

Data-Driven Tube-Based Stochastic Predictive Control

SEBASTIAN KERZ ^{id} (Graduate Student Member, IEEE), JOHANNES TEUTSCH ^{id},
TIM BRÜDIGAM ^{id} (Member, IEEE), MARION LEIBOLD ^{id}, AND
DIRK WOLLHERR ^{id} (Senior Member, IEEE)

(Intersection of Machine Learning With Control)

Technical University Munich, 80333 Munich, Germany

CORRESPONDING AUTHOR: SEBASTIAN KERZ (e-mail: s.kerz@tum.de)

ABSTRACT A powerful result from behavioral systems theory known as the fundamental lemma allows for predictive control akin to Model Predictive Control (MPC) for linear time-invariant (LTI) systems with unknown dynamics purely from data. While most data-driven predictive control literature focuses on robustness with respect to measurement noise, only a few works consider exploiting probabilistic information of disturbances for performance-oriented control as in stochastic MPC. This work proposes a novel data-driven stochastic predictive control scheme for chance-constrained LTI systems subject to measurement noise and additive stochastic disturbances. In order to render the otherwise stochastic and intractable optimal control problem deterministic, our approach leverages ideas from tube-based MPC by decomposing the state into a deterministic nominal state driven by inputs and a stochastic error state affected by disturbances. Satisfaction of original chance constraints is guaranteed by tightening nominal constraints probabilistically with respect to additive disturbances and robustly with respect to measurement noise. The resulting data-driven receding horizon optimal control problem is lightweight, recursively feasible, and renders the closed loop input-to-state stable in the presence of both additive disturbances and measurement noise. We demonstrate the effectiveness of the proposed approach in a simulation example.

INDEX TERMS Data-driven control, predictive control for linear systems, stochastic optimal control, uncertain systems.

I. INTRODUCTION

As Sensor data and computational power become more widely available, data-driven approaches are increasingly relevant in modern control applications. Recently, direct data-driven control design based on Willems' lemma [1] has emerged as an appealing alternative to the classical approach of first constructing a model from data based on system identification methods. Informally, the so-called *fundamental lemma* states that for discrete-time linear time-invariant (LTI) systems, the time-shifted vectors of any input-output trajectory generated by a persistently exciting input signal span the vector space of all (fixed length) input-output trajectories of the system. As a consequence, discrete-time LTI systems can be represented by a single measured trajectory, enabling control and analysis problems to be solved directly from trajectory

data [2], [3], [4], [5], [6]. A concise and comprehensive recent review is provided in [7]. By replacing the model with trajectory data, one effectively obtains a non-parametric representation of the subspace spanning the system behavior. This behavioral subspace can be directly searched by varying the coefficients of the linear combination of the basis (or *library of trajectories* [7]), making the framework naturally well-suited for finding optimal future input-output sequences within data-driven (or *data-enabled*) predictive control [8], [9]. Data-driven predictive controllers have since been modified for different classes of systems [10], [11], [12] with theoretical guarantees in various settings. Several works consider data perturbed by measurement noise and provide robust extensions of data-driven predictive control, akin to robust MPC (RMPC) [13], [14], [15], [16],

[17], [18]. Few works, however, consider a setting akin to stochastic MPC (SMPC), i.e., avoid overly conservative performance by exploiting probabilistic knowledge of stochastic disturbances.

A. RELATED WORK

Systems subject to stochastic disturbances are considered in [19], [20] for a very general setting in which no knowledge about the probability distribution of the disturbance is assumed, and the data available are inexact. Strong probabilistic guarantees on out-of-sample performance are given for the resulting open-loop data-driven optimal control problem. Probabilistic constraint satisfaction is enforced for the worst-case probability distribution that would explain the data. However, closed-loop properties and recursive feasibility of the optimal control problem are not considered. In [21], the authors present a data-driven predictive control scheme for systems affected by zero-mean white Gaussian process and measurement noise. By extending Willems' fundamental lemma to incorporate innovation data, the innovation form of the underlying state space system is represented from data. If the innovation data are exact, the resulting predictor is optimal, but ensuring stability and robustness of the closed-loop system remains an open research challenge.

In literature, only one line of work [22], [23], [24] falls into the category of data-driven SMPC and comes with certificates for recursive feasibility and stability. In [22], Pan et al. presented an extension of the fundamental lemma for stochastic LTI systems that leverages polynomial chaos expansion (PCE). Due to the linearity of PCE coefficients, stochastic variables described by those coefficients can be propagated through the dynamics in full. Applied to systems subject to stochastic additive disturbances with a known probability distribution, this allows for a deterministic reformulation of the otherwise intractable stochastic optimal control problem. The resulting predictive control scheme is rendered recursively feasible and practically stable with backup initial conditions [23] and is extended to the input-output case in [24]. Since distributional knowledge informs predictions in full, the approach is very appealing for settings in which the distribution of the disturbance is known exactly, but comes at the cost of relatively heavy online computations.

B. CONTRIBUTION

This work presents a novel data-driven SMPC scheme for the performance-oriented control of unknown LTI systems subject to stochastic disturbances based on trajectory data. By employing ideas from tube-based MPC, we provide a deterministic reformulation of the stochastic optimal control problem that is lightweight, recursively feasible, and leads to a stable closed loop. Similarly to [23], we assume that disturbance realizations can be measured or estimated offline before the control phase. In contrast to [23], we consider a setting in which no further distributional information is available, and an additional uncertainty, in the form of online state measurement noise, needs to be accounted for.

In order to combine the deterministic fundamental lemma with stochastic dynamics, we split the state into a deterministic nominal part affected by inputs and a stochastic error part affected by disturbances. By formulating the optimal control problem based on nominal states, it is rendered deterministic. Offline, before the actual control phase, probabilistic error predictions are used to appropriately tighten constraints on the nominal state such that original state constraints are met with a pre-specified probability level. We represent the decomposed dynamics in the data-driven setting and show how both the error predictions for the probabilistic constraint tightening and an additional robust constraint tightening with respect to bounded online measurement noise can be computed from data. To render the closed loop stable and the optimal control problem recursively feasible, we employ ideas from classical SMPC in the design of a data-driven robust first step constraint [25], stochastic tubes [26] with tightened constraint sets [27], terminal ingredients [28], [29], and a pre-stabilizing data-driven state feedback, with gains derived from the initially measured trajectory [3].

The main contribution is a recursively feasible and stable tube-based data-driven predictive control scheme that leverages chance constraints to efficiently control against process disturbances in the presence of measurement noise.

C. OUTLINE

The remainder of this work is structured as follows. Section II states the problem setting and all standing assumptions. Section III introduces preliminaries on tube-based MPC. Section IV proposes the data-driven predictive control algorithm, and Section V presents ingredients for closed loop certificates. Section VII tests the algorithm in simulation. Section VI discusses the case of noisy offline data, and we conclude the article in Section VIII.

D. NOTATION

Boldface uppercase (resp., lowercase) letters denote matrices (resp., vectors), \mathbf{I}_n is the $n \times n$ identity matrix, \mathbb{N}_a^b abbreviates the integer sequence $\{a, \dots, b\}$, and $(\cdot)^\dagger$ denotes a pseudoinverse. $\mathbf{A} > 0$ ($\mathbf{A} \geq 0$) means matrix \mathbf{A} is positive (semi-)definite. Sequences of vectors are written as $\mathbf{s}_{[1,N]} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$, and $\underline{\mathbf{s}}_{[1,N]} = [\mathbf{s}_1^\top, \dots, \mathbf{s}_N^\top]^\top$ denotes a stacked column vector of the sequence. Let $\|\mathbf{s}_{[1,N]}\|_\infty = \sup\{|\mathbf{s}_1|, \dots, |\mathbf{s}_N|\}$, and $\mathbf{s}_{[1,N]} \subset \mathcal{S}$ denote that $\mathbf{s}_i \in \mathcal{S}$ for all $i \in \mathbb{N}_1^N$. The weighted 2-norm $\sqrt{\mathbf{x}^\top \mathbf{Q} \mathbf{x}}$ of \mathbf{x} is abbreviated by $\|\mathbf{x}\|_{\mathbf{Q}}$. The probability of an event X is denoted by $\Pr(X)$. For a matrix \mathbf{H} and appropriate integers a, b , $[\mathbf{H}]_a$ is the a -th row vector of \mathbf{H} and $[\mathbf{H}]_{[a,b]}$ represents the submatrix of \mathbf{H} that is comprised of rows a, \dots, b . For any sets $\mathcal{A}, \mathcal{B} \subset \mathbb{R}^n$, we write the Minkowski set addition as $\mathcal{A} \oplus \mathcal{B} = \{a + b \mid a \in \mathcal{A}, b \in \mathcal{B}\}$, the Pontryagin set difference as $\mathcal{A} \ominus \mathcal{B} = \{a \in \mathcal{A} \mid (\forall b \in \mathcal{B}) a + b \in \mathcal{A}\}$ and set multiplication as $\mathbf{K}\mathcal{A} = \{\mathbf{K}a \mid a \in \mathcal{A}\}$. In the context of predictive control, we write $\mathbf{x}_{l|k}$ for the predicted state l steps ahead of \mathbf{x}_k . For any sequence of vectors $\mathbf{s}_{[0,N]}$, the Hankel matrix $\mathbf{H}_L(\mathbf{s}_{[0,N]})$ is

