hat{\Sigma}_{\text{pred}}^{-1} (y_{\text{f}} - \hat{\mu}_{\text{pred}}) + c$$

$$= \frac{2}{D} D_{KL} \left(\mathcal{N}(y_{\text{f}}, \Sigma) \parallel \mathcal{N}(\hat{\mu}_{\text{pred}}, \hat{\Sigma}_{\text{pred}}) \right) + c.$$

Thus, for any optimal value of g, the regularizer h(g) is proportional to the KL divergence plus a constant term.

C. Distributionally robust control

j

While optimism in the face of uncertainty often works well, sometimes one needs to be conservative to minimize risks. For this reason, we propose a novel distributionally robust control formulation that minimizes the expected control cost under the worst-case deviation between the mean of the true and estimated predictive distributions. Let us again fix the covariance Σ , and formulate the distributionally robust control problem as

$$\min_{u_{f} \in \mathcal{U}} \max_{\mu \in \mathbb{R}^{d_{y}}} \mathbb{E}_{\mathcal{N}(\mu, \Sigma)}[\tilde{c}_{\text{ctrl}}(u_{f})] \\
\text{s.t. } D_{KL}\left(\mathcal{N}(\mu, \Sigma) \parallel \mathcal{N}\left(\hat{\mu}_{\text{pred}}, \hat{\Sigma}_{\text{pred}}\right)\right) \leq \epsilon.$$
(11)

By writing the dual of the inner problem, we obtain an upper bound on the cost for any feasible point, which is expressed as a minimization over the Lagrange multiplier λ of the constraint on the KL divergence.

Theorem 3: For any feasible pair $(\tilde{u}_f, \tilde{\mu})$, the cost of problem (11) is upper bounded by

$$\min_{\lambda \ge \lambda_0} \mathbb{E}_{\mathcal{N}(\mu^{\star}, \Sigma)} [\tilde{c}_{\text{ctrl}}(\tilde{u}_{\text{f}})] - \lambda \left(D_{KL} \left(\mathcal{N}(\mu^{\star}, \Sigma) \parallel \mathcal{N} \left(\hat{\mu}_{\text{pred}}, \hat{\Sigma}_{\text{pred}} \right) \right) - \epsilon \right).$$
(12)

where $\lambda_0 > 0$ is such that $\lambda_0 \hat{\Sigma}_{pred}^{-1} - Q \in \mathbb{S}_{>0}$, and μ^{\star} is defined as

$$\mu^{\star} := (\lambda \hat{\Sigma}_{\text{pred}}^{-1} - Q)^{-1} (\lambda \hat{\Sigma}_{\text{pred}}^{-1} \hat{\mu}_{\text{pred}} - Q y_{\text{ref}}).$$
(13)

The mean μ^* maximizes the Lagrangian of the inner problem in (11) by balancing the term $\|\mu - y_{\text{ref}}\|_Q^2$ from the control cost and the term $\lambda \|\mu - \hat{\mu}\|_{\hat{\Sigma}_{\text{pred}}^{-1}}^2$ from the KL divergence. Similarly to the distributionally optimistic formulation (10), we consider the value of $\lambda \ge \lambda_0$ in (12) to be a tuning parameter, instead of minimizing over it. Furthermore, we choose the input u_{f} that minimizes the upper bound (12) in Theorem 3 leading to the control problem

$$\min_{u_{\rm f}\in\mathcal{U}} \|\lambda\hat{\Sigma}_{\rm pred}^{-1}\hat{\mu}_{\rm pred} - Qy_{\rm ref}\|_{(\lambda\hat{\Sigma}_{\rm pred}^{-1} - Q)^{-1}}^{2} - \lambda\|\hat{\mu}_{\rm pred}\|_{\hat{\Sigma}_{\rm pred}^{-1}}^{2} + \|u_{\rm f} - u_{\rm ref}\|_{R}^{2} + c,$$
(14)

where c denotes the terms independent of $u_{\rm f}$. Since λ is the Lagrange multiplier of the constraint with the KL divergence in (11), increasing its value is directly related to considering a smaller uncertainty set, i.e., smaller ϵ . As $\lambda \to \infty$, we have $\mu^* \to \hat{\mu}_{\rm pred}$, and thus, we recover the certainty-equivalence formulation (8).