defined as

$$\mathbf{H}_L(\mathbf{s}_{[0, N]}) = \begin{bmatrix} s_0 & s_1 & \cdots & s_{N-L+1} \\ s_1 & s_2 & \cdots & s_{N-L+2} \\ \vdots & \vdots & \ddots & \vdots \\ s_{L-1} & s_L & \cdots & s_N \end{bmatrix}. \quad (1)$$

A continuous function $\gamma: \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is of class \mathcal{K} if γ is strictly increasing and $\gamma(0) = 0$. If $\gamma \in \mathcal{K}$ and γ is unbounded, then $\gamma \in \mathcal{K}_\infty$. If $\gamma: \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is continuous, strictly decreasing and $\lim_{t \rightarrow \infty} \gamma(t) = 0$, then it is of class \mathcal{L} . If $\delta \in \mathbb{R}_0^+ \times \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is such that $\delta(\cdot, t) \in \mathcal{K}$ and $\delta(r, \cdot) \in \mathcal{L}$, δ is of class \mathcal{KL} .

II. PROBLEM SETUP

We consider a linear, discrete-time time-invariant system

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

$$\hat{\mathbf{x}}_k = \mathbf{x}_k + \boldsymbol{\mu}_k \quad (2b)$$

with state $\mathbf{x}_k \in \mathbb{R}^{n_x}$, input $\mathbf{u}_k \in \mathbb{R}^{n_u}$ and $(\mathbf{A}, [\mathbf{B} \ \mathbf{E}])$ controllable. System (2) is subject to two kinds of uncertainty: A disturbance $\mathbf{d}_k \in \mathbb{R}^{n_d}$, akin to process noise, and a state measurement error $\boldsymbol{\mu}_k \in \mathbb{R}^{n_x}$, akin to measurement noise. Both disturbances \mathbf{d}_k and measurement errors $\boldsymbol{\mu}_k$ are realizations of a stochastic process whose distribution is fixed but unknown. The problem is to design a data-driven stochastic predictive control algorithm that stabilizes the origin while respecting input and state constraints. The predictive control algorithm centers around repeated solutions of a stochastic optimal control problem (OCP)

$$\underset{\mathbf{U}_k}{\text{minimize}} \mathbb{E} \left(\sum_{l=0}^{L-1} J_s(\mathbf{x}_{l|k}, \mathbf{u}_{l|k}) + J_f(\mathbf{x}_{L|k}) \right) \quad (3a)$$

$$\text{s.t. } \mathbf{x}_{0|k} = \mathbf{x}_k, \quad (3b)$$

$$\mathbf{x}_{l|k} \text{ evolves according to (2a),} \quad (3c)$$

$$\mathbf{u}_{l|k} \in \mathcal{U} \quad \forall l \in \mathbb{N}_0^{L-1}, \quad (3d)$$

$$\Pr(\mathbf{x}_{l|k} \in \mathcal{X}) \geq p \quad \forall l \in \mathbb{N}_1^L, \quad (3e)$$

where $\mathbb{E}(\cdot)$ denotes the expected value, and J_s and J_f represent stage and terminal costs, respectively. State constraints must be satisfied in probability, as specified by the chance constraint (3e) with user-specified risk parameter $p \in (0, 1]$. Both \mathcal{U} and \mathcal{X} are user-specified convex polytopic sets

$$\mathcal{U} = \{ \mathbf{u} \in \mathbb{R}^{n_u} \mid \mathbf{G}_u \mathbf{u} \leq \mathbf{g}_u \}, \quad (4a)$$

$$\mathcal{X} = \{ \mathbf{x} \in \mathbb{R}^{n_x} \mid \mathbf{G}_x \mathbf{x} \leq \mathbf{g}_x \}. \quad (4b)$$

Problem (3) is solved in a receding horizon fashion, where in each time step k , the first input $\mathbf{u}_{0|k}$ of the minimizing solution $\mathbf{U}_k = (\mathbf{u}_{0|k}, \dots, \mathbf{u}_{L-1|k})$ is applied to the actual system (2). Due to the probabilistic nature of both the system evolution (3c) and the chance constraints (3e), problem (3) is

intractable. This work aims to provide a tractable reformulation based on data.

A. AVAILABLE DATA

The system matrices \mathbf{A} , \mathbf{B} , and \mathbf{E} of system (2a) are unknown. In their stead, we use persistently exciting trajectory data to design the predictive control scheme.

Definition 1 (Persistency of excitation): A sequence of vectors $\mathbf{s}_{[0, N-1]} \in \mathbb{R}^m$ is persistently exciting of order L if the Hankel matrix $\mathbf{H}_L(\mathbf{s}_{[0, N-1]})$ has full row rank mL .

Assumption 1: A trajectory $(\mathbf{u}^d_{[0, N-1]}, \mathbf{d}^d_{[0, N-1]}, \mathbf{x}^d_{[0, N-1]})$ generated by (2a) is available, where the sequence of joint vectors $\mathbf{s}_{[0, N-1]}$, $\mathbf{s}_i = \begin{bmatrix} \mathbf{u}_i^d \\ \mathbf{d}_i^d \end{bmatrix}$, is persistently exciting of order $L + n_x + 1$.

Remark 1: For ease of exposition, and as in related works [21], [23], we consider full state measurements, instead of output measurements. An extension to the input-output setting can follow from considering an extended state of past outputs, as in for example [16], [24]. If disturbance recordings cannot be accessed, they may be estimated from input-state data as in [22, Section IV.D]. Note that if $\mathbf{E} \neq \mathbf{I}_{n_x}$, further structure on \mathbf{E} needs to be assumed. The availability of noise-free offline data is restrictive, especially in a setting with online noise. Nevertheless, we include online noise in order to narrow the gap to a fully realistic setting and show the flexibility of the proposed approach in dealing with additional uncertainties. We give a more detailed discussion on the influence and challenges of noisy offline data in Section VI, and show its effect in a simulation example in Section VII.

In order to guarantee the satisfaction of state constraints (3e), we employ a stochastic constraint tightening with respect to the disturbance \mathbf{d} , and a robust constraint tightening with respect to the measurement noise $\boldsymbol{\mu}$. To this end, we make the following assumption.

Assumption 2:

(a) In addition to the data in Assumption 1, N_S samples of disturbance sequences $\mathbf{d}_{[0, L-1]}^{(i)}$, $i \in \mathbb{N}_1^{N_S}$, that is $N_S L$ realizations of the disturbance process, are available.

(b) The noise $\boldsymbol{\mu}$ is bounded by a known polytopic set

$$\mathcal{M} = \{ \boldsymbol{\mu} \in \mathbb{R}^{n_x} \mid \mathbf{G}_\mu \boldsymbol{\mu} \leq \mathbf{g}_\mu \}. \quad (5)$$

that contains the origin.

(c) The disturbance \mathbf{d} is bounded by a known polytopic set

$$\mathcal{D} = \{ \mathbf{d} \in \mathbb{R}^{n_d} \mid \mathbf{G}_d \mathbf{d} \leq \mathbf{g}_d \} \quad (6)$$

that contains the origin.