Notice that the estimated mean $\hat{\mu}_{\text{pred}}$ calculated as in (6) is affine in u_{f} , and thus, the cost of (14) is quadratic in u_{f} . Let us write

$$\hat{\mu}_{\text{pred}} = \hat{M}_u u_{\text{f}} + \hat{M}_{\text{ini}} w_{\text{ini}}.$$

Then, the Hessian of the cost in (14) with respect to $u_{\rm f}$ is

$$H(\lambda) := \lambda^2 \hat{M}_u^{\top} \hat{\Sigma}_{\text{pred}}^{-1} \left(\lambda \hat{\Sigma}_{\text{pred}}^{-1} - Q \right)^{-1} \hat{\Sigma}_{\text{pred}}^{-1} \hat{M}_u$$
$$- \lambda \hat{M}_u^{\top} \hat{\Sigma}_{\text{pred}}^{-1} \hat{M}_u + R.$$

If $\lambda \geq \lambda_0$ is chosen such that $H(\lambda) \succeq 0$, problem (14) is convex, and hence, can be solved efficiently. Furthermore, since $H(\lambda) \to R$ as $\lambda \to \infty$, the convexity of the problem is guaranteed for large enough λ . Preliminary numerical simulations showed that as λ is increased, the optimal solution of (14) converges to that of (8), as expected. Furthermore, the distributionally optimistic and robust formulations in (10) and (14) yield similar realized control costs for the LTI system considered in [10]. Thorough numerical testing of the algorithms including more challenging nonlinear case studies is the subject of future research.

Proof: For a fixed $\tilde{u}_{\rm f}$, the inner problem in (11) is

$$\begin{array}{ll} \min_{\mu \in \mathbb{R}^{d_{y}}} & -\mathbb{E}_{\mathcal{N}(\mu,\Sigma)}[\tilde{c}_{\mathrm{ctrl}}(\tilde{u}_{\mathrm{f}})] \\ \text{s.t.} & D_{KL}\left(\mathcal{N}(\mu,\Sigma) \parallel \mathcal{N}\left(\hat{\mu}_{\mathrm{pred}},\hat{\Sigma}_{\mathrm{pred}}\right)\right) \leq \epsilon. \end{array} \tag{15}$$

The Lagrangian associated with (15) is

$$\begin{aligned} \mathcal{L}(\mu,\lambda) &= -\mathbb{E}_{\mathcal{N}(\mu,\Sigma)}[\tilde{c}_{\mathrm{ctrl}}(\tilde{u}_{\mathrm{f}})] \\ &+ \lambda \left(D_{KL} \left(\mathcal{N}(\mu,\Sigma) \parallel \mathcal{N} \left(\hat{\mu}_{\mathrm{pred}}, \hat{\Sigma}_{\mathrm{pred}} \right) \right) - \epsilon \right) \\ &= \mu^{\top} (\lambda \hat{\Sigma}_{\mathrm{pred}}^{-1} - Q) \mu - 2(\lambda \hat{\mu}_{\mathrm{pred}}^{\top} \hat{\Sigma}_{\mathrm{pred}}^{-1} - y_{\mathrm{ref}}^{\top} Q) \mu + c(\lambda), \end{aligned}$$

where λ denotes the Lagrange multiplier of the inequality constraint, μ is the original optimization variable in (15), and $c(\lambda)$ denotes the terms that do not depend on μ . To ensure that the Lagrangian is strictly convex in μ , and thus, the dual function $g(\lambda) = \inf_{\mu} \mathcal{L}(\mu, \lambda)$ is bounded, we require $\lambda \ge \lambda_0$. Then, $g(\lambda) = \mathcal{L}(\mu^*, \lambda)$, with $\mu^* = \arg \min_{\mu} \mathcal{L}(\mu, \lambda)$ given in (13). For any pair $(\tilde{u}_f, \tilde{\mu})$, for which the original problem (11) is feasible and any $\lambda \ge \lambda_0 > 0$, we have that

$$g(\lambda) = \inf_{\mu} \mathcal{L}(\mu, \lambda) \le \mathcal{L}(\tilde{\mu}, \lambda) \le -\mathbb{E}_{\mathcal{N}(\tilde{\mu}, \Sigma)}[\tilde{c}_{\mathrm{ctrl}}(\tilde{u}_{\mathrm{f}})].$$

One can find the tightest lower bound by maximizing $g(\lambda)$ yielding the dual problem $\max_{\lambda \ge \lambda_0} g(\lambda)$. Finally, we arrive at (12) by writing the dual problem as a minimization of $-g(\lambda)$, leading to an upper bound on the cost of (11).

Remark 5: Note that the dual of the inner maximization problem in (11) can also be expressed as a semidefinite programming [39, App. B.1]. Then, strong duality holds if

$$\operatorname{tr}(\hat{\Sigma}_{\operatorname{pred}}^{-1}\Sigma) - d_{\operatorname{y}} + \ln\left(\frac{\operatorname{det}(\hat{\Sigma}_{\operatorname{pred}})}{\operatorname{det}(\Sigma)}\right) < 2\epsilon$$

that is, the constraint on the KL divergence satisfies Slater's condition. However, to ensure that the upper bound in Theorem 3 attains a simple form, we formulate the dual problem as in (12) and restrict λ such that the matrix $\lambda \hat{\Sigma}_{\text{pred}}^{-1} - Q$ is invertible.

V. CONCLUSION

We proposed a modeling framework for stochastic systems termed Gaussian behaviors, which is consistent with existing LTI system models. This system description leads to a new perspective on data-driven predictive control methods from the literature, by interpreting regularization as accounting for errors in the estimated distribution in an optimistic fashion. Furthermore, we proposed a novel distributionally robust control formulation that leads to a convex optimization problem. Future work includes extending the results beyond Gaussian distributions allowing us to model a larger class of systems and address nonlinearities.

REFERENCES

- J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," in *Proc. 18th European Control Conference*, pp. 307–312, IEEE, 2019.
- [2] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," *IEEE Trans. Automatic Control*, vol. 67, no. 7, pp. 3289–3304, 2021.
- [3] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE Trans. Automatic Control*, vol. 66, no. 4, pp. 1702–1717, 2020.
- [4] V. Breschi, A. Chiuso, and S. Formentin, "Data-driven predictive control in a stochastic setting: A unified framework," *Automatica*, vol. 152, p. 110961, 2023.