Items (a) and (b) of Assumption 2 are necessary by virtue of stochastic and robust constraint handling, respectively. The addition of item (c) allows for recursive feasibility and stability in a robust sense, i.e., both properties can only be guaranteed in probability if either \mathbf{d} or $\boldsymbol{\mu}$ are unbounded.

B. DESIRED PROPERTIES

We want to guarantee that the reformulated OCP (3) remains solvable while the system is steered by the predictive controller. Given that the OCP is feasible at initial time, *recursive feasibility* guarantees that constraint satisfaction is possible for all time.

Definition 2 (Recursive Feasibility): A receding horizon OCP is *recursively feasible* if the existence of an admissible solution \mathbf{U}_k implies the existence of an admissible solution \mathbf{U}_{k+1} at the next time step.

Additionally, we desire input-to-state stability of the closed-loop system

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{w}_k), \quad (7)$$

with respect to $\mathbf{w}_k \in \mathbb{R}^{n_w}$, which is an extended disturbance term comprising all uncertainty (both from the measurement noise and the additive disturbance). For the following definitions, let the extended disturbance be bounded by a compact support set \mathcal{W} that contains the origin. Furthermore, let the origin be an equilibrium of the closed loop system for zero disturbance, $\mathbf{f}(\mathbf{0}, \mathbf{0}) = \mathbf{0}$, as is the case for system (2a). Denote by $\phi(k, \mathbf{x}_0, \mathbf{w}_{[0, k]})$ the solution to (7) at time step k for initial state \mathbf{x}_0 and disturbance sequence $\mathbf{w}_{[0, k]}$.

Definition 3 (Robust positive invariant set): A set \mathcal{X}_0 is *robust positive invariant* (RPI) for system (7) if $\mathbf{x}_0 \in \mathcal{X}_0 \Rightarrow \phi(k, \mathbf{x}_0, \mathbf{w}_{[0, k]}) \in \mathcal{X}_0$ for all $k \in \mathbb{N}$, and for all disturbance realizations $\mathbf{w}_{[0, k]} \subset \mathcal{W}$.

Definition 4 (Input-to-state stability [28, Def. 19]): Let $\mathcal{X}_0 \subseteq \mathbb{R}^{n_x}$ be a closed robust positive invariant set for system (7) with the origin in its interior. System (7) is *input-to-state stable* (ISS) with respect to the disturbance and with region of attraction \mathcal{X}_0 , if there exist functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that for all $k \in \mathbb{N}_0$, $\mathbf{x}_0 \in \mathcal{X}_0$, and $\mathbf{w}_k \in \mathcal{W}$

$$\|\phi(k, \mathbf{x}_0, \mathbf{w}_{[0, k]})\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\|\mathbf{w}_{[0, k]}\|_\infty). \quad (8)$$

By causality, the same definition would result if $\mathbf{w}_{[0, k]}$ is replaced by $\mathbf{w}_{[0, k-1]}$ [30]. Definition 4 is equivalent to *robust asymptotic stability* of the origin as defined in [29, Definition 4.43]. It implies that the origin of the undisturbed system $\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{0})$ is asymptotically stable with region of attraction \mathcal{X}_0 .

III. TUBE-BASED MPC

In this work, we present a tractable deterministic reformulation of the stochastic OCP (3) based on trajectory data. This reformulation uses ideas from stochastic and robust tube-based MPC [26], [29], which are introduced in the following. In tube-based MPC, a key idea is to decompose the state \mathbf{x} into a deterministic nominal part \mathbf{z} and a (possibly) stochastic error part \mathbf{e} ,

$$\mathbf{x}_k = \mathbf{z}_k + \mathbf{e}_k. \quad (9)$$

Given the underlying system (2a), the resulting dynamics are

$$\mathbf{z}_{k+1} = \mathbf{A}\mathbf{z}_k + \mathbf{B}\mathbf{u}_k, \quad \mathbf{x}_0 = \mathbf{z}_0, \quad (10a)$$

$$\mathbf{e}_{k+1} = \mathbf{A}\mathbf{e}_k + \mathbf{E}\mathbf{d}_k, \quad \mathbf{e}_0 = \mathbf{0}. \quad (10b)$$

By only considering the deterministic nominal dynamics (10a) in the predictive controller, the OCP (3) is rendered deterministic. The expected value in the cost function (3a) can then be computed explicitly, and all terms depending on the state error can be neglected since they are constant.

For any finite time step k , the error \mathbf{e} is bounded, and the actual state \mathbf{x} is in a neighborhood of the nominal state \mathbf{z} . For an L -step trajectory, these neighborhoods form a tube around the nominal state predictions. The original state \mathbf{x} satisfies user-specified constraints \mathcal{X} if the constraints on the nominal state \mathbf{z} are tightened with respect to that tube. Let us denote the tightened constraint sets for the nominal states by \mathcal{Z}_l , i.e., $\mathbf{z}_{l|k} \in \mathcal{Z}_l$ replaces (3e) and $\{\mathcal{Z}_0, \dots, \mathcal{Z}_{L-1}\}$ describes the nominal state tube.

A robust constraint tightening (as in RMPC) computes \mathcal{Z}_l such that for all possible realizations of uncertainty

$$\mathbf{z}_{l|k} \in \mathcal{Z}_l \Rightarrow \mathbf{x}_{l|k} \in \mathcal{X}. \quad (11)$$

Naturally, \mathcal{Z}_l should be as large as possible, as every other choice introduces unnecessary conservatism. The maximal sets that satisfy (11) are $\mathcal{Z}_l = \mathcal{X} \ominus \mathcal{E}_l$, where $\mathcal{E}_l = \{\mathbf{e} \in \mathbb{R}^{n_x} \mid (\exists \mathbf{d}_0, \dots, \mathbf{d}_{l-1} \in \mathcal{D}) \mathbf{e} = \mathbf{e}_l \text{ from (10b)}\}$ encompasses the resulting state errors $\mathbf{e}_{l|k}$ for all realizations of the disturbance.

In a probabilistic constraint tightening (as in SMPC), the state constraints (4b) are required to hold with probability level p for each future predicted state $\mathbf{x}_{l|k}$, that is

$$\mathbf{z}_{l|k} \in \mathcal{Z}_l \Rightarrow \Pr(\mathbf{x}_{l|k} \in \mathcal{X}) > p, \quad (12)$$

where the conditional dependency on $\mathbf{x}_{0|k} = \mathbf{x}_k$ is understood and omitted in the following. With \mathcal{X} defined as in (4b) and the state decomposed as in (9), the right-hand-side of (12) can be split into two separate expressions

$$\mathbf{G}_x \mathbf{z}_{l|k} \leq \tilde{\boldsymbol{\eta}}, \quad (13a)$$

$$\Pr(\tilde{\boldsymbol{\eta}} \leq \mathbf{g}_x - \mathbf{G}_x \mathbf{e}_{l|k}) \geq p, \quad (13b)$$

with the introduction of a new parameter $\tilde{\boldsymbol{\eta}} \in \mathbb{R}^{r_x}$. The deterministic constraint (13a) is used online, i.e., replaces the chance constraints (3e) in the optimal control problem. Offline, before the actual control phase, the parameter $\tilde{\boldsymbol{\eta}}$ is computed such that (13b) holds. Tightened nominal constraints are defined as

$$\mathcal{Z}_l = \{\mathbf{z} \in \mathbb{R}^{n_x} \mid \mathbf{G}_x \mathbf{z} \leq \boldsymbol{\eta}_l\}, \quad (14)$$

and to minimize conservatism, $\boldsymbol{\eta}_l$ is determined by solving

$$\boldsymbol{\eta}_l = \max_{\tilde{\boldsymbol{\eta}}} \tilde{\boldsymbol{\eta}} \quad (15a)$$

$$\text{s.t. } \Pr(\tilde{\boldsymbol{\eta}} \leq \mathbf{g}_x - \mathbf{G}_x \mathbf{e}_l) \geq p, \quad (15b)$$

with $\max(\cdot)$ applied element-wise. Suppose the true disturbance distribution is unknown, and a finite number of N_S disturbance sequences $\mathbf{d}_{[0, l-1]}$ are available. Then, the chance-constrained optimization problem (15) may be solved approximately by reformulating it based on computed state error samples \mathbf{e}_l . To that end, a multitude of methods are available, see the recent survey [31]. For simplicity, we employ

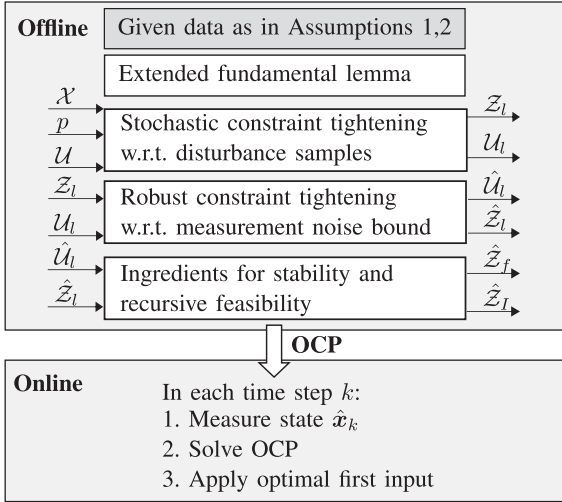


FIGURE 1. Overview of the proposed approach. The extended fundamental lemma, data-driven stochastic and robust constraint tightening, and the overall algorithm are presented in Section IV. Additional constraints that ensure stability and recursive stability are presented in Section V.

a scenario approximation [32], in which (15) is reformulated into a large-scale linear program where the chance constraint (15b) is replaced by deterministic constraints

$$\tilde{\eta} \leq \mathbf{g}_x - \mathbf{G}_x \mathbf{e}_l \quad (16)$$

and required to hold for all but N_D , $N_D < N_S$, samples, which are discarded. For the resulting η_l , (15b) holds with confidence $1 - \beta$ for a risk parameter $p \in [p_{\min}, p_{\max}]$. A more detailed description of the sampling-based solution to (15) is provided in Section A of the supplementary material.

Remark 2: If the probability distribution of the error is known, the chance-constrained optimization problem (15) may be solved numerically to arbitrary precision. If the probability distribution of the error is Gaussian, the chance constraint may be reformulated into an analytic expression [33]. If the probability distribution of the disturbance is known, the probability distribution of the error may be obtained by propagating the disturbance distribution through the error dynamics (10b) via the stochastic fundamental lemma [22].

In Section IV-IV-B, we show how state error samples $\mathbf{e}_l^{(i)}$ are computed from disturbance samples $\mathbf{d}_{[0, l-1]}^{(i)}$ in the data-driven setting, and present a data-driven formulation of robust tubes. An overview of the proposed approach and the contents of the remainder of the article is given in Fig. 1.

IV. DATA-DRIVEN STOCHASTIC PREDICTIVE CONTROL

Trajectory data generated by a persistently exciting input signal allow for a data-driven representation of discrete-time LTI systems, based on a powerful result of behavioral system theory [1], which we state in input-state-space [34].

Lemma 1 (Fundamental Lemma): Consider system (2a) with $\mathbf{d}_k = \mathbf{0}$. If the input sequence $\mathbf{u}_{[0, N-1]}^d$ is persistently exciting of order $L + n_x + 1$, then any $(L + 1)$ -long input-state sequence $(\mathbf{u}_{[k, k+L]}, \mathbf{x}_{[k, k+L]})$ is a valid trajectory

of the system if and only if there exists $\alpha \in \mathbb{R}^{N-L}$ such that

$$\begin{bmatrix} \mathbf{u}_{[k, k+L]} \\ \mathbf{x}_{[k, k+L]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{L+1}(\mathbf{u}_{[0, N-1]}^d) \\ \mathbf{H}_{L+1}(\mathbf{x}_{[0, N-1]}^d) \end{bmatrix} \alpha. \quad (17)$$

Lemma 1 describes a non-parametric system representation of a discrete-time LTI system, where varying α on the right-hand-side of (17) returns different $(L + 1)$ -step input-state trajectories on the left-hand-side. In the OCP of a data-driven predictive control scheme [8], [9], [13], equation (17) (or the equivalent input-output formulation) replaces the model and α acts as the new decision variable.

A. DATA-DRIVEN REPRESENTATION OF STABILIZED NOMINAL AND ERROR DYNAMICS

In order to represent system (2a) with past data, we extend Lemma 1 to the case of systems with additive disturbance. In the following, data Hankel matrices will be abbreviated as $\mathbf{H}_u := \mathbf{H}_{L+1}(\mathbf{u}_{[0, N-1]}^d)$, $\mathbf{H}_d := \mathbf{H}_{L+1}(\mathbf{d}_{[0, N-1]}^d)$, and $\mathbf{H}_x := \mathbf{H}_{L+1}(\mathbf{x}_{[0, N-1]}^d)$.

Lemma 2 (Extended Fundamental Lemma): Consider data as in Assumption 1. Any $(L + 1)$ -long input-disturbance-state sequence $(\mathbf{u}_{[k, k+L]}, \mathbf{d}_{[k, k+L]}, \mathbf{x}_{[k, k+L]})$ is a valid trajectory of system (2a) if and only if there exists $\alpha \in \mathbb{R}^{N-L}$ such that

$$\begin{bmatrix} \mathbf{u}_{[k, k+L]} \\ \mathbf{d}_{[k, k+L]} \\ \mathbf{x}_{[k, k+L]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha. \quad (18)$$

Proof: Lemma 2 follows from Lemma 1 by considering an extended input $\begin{bmatrix} \mathbf{u}^d \\ \mathbf{d}^d \end{bmatrix}$ and reordering rows accordingly. ■

Remark 3: Note that trajectory data of length $N \geq (n_u + n_d + 1)(L + 1) + n_x$ is required, such that the system of equations (18) is not overdetermined for given initial state \mathbf{x}_0 , input sequence $\mathbf{u}_{[k, k+L]} \in \mathbb{R}^{n_u(L+1)}$, and disturbance sequence $\mathbf{d}_{[k, k+L]} \in \mathbb{R}^{n_d(L+1)}$.

In order to counteract an inflation of the error state due to the dynamics in \mathbf{A} , we pre-stabilize the system by introducing a stabilizing state feedback controller. Accordingly, the input \mathbf{u}_k is decomposed into a state feedback component and a new artificial input \mathbf{v}_k , $\mathbf{u}_k = \mathbf{K}\mathbf{x}_k + \mathbf{v}_k$.

Remark 4: Stabilizing and LQR-optimal state feedback gains $\mathbf{K} \in \mathbb{R}^{n_u \times n_x}$ for nominal LTI systems can be computed from data based on [3, Theorem 3] and [3, Theorem 4], respectively (cf. Section B of the supplementary material). To that end, we retrieve a nominal (undisturbed) input-state trajectory with Lemma 2 by setting $\mathbf{d}_{[k, k+L]} = \mathbf{0}$, choosing arbitrary inputs $\mathbf{u}_{[k, k+L]}$, fixing an arbitrary initial state \mathbf{x}_k , and solving (18) for the resulting state sequence $\mathbf{x}_{[k+1, k+L]}$.

In order to include the state feedback in the data-driven system representation (18), i.e., represent the closed-loop behavior with open-loop data, we pretend the data generating input sequence was already given by $\mathbf{u}_k = \mathbf{K}\mathbf{x}_k + \mathbf{v}_k$. Then,

we rearrange (18) appropriately, to change the input variable of interest from \mathbf{u}_k to $\mathbf{v}_k = \mathbf{u}_k - \mathbf{K}\mathbf{x}_k$.