- [5] C. Verhoek, H. S. Abbas, R. Tóth, and S. Haesaert, "Data-driven predictive control for linear parameter-varying systems," in *Proc. 4th IFAC Workshop on Linear Parameter Varying Systems LPVS*, vol. 54, pp. 101–108, Elsevier, 2021.
- [6] H. J. van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, "Data informativity: A new perspective on data-driven analysis and control," *IEEE Trans. Automatic Control*, vol. 65, no. 11, pp. 4753– 4768, 2020.
- [7] C. De Persis and P. Tesi, "Formulas for data-driven control: Stabilization, optimality, and robustness," *IEEE Trans. Automatic Control*, vol. 65, no. 3, pp. 909–924, 2019.
- [8] J. C. Willems and J. W. Polderman, Introduction to mathematical systems theory: A behavioral approach. Springer Media, 1997.
- [9] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. De Moor, "A note on persistency of excitation," *Systems & Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.
- [10] F. Dörfler, J. Coulson, and I. Markovsky, "Bridging direct and indirect data-driven control formulations via regularizations and relaxations," *IEEE Trans. Automatic Control*, vol. 68, no. 2, pp. 883–897, 2022.
- [11] 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*, vol. 43, pp. 28–68, 2023.
- [12] I. Markovsky and F. Dörfler, "Behavioral systems theory in data-driven analysis, signal processing, and control," *Annual Reviews in Control*, vol. 52, pp. 42–64, 2021.
- [13] L. Huang, J. Zhen, J. Lygeros, and F. Dörfler, "Quadratic regularization of data-enabled predictive control: Theory and application to power converter experiments," in *Proc. 19th IFAC Symposium on System Identification*, pp. 192–197, Elsevier, 2021.
- [14] J. C. Willems, "Open stochastic systems," *IEEE Trans. Automatic Control*, vol. 58, no. 2, pp. 406–421, 2012.
- [15] G. Baggio and R. Sepulchre, "LTI stochastic processes: A behavioral perspective," in *Proc. 20th IFAC World Congress*, Elsevier, 2017.
- [16] T. Faulwasser, R. Ou, G. Pan, P. Schmitz, and K. Worthmann, "Behavioral theory for stochastic systems? A data-driven journey from Willems to Wiener and back again," *Annual Reviews in Control*, vol. 55, pp. 92–117, 2023.
- [17] G. Pan, R. Ou, and T. Faulwasser, "On a stochastic fundamental lemma and its use for data-driven optimal control," *IEEE Trans. Automatic Control*, vol. 68, no. 10, pp. 5922–5937, 2022.
- [18] A. Chiuso, M. Fabris, V. Breschi, and S. Formentin, "Harnessing uncertainty for a separation principle in direct data-driven predictive control," *Automatica*, vol. 173, p. 112070, 2025.
- [19] F. Zhao, F. Dörfler, A. Chiuso, and K. You, "Data-enabled policy optimization for direct adaptive learning of the LQR," arXiv preprint arXiv:2401.14871, 2024.
- [20] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part i," *IEEE Trans. Automatic Control*, vol. 61, no. 5, pp. 1158–1169, 2015.
- [21] F. Liu, G. Rapakoulias, and P. Tsiotras, "Optimal covariance steering for discrete-time linear stochastic systems," *IEEE Trans. Automatic Control*, vol. 70, no. 4, pp. 2289–2304, 2025.
- [22] J. Pilipovsky and P. Tsiotras, "Data-driven covariance steering control design," in *Proc. 62nd Conference on Decision and Control*, pp. 2610– 2615, IEEE, 2023.
- [23] L. Aolaritei, N. Lanzetti, H. Chen, and F. Dörfler, "Distributional uncertainty propagation via optimal transport," *arXiv preprint* arXiv:2205.00343, 2023.
- [24] J. L. Doob, Stochastic Processes. John Wiley, New York, 1953.
- [25] A. Papoulis, Probability, Random Variables, and Stochastic Processes. McGraw-Hill, New York, third edition ed., 1991.
- [26] H. J. Van Waarde, C. De Persis, M. K. Camlibel, and P. Tesi, "Willems' fundamental lemma for state-space systems and its extension to multiple datasets," *IEEE Contr. Syst. Lett.*, vol. 4, no. 3, pp. 602–607, 2020.
- [27] W. Favoreel, B. De Moor, and M. Gevers, "SPC: Subspace predictive control," *IFAC Proceedings Volumes*, vol. 32, no. 2, pp. 4004–4009, 1999.
- [28] G. Picci, An Introduction to Statistical Data Science. Springer, 2024.
- [29] C. K. Williams and C. E. Rasmussen, Gaussian processes for machine learning, vol. 2. MIT press Cambridge, MA, 2006.
- [30] K. V. Mardia, J. T. Kent, and C. C. Taylor, *Multivariate analysis*. Academic Press, 1995.

- [31] A. Padoan, F. Dörfler, and J. Lygeros, "Data-driven representations of conical, convex, and affine behaviors," in *Proc. 62nd Conference on Decision and Control*, pp. 596–601, IEEE, 2023.
- [32] R. J. Muirhead, Aspects of Multivariate Statistical Theory. John Wiley & Sons, Hoboken, NJ, 1982.
- [33] Z. Burda, A. Jarosz, M. A. Nowak, J. Jurkiewicz, G. Papp, and I. Zahed, "Applying free random variables to random matrix analysis of financial data. part i: The gaussian case," *Quantitative Finance*, vol. 11, no. 7, pp. 1103–1124, 2011.
- [34] W. Cui, X. Zhang, and Y. Liu, "Covariance matrix estimation from linearly-correlated Gaussian samples," *IEEE Trans. Signal Processing*, vol. 67, no. 8, pp. 2187–2195, 2019.
- [35] R. Vershynin, High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge University Press, 2018.
- [36] T. Kailath, "The divergence and bhattacharyya distance measures in signal selection," *IEEE Trans. Communication Technology*, vol. 15, no. 1, pp. 52–60, 2003.
- [37] Y. Abbasi-Yadkori and C. Szepesvári, "Regret bounds for the adaptive control of linear quadratic systems," in *Proc. 24th Annu. Conf. Learn. Theory*, pp. 1–26, 2011.
- [38] M. Abeille and A. Lazaric, "Efficient optimistic exploration in linearquadratic regulators via lagrangian relaxation," in *International Conference on Machine Learning*, pp. 23–31, 2020.
- [39] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.