Lemma 3 (Pre-stabilized fundamental lemma): Consider open-loop data of system (2a) as in Assumption 1 and let $\mathbf{K} \in \mathbb{R}^{n_u \times n_x}$ be chosen such that $\mathbf{u}^d_{[0, N-1]} - \mathbf{K}\mathbf{x}^d_{[0, N-1]}$ is persistently exciting of order $L + n_x + 1$. Let $\tilde{\mathbf{K}} \in \mathbb{R}^{n_u(L+1) \times n_x(L+1)}$ be a block-diagonal expansion of \mathbf{K} . Any $(L + 1)$ -long sequence $(\mathbf{v}_{[k, k+L]}, \mathbf{d}_{[k, k+L]}, \mathbf{x}_{[k, k+L]})$ is a valid trajectory of $\mathbf{x}_{k+1} = (\mathbf{A} + \mathbf{BK})\mathbf{x}_k + \mathbf{B}\mathbf{v}_k + \mathbf{E}\mathbf{d}_k$ if and only if there exists $\alpha \in \mathbb{R}^{N-L}$ such that

$$\begin{bmatrix} \mathbf{v}_{[k, k+L]} \\ \mathbf{d}_{[k, k+L]} \\ \mathbf{x}_{[k, k+L]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha. \quad (19)$$

Proof: The proof follows from Lemma 18, by replacing \mathbf{H}_u with $\mathbf{H}_v = \mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x$, reflecting the change of input from \mathbf{u}_k to \mathbf{v}_k . ■

Since the dynamics of the underlying system are stochastic in nature, we split the system state into nominal state \mathbf{z} and error state \mathbf{e} , as motivated in Section III. In the data-driven setting, the state decomposition (9) translates to nominal state trajectories $\mathbf{z}_{[k, k+L]}$ only influenced by inputs $\mathbf{u}_{[0, L]}$, and state error trajectories $\mathbf{e}_{[k, k+L]}$, only influenced by disturbances $\mathbf{d}_{[0, L]}$. With either the disturbance sequence or the input sequence set to zero in (19), we obtain the data-driven representation of the pre-stabilized nominal (10a) or error (10b) dynamics as

$$\begin{bmatrix} \mathbf{v}_{[k, k+L]} \\ \mathbf{0} \\ \mathbf{z}_{[k, k+L]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha_z, \quad (20a)$$

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{d}_{[k, k+L]} \\ \mathbf{e}_{[k, k+L]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha_e, \quad (20b)$$

where $\alpha_e \in \mathbb{R}^{N-L}$, $\alpha_z \in \mathbb{R}^{N-L}$, and the true state sequence is given by $\mathbf{x}_{[k, k+L]} = \mathbf{z}_{[k, k+L]} + \mathbf{e}_{[k, k+L]} = \mathbf{H}_x(\alpha_z + \alpha_e)$. Equation (20b) allows for the computation of state errors $\mathbf{e}_{[k, k+L]}$ from a sequence of disturbances $\mathbf{d}_{[k, k+L]}$, which will be used to compute tightened constraints for the nominal state \mathbf{z} in Section IV-IV-B. In practice, this computation takes place offline, before the actual control phase. Equation (20a) allows for prediction of nominal state sequences and replaces (3c) in the OCP (formulated in Section IV-IV-D), to find optimal input sequences $\mathbf{v}_{[k, k+L]}$ during the control phase.

B. DATA-DRIVEN TIGHTENING OF STATE CONSTRAINTS

In order to guarantee chance constraint satisfaction of the states $\mathbf{x}_{l|k}$ in the prediction horizon, we subject the nominal states $\mathbf{z}_{l|k}$ to tightened constraints. We first tighten constraints probabilistically with respect to the additive disturbance, before further tightening constraints robustly with respect to the measurement noise.

The probabilistically tightened constraint sets \mathcal{Z}_l are defined by (14) based on parameters η_l . To find η_l , the chance constrained optimization problem (15) is solved using sampled state error sequences $\mathbf{e}_{[0, L]}$. Given the initial state error $\mathbf{e}_0 = \mathbf{0}$ and a sampled disturbance sequence $\mathbf{d}_{[0, L]}$, the state error sequence $\mathbf{e}_{[0, L]}$ follows from (20b) as

$$\mathbf{e}_{[0, L]} = \mathbf{H}_x \begin{bmatrix} \mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x \\ \mathbf{H}_d \\ [\mathbf{H}_x]_{[1, \mathbb{R}^{n_x}]} \end{bmatrix}^\dagger \begin{bmatrix} \mathbf{0} \\ \mathbf{d}_{[0, L]} \\ \mathbf{0} \end{bmatrix}. \quad (21)$$

With the constraints \mathcal{X} tightened to \mathcal{Z}_l for the nominal state, (12) holds (with confidence β) and the pre-specified chance constraints are satisfied for system (2a).

Next, we account for inexact state measurements $\hat{\mathbf{x}}_k = \mathbf{x}_k + \boldsymbol{\mu}_k$ (2b) during the control phase. Let $\hat{\mathbf{z}}_{[0|k, L|k]}$ denote the nominal state predictions initialized at the measured state $\hat{\mathbf{z}}_{[0|k]} = \hat{\mathbf{x}}_k$. In the following, we further tighten the tube sets \mathcal{Z}_l robustly with respect to all possible realizations of the measurement noise $\boldsymbol{\mu}_k \in \mathcal{M}$. That is, we construct robust tube sets $\hat{\mathcal{Z}}_l$ as in (11) such that

$$\hat{\mathbf{z}}_{l|k} \in \hat{\mathcal{Z}}_l \Rightarrow \mathbf{z}_{l|k} \in \mathcal{Z}_l. \quad (22)$$

In order to obtain a maximal tube, we set $\hat{\mathcal{Z}}_l := \mathcal{Z}_l \ominus \mathcal{E}_{\mu, l}$, where $\mathcal{E}_{\mu, l}$ encompasses all possible deviations from predicted nominal states $\hat{\mathbf{z}}_{l|k}$ to noise-free nominal states $\mathbf{z}_{l|k}$,

$$\mathbf{z}_{l|k} \in \{\hat{\mathbf{z}}_{l|k}\} \oplus -\mathcal{E}_{\mu, l} \quad \forall \boldsymbol{\mu}_{[0|k, l|k]} \in \mathcal{M}. \quad (23)$$

The smallest possible sets $\mathcal{E}_{\mu, l}$ which satisfy (23) are given by $\mathcal{E}_{\mu, l} = \{\hat{\mathbf{z}}_{l|k} - \mathbf{z}_{l|k} \mid \boldsymbol{\mu}_k \in \mathcal{M}\}$, and are bounded, since the measurement noise $\boldsymbol{\mu}_k \in \mathcal{M}$ is bounded. As the deviation of the two nominal trajectories evolves with \mathbf{A} via

$$\hat{\mathbf{z}}_{l|k} - \mathbf{z}_{l|k} = \mathbf{A}^l \boldsymbol{\mu}_k, \quad \forall l \in \mathbb{N}_{0}^L, \quad (24)$$

the difference of the two input-disturbance-state trajectories $(\mathbf{u}_{[k, k+L]}, \mathbf{0}, \hat{\mathbf{z}}_{[0|k, L|k]})$ and $(\mathbf{u}_{[k, k+L]}, \mathbf{0}, \mathbf{z}_{[0|k, L|k]})$ is again a valid input-disturbance-state trajectory of the same underlying system. Consequently, the data-driven representation in Lemma 2 applies, yielding

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \hat{\mathbf{z}}_{[0|k, L|k]} - \mathbf{z}_{[0|k, L|k]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha_\mu, \quad (25)$$

with $\alpha_\mu \in \mathbb{R}^{N-L}$, and where the noise $\boldsymbol{\mu}_k = \hat{\mathbf{z}}_{0|k} - \mathbf{z}_{0|k}$ appears as the first n_x entries of $\hat{\mathbf{z}}_{[0|k, L|k]} - \mathbf{z}_{[0|k, L|k]}$, i.e.,

$$\boldsymbol{\mu}_k = [\mathbf{H}_x]_{[1, n_x]} \alpha_\mu. \quad (26)$$

Based on (25), (26) and with $\boldsymbol{\mu}_k \in \mathcal{M}$ as in Assumption 2, the sets $\mathcal{E}_{\mu, l}$ may be constructed from data.

Lemma 4: Consider system (2) and data persistently exciting of order $L + n_x + 1$ as in Assumption 1. Let

$$\mathbb{A}_\mu = \left\{ \alpha \in \mathbb{R}^{N-L} \mid \begin{bmatrix} \mathbf{H}_u \\ \mathbf{H}_d \end{bmatrix} \alpha = \mathbf{0}, \mathbf{G}_\mu [\mathbf{H}_x]_{[1, n_x]} \alpha \leq \mathbf{g}_\mu \right\}$$

and define for each $l \in \mathbb{N}_0^L$ the tube sets

$$\mathcal{E}_{\mu,l} = [\mathbf{H}_x]_{[l n_x + 1, l n_x + n_x]} \mathbb{A} \mu. \quad (27)$$

Then, $\mathbf{z}_{l|k} \in \{\hat{\mathbf{z}}_{l|k}\} \oplus -\mathcal{E}_{\mu,l}$ for all μ_k and the sets $\mathcal{E}_{\mu,l}$ define the smallest possible tube that satisfies (22).

Proof: We have to prove that $\hat{\mathbf{z}}_{l|k} - \mathbf{z}_{l|k} \in \mathcal{E}_{\mu,l}$ for all $l \in \mathbb{N}_1^L$ and all $\mu_k \in \mathcal{M}$. Denote with $\mathbf{H}_{x,l} = [\mathbf{H}_x]_{[l n_x + 1, l n_x + n_x]}$ the rows of \mathbf{H}_x in (25) corresponding to $\hat{\mathbf{z}}_{l|k} - \mathbf{z}_{l|k}$. First, assume μ_k fixed and note that by Lemma 2 and (25), $\hat{\mathbf{z}}_{l|k} - \mathbf{z}_{l|k} = \mathbf{H}_{x,l} \alpha_\mu = \mathbf{A}^l \mu_k$ if $\mathbf{H}_{x,0} \alpha = \mu_k$ and $\mathbf{H}_d |_{\alpha=0}$. Thus, as $\mu_k \in \mathcal{M}$ per Assumption 2, $\mathbf{A}^l \mathcal{M} = \mathbf{H}_{x,l} \{\alpha \in \mathbb{R}^{N-L} \mid \mathbf{H}_d |_{\alpha=0}, \mathbf{H}_{x,0} \alpha \in \mathcal{M}\} = \mathcal{E}_{\mu,l}$. If any element of $\mathcal{E}_{\mu,l}$ is omitted, $\mathcal{E}_{\mu,l} \subsetneq \mathbf{A}^l \mathcal{M}$ and there exists a sequence of measurement noise realizations such that (22) is violated. As a consequence, the tube sets $\mathcal{E}_{\mu,l}$ are minimal. ■

The nominal state within the OCP is thus constrained as

$$\hat{\mathbf{z}}_{l|k} \in \hat{\mathcal{Z}}_l := \mathcal{Z}_l \ominus \mathcal{E}_{\mu,l} \quad \forall l \in \mathbb{N}_1^L. \quad (28)$$

Remark 5: Note that with $l \in \mathbb{N}_1^L$ we omit constraints on the initial state $\hat{\mathbf{z}}_{0|k} = \hat{\mathbf{x}}_k$ as it is not affected by future inputs, but fixed. For recursive feasibility and stability proofs or for the construction of a feasible candidate solution, it is often desired that the tightened nominal state constraints $\hat{\mathcal{Z}}$ contract, such that $\hat{\mathcal{Z}}_{l+1} \subseteq \hat{\mathcal{Z}}_l \quad \forall l \in \mathbb{N}_1^{L-1}$. This can be guaranteed by substituting $\mathcal{E}_{\mu,l}$ in (28) with the convex hull of $\bigcup_{i=1}^l \mathcal{E}_{\mu,i}$, albeit the tube may become more conservative.

C. TIGHTENED INPUT CONSTRAINTS

Both the additive disturbance and the measurement noise influence the measured state, which introduces uncertainty into actually applied inputs via the state feedback component. Similar to the constraints on the state in the previous subsection, we tighten the original input constraint set $\mathcal{U} = \{\mathbf{u} \in \mathbb{R}^{n_u} \mid \mathbf{G}_u \mathbf{u} \leq \mathbf{g}_u\}$ (4a) twice. As $\mathbf{z}_{0|k} = \hat{\mathbf{x}}_k = \mathbf{A} \mathbf{x}_{k-1} + \mathbf{B} \mathbf{u}_{k-1} + \mathbf{E} \mathbf{d}_{k-1} + \mu_k$, the additive disturbance \mathbf{d}_k does not influence the actually applied input $\mathbf{u}_{0|k}^*$, but only succeeding inputs $\mathbf{u}_{1|k}^*, \dots, \mathbf{u}_{L|k}^*$. Thus, with respect to the disturbance, a probabilistic constraint tightening may be employed to reduce conservatism. For all steps in the prediction horizon $l \in \mathbb{N}_0^L$, the probabilistically tightened input constraint sets \mathcal{U}_l are computed such that

$$\mathbf{v}_{l|k} + \mathbf{K} \mathbf{z}_{l|k} \in \mathcal{U}_l \Rightarrow \Pr(\mathbf{v}_{l|k} + \mathbf{K} \mathbf{x}_{l+k} \in \mathcal{U}) \geq p. \quad (29)$$

To that end, assume the dynamics (10b) and note that $\mathbf{K} \mathbf{x}_l = \mathbf{K} \mathbf{z}_l + \mathbf{K} \mathbf{e}_l$, which lets us split the expression on the right hand side of (29) into deterministic and probabilistic part by introducing a tightening parameter σ_l analogously to (13). As in (14), (15), the input constraint sets are then given by

$$\mathcal{U}_l = \{\mathbf{u} \in \mathbb{R}^{n_u} \mid \mathbf{G}_u \mathbf{u} \leq \sigma_l\}, \quad (30)$$

with the tightening parameter σ_l computed by solving the chance-constrained optimization problems

$$\sigma_l = \max_{\tilde{\sigma}} \tilde{\sigma} \quad (31a)$$

$$\text{s.t. } \Pr(\tilde{\sigma} \leq \mathbf{g}_u - \mathbf{G}_u \mathbf{K} \mathbf{e}_l) \geq p \quad (31b)$$

based on error samples \mathbf{e}_l computed from available disturbance samples via (21) as in the previous subsection.

Since the measurement noise μ_k directly influences the actually applied input $\mathbf{v}_{0|k} + \mathbf{K} \hat{\mathbf{x}}_k$, we further tighten the sets \mathcal{U}_l with respect to all possible realizations of μ_k . That is, for all steps in the prediction horizon, we employ input constraint sets $\hat{\mathcal{U}}_l$ constructed via

$$\hat{\mathcal{U}}_l = \mathcal{U}_l \ominus \mathbf{K} \mathcal{E}_{\mu,l} \quad \forall l \in \mathbb{N}_0^L, \quad (32)$$

which enforce $\mathbf{v}_{l|k} + \mathbf{K} \hat{\mathbf{z}}_{l|k} \in \hat{\mathcal{U}}_l \Rightarrow \mathbf{v}_{l|k} + \mathbf{K} \mathbf{z}_{l|k} \in \mathcal{U}_l$.

Remark 6: The tightened input constraint sets are such that $\hat{\mathcal{U}}_{l+1} \subseteq \hat{\mathcal{U}}_l \quad \forall l \in \mathbb{N}_0^{L-2}$ and the actual applied input always satisfies the specified constraints $\mathbf{u}_k \in \mathcal{U}$. Note that since $\mathbf{e}_0 = \mathbf{0}$, we have $\mathcal{U}_0 = \mathcal{U}$ and the constraints on the actually applied input only take into account the measurement noise.

D. PROPOSED OPTIMAL CONTROL PROBLEM AND ALGORITHM

Given the evolution of the nominal state (20a), tightened state and input constraint sets (28), (32), and the explicit expectation of the cost function, we obtain a tractable deterministic data-driven reformulation of the stochastic OCP (3)

$$\text{minimize}_{\hat{\mathbf{z}}_k, \mathbf{V}_k, \alpha_k} \sum_{l=0}^{L-1} J_s(\hat{\mathbf{z}}_{l|k}, \mathbf{v}_{l|k} + \mathbf{K} \hat{\mathbf{z}}_{l|k}) + J_f(\hat{\mathbf{z}}_{L|k}) \quad (33a)$$

$$\text{s.t. } \hat{\mathbf{z}}_{0|k} = \hat{\mathbf{x}}_k, \quad (33b)$$

$$\begin{bmatrix} \mathbf{v}_{[0|k, L|k]} \\ \mathbf{0} \\ \hat{\mathbf{z}}_{[0|k, L|k]} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_u - \tilde{\mathbf{K}} \mathbf{H}_x \\ \mathbf{H}_d \\ \mathbf{H}_x \end{bmatrix} \alpha_k, \quad (33c)$$

$$\hat{\mathbf{z}}_{l|k} \in \hat{\mathcal{Z}}_l \quad \forall l \in \mathbb{N}_1^L, \quad (33d)$$

$$\mathbf{v}_{l|k} + \mathbf{K} \hat{\mathbf{z}}_{l|k} \in \hat{\mathcal{U}}_l \quad \forall l \in \mathbb{N}_0^{L-1}, \quad (33e)$$

$$\hat{\mathbf{z}}_{L|k} \in \hat{\mathcal{Z}}_f, \quad (33f)$$

$$\hat{\mathbf{z}}_{1|k} \in \hat{\mathcal{Z}}_I, \quad (33g)$$

with prediction horizon L , predicted nominal state sequence $\hat{\mathcal{Z}}_k := \hat{\mathbf{z}}_{[0|k, L|k]}$, input sequence $\mathbf{V}_k := \mathbf{v}_{[0|k, L-1|k]}$, stage cost

$$J_s(\hat{\mathbf{z}}_{l|k}, \mathbf{u}_{l|k}) = \|\hat{\mathbf{z}}_{l|k}\|_{\mathbf{Q}}^2 + \|\mathbf{u}_{l|k}\|_{\mathbf{R}}^2, \quad (34)$$

with positive definite weighting matrices $\mathbf{Q} \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{R} \in \mathbb{R}^{n_u \times n_u}$, and a terminal cost function

$$J_f(\hat{\mathbf{z}}_{L|k}) = \|\hat{\mathbf{z}}_{L|k}\|_{\mathbf{P}}^2, \quad (35)$$

with positive definite $\mathbf{P} \in \mathbb{R}^{n_x \times n_x}$. (33f), (33g) are terminal and first-step constraints, which play an important role for recursive feasibility and stability of the proposed receding horizon control scheme in Section V. The OCP (33) has polyhedral constraints and quadratic costs, and can thus be written as a convex quadratic program. It is parameterized by the measured state $\hat{\mathbf{x}}_k$ and its decision variable is $\alpha_k \in \mathbb{R}^{N-L}$. As a consequence, and since \mathbf{R} is positive definite, it admits a unique solution α_k^* (and thereby \mathbf{V}_k^*) that depends entirely

Algorithm 1: Data-Driven Stochastic Predictive Control.

- Offline:** Given data as in Assumptions 1, 2.
- 1: Compute LQR feedback gain \mathbf{K} as in Remark 4.
 - 2: For each $l \in \mathbb{N}_1^L$, compute state errors $\mathbf{e}_l^{(i)}$ (21) from disturbance data $\mathbf{d}_{[0, l-1]}^{(i)}$ for all $i = 1, \dots, N_S$.
 - 3: Compute stochastic tube constraint sets for nominal state \mathcal{Z}_l (14) and input \mathcal{U}_l (30) by solving (15), (31).
 - 4: Robustly tighten $\mathcal{Z}_l, \mathcal{U}_l$ with respect to the bounded measurement noise, obtaining $\hat{\mathcal{Z}}_l$ (28) and $\hat{\mathcal{U}}_l$ (32).
 - 5: Compute \mathbf{P} (as in Section D of the supplementary material) and define terminal costs (35).
 - 6: Compute tightened terminal set $\hat{\mathcal{Z}}_f$ as in (37), (41).
 - 7: Compute the first step constraint $\hat{\mathcal{Z}}_1$ set as in (50).
 - 8: Construct the OCP $\mathbb{P}(\cdot)$ as in (33).
- Online:** For all time steps k :
- 9: Obtain noisy state measurement $\hat{\mathbf{x}}_k$.
 - 10: Solve the OCP $\mathbb{P}(\hat{\mathbf{x}}_k)$ (33) to retrieve $\mathbf{v}_{0|k}^*$.
 - 11: Apply $\mathbf{u}_k = \kappa(\hat{\mathbf{x}}_k) = \mathbf{K}\hat{\mathbf{x}}_k + \mathbf{v}_{0|k}^*$ (36) to the system.
 - 12: Set $k \leftarrow k + 1$ and go back to Step 9.

on $\hat{\mathbf{x}}_k$ (see for example [29, Chapter 7] for details). In the following, we denote the OCP (33) as $\mathbb{P}(\hat{\mathbf{x}}_k)$ and its implicit control law by

$$\kappa(\hat{\mathbf{x}}_k) := \mathbf{u}_k^* = \mathbf{K}\hat{\mathbf{x}}_k + \mathbf{v}_{0|k}^*. \quad (36)$$

Since the OCP is a convex program, small changes in the initial state $\hat{\mathbf{x}}_k$ lead to small changes in the optimal decision variable α_k^* and therefore the associated optimal input sequence \mathbf{V}_k^* .

Proposition 1: Let \mathcal{X}_0 be the set of initial states $\hat{\mathbf{x}}_k \in \mathcal{X}_0$ for which $\mathbb{P}(\hat{\mathbf{x}}_k)$ is solvable, i.e., \mathcal{X}_0 is the set of feasible initial states. For all $\hat{\mathbf{x}}_k \in \mathcal{X}_0$, the implicit control law $\kappa(\hat{\mathbf{x}}_k)$ is Lipschitz continuous.

Proof: Since the OCP itself is uncertainty free, has quadratic costs and polytopic constraints, Lipschitz continuity of its implicit control law is a well known property and follows for example from [28, Proposition 17]. Lipschitz continuity of $\kappa(\hat{\mathbf{x}}_k)$ follows from Lipschitz continuity of $\mathbf{K}\hat{\mathbf{x}}_k$. ■

The complete data-driven tube-based stochastic predictive control scheme is summarized in Algorithm 1.

V. RECURSIVELY FEASIBLE AND CLOSED-LOOP STABLE DESIGN

In this section, we design terminal costs, terminal constraints, and first-step constraints such that the OCP remains feasible and the resulting closed-loop system is stable. Stability, discussed in Section V-B, follows similarly to classical arguments in MPC by choosing the terminal constraint set as a robust positive invariant set under the local control law and the terminal cost function as a local control Lyapunov function in that terminal set. Recursive feasibility, discussed in Section V-A, follows from an additional first step constraint that guarantees robust positive invariance of the feasible set.

In the following, we will elaborate on the terminal constraint set, which plays a role in both feasibility and stability. Inspired by [27, Prop. 2] we first define a RPI feasible set \mathcal{X}_f under the control law $\mathbf{u}_k = \mathbf{K}\hat{\mathbf{x}}_k$. To that end, consider the set $\tilde{\mathcal{X}}_f$ of all states $\hat{\mathbf{x}}$ for which $\mathbf{K}\hat{\mathbf{x}}$ satisfies tightened input constraints, and the nominal successor state $\hat{\mathbf{x}}^+ = (\mathbf{A} + \mathbf{B}\mathbf{K})\hat{\mathbf{x}}$ is inside the tube (28).

$$\tilde{\mathcal{X}}_f = \left\{ \hat{\mathbf{x}} \in \mathbb{R}^{n_x} \left| \begin{array}{l} (\exists \alpha_f \in \mathbb{R}^{N-L}) \\ \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \hat{\mathbf{x}} \\ \hat{\mathbf{x}}^+ \end{bmatrix} = \begin{bmatrix} [\mathbf{H}_u - \tilde{\mathbf{K}}\mathbf{H}_x]_{[1, n_u]} \\ [\mathbf{H}_d]_{[1, n_d]} \\ [\mathbf{H}_x]_{[1, n_x]} \\ [\mathbf{H}_x]_{[n_x+1, 2n_x]} \end{bmatrix} \alpha_f \\ \hat{\mathbf{x}}^+ \in \hat{\mathcal{Z}}_1, \mathbf{K}\hat{\mathbf{x}} \in \hat{\mathcal{U}}_0 \end{array} \right. \right\}. \quad (37)$$

Proposition 2 (RPI feasible set under local control): Let $\mathcal{X}_f \subseteq \tilde{\mathcal{X}}_f$ be a RPI polytope for system (2) controlled via $\mathbf{u}_k = \mathbf{K}\hat{\mathbf{x}}_k$, i.e., $\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{K}(\mathbf{x}_k + \boldsymbol{\mu}_k) + \mathbf{d}_k$. Then, for any measured initial state $\hat{\mathbf{x}}_0 \in \mathcal{X}_f$, the original state constraints (3e) and input constraints (3d) are satisfied for all $k > 0$ in closed loop.

Proof: By definition of $\tilde{\mathcal{X}}_f$, a measured state $\hat{\mathbf{x}} \in \mathcal{X}_f$ implies $\hat{\mathbf{x}}^+ \in \hat{\mathcal{Z}}_1$, which in turn implies $\mathbf{x}^+ \in \mathcal{Z}_1$ for the actual (not measured) state by definition of the tube set (28). By construction of \mathcal{Z}_1 in (14), $\mathbf{x}^+ \in \mathcal{Z}_1$ implies satisfaction of the state constraints $\Pr(\mathbf{x}^+ \in \mathcal{X}) \geq p$. By robust positive invariance of \mathcal{X}_f , $\hat{\mathbf{x}}^+ \in \mathcal{X}_f$ and the argument holds for all future time steps by induction. Similarly, $\hat{\mathbf{x}} \in \mathcal{X}_f$ implies $\mathbf{K}\hat{\mathbf{x}} \in \hat{\mathcal{U}}_0$ so that input constraints $\mathbf{u} = \mathbf{K}\hat{\mathbf{x}} \in \mathcal{U}$ are satisfied since $\hat{\mathcal{U}}_0 \subseteq \mathcal{U}$. Again, robust positive invariance of \mathcal{X}_f guarantees that $\hat{\mathbf{x}}^+ \in \mathcal{X}_f$ and the constraints hold for all time. ■

We use \mathcal{X}_f to specify the terminal constraint set $\hat{\mathcal{Z}}_f$ for the last nominal state $\hat{\mathbf{z}}_{L|k}$ in the OCP such that

$$\hat{\mathbf{z}}_{L|k} \in \hat{\mathcal{Z}}_f \Rightarrow \Pr(\mathbf{x}_{L|k} \in \mathcal{X}_f) \geq p \quad (38)$$

for all realizations of the disturbance and measurement noise. $\hat{\mathcal{Z}}_f$ is computed by first probabilistically tightening \mathcal{X}_f with respect to the additive disturbance as in (14), and then robustly tightening with respect to the measurement noise. Let $\mathcal{X}_f = \{\mathbf{x} \in \mathbb{R}^{n_x} \mid \mathbf{G}_{x,f}\mathbf{x} \leq \boldsymbol{\eta}_{x,f}\}$, then

$$\mathcal{Z}_f = \{\mathbf{z} \in \mathbb{R}^{n_x} \mid \mathbf{G}_{x,f}\mathbf{z} \leq \boldsymbol{\eta}_f\}, \quad (39)$$

where $\boldsymbol{\eta}_f$ solves the chance constrained optimization

$$\boldsymbol{\eta}_f = \max_{\tilde{\boldsymbol{\eta}}} \tilde{\boldsymbol{\eta}} \quad (40a)$$

$$\text{s.t.} \quad \Pr(\tilde{\boldsymbol{\eta}} \leq \mathbf{g}_{x,f} - \mathbf{G}_{x,f}\mathbf{e}_L) \geq p, \quad (40b)$$

and the robustly tightened terminal set is then defined as

$$\hat{\mathcal{Z}}_f := \mathcal{Z}_f \ominus \mathcal{E}_{\mu,L}, \quad (41)$$

where $\mathcal{E}_{\mu,L}$ encompasses all uncertainty induced by the measurement noise and is given by (27).

A. RECURSIVE FEASIBILITY

In order to guarantee recursive feasibility (Definition 2), we need to show that the OCP $\mathbb{P}(\hat{\mathbf{x}}_{k+1})$ at the next time step is feasible if the OCP $\mathbb{P}(\hat{\mathbf{x}}_k)$ at the current time step is feasible. Since input and state sequences within the OCP (33) are uniquely determined by the decision vector $\boldsymbol{\alpha}_k$ in (33c), we frame recursive feasibility as the problem of ensuring the existence of a feasible $\boldsymbol{\alpha}_{k+1}$ at the next time step. To that end, let the set $\mathcal{A}_F(\hat{\mathbf{x}}_k)$ denote all feasible decision variables $\boldsymbol{\alpha}_k$ for $\mathbb{P}(\hat{\mathbf{x}}_k)$, i.e., $\mathcal{A}_F(\hat{\mathbf{x}}_k) = \{\boldsymbol{\alpha}_k \in \mathbb{R}^{N-L} \mid (33b)–(33g) \text{ are satisfied}\}$.

Proposition 3 (Recursive feasibility in the space of $\boldsymbol{\alpha}$): The receding horizon OCP (33) for system (2) under the proposed control law (36) is recursively feasible according to Definition 2 if and only if

$$\mathcal{A}_F(\hat{\mathbf{x}}_k) \neq \emptyset \Rightarrow \mathcal{A}_F(\hat{\mathbf{x}}_{k+1}) \neq \emptyset. \quad (42)$$

Remark 7: Next to the disturbance \mathbf{d}_k , both measurement noise realizations $\boldsymbol{\mu}_k, \boldsymbol{\mu}_{k+1}$ influence $\mathcal{A}_F(\hat{\mathbf{x}}_{k+1})$, since $\boldsymbol{\mu}_k$ influences the control input $\boldsymbol{\kappa}(\hat{\mathbf{x}}_k) = \boldsymbol{\kappa}(\mathbf{x}_k + \boldsymbol{\mu}_k)$, which in turn influences the successor state \mathbf{x}_{k+1} , and thereby the initial state $\hat{\mathbf{x}}_{k+1} = \mathbf{x}_{k+1} + \boldsymbol{\mu}_{k+1}$ of the next OCP.

In nominal MPC, recursive feasibility can be guaranteed with a control invariant terminal constraint set $\hat{\mathcal{Z}}_f$. The usual argument, for example presented in [35], is based on a set \mathcal{C}_l that denotes all states from which the terminal set can be reached in l steps without violation of input or state constraints. Since the terminal set is control invariant, $\mathbf{x} \in \mathcal{C}_l$ implies $\mathbf{x} \in \mathcal{C}_{l+1}$, and therefore $\mathcal{C}_l \subseteq \mathcal{C}_{l+1}$. If $\mathbf{x} \in \mathcal{C}_l$ and we apply the first input of any admissible input sequence, the successor state \mathbf{x}^+ is in $\mathcal{C}_{l-1} \subseteq \mathcal{C}_l$. Thus \mathcal{C}_l is positive invariant under the MPC law (yielding admissible inputs), and for all $\mathbf{x}_0 \in \mathcal{C}_l$ the OCP remains feasible for all time.

In the case of disturbed (as opposed to nominal) systems, it is not guaranteed that $\mathbf{z} = \mathbf{x} \in \mathcal{C}_l \Rightarrow \mathbf{z}^+ = \mathbf{x}^+ \in \mathcal{C}_{l-1}$ holds for the initial states of consecutive OCPs, since the actual next state \mathbf{x}^+ is not equal to the next predicted nominal state. For probabilistically tightened constraint sets, this implication is nontrivial to restore [36]. In the model-based setting, a remedy presented in [27] is to directly ensure robust positive control invariance of \mathcal{C}_L (or equivalently robust positive invariance of \mathcal{C}_L for the closed loop system under MPC law) by introducing an additional constraint on the first step of the prediction horizon.

If the initial state of the OCP is additionally perturbed by measurement noise, the above implication $\hat{\mathbf{x}}_k \in \mathcal{C}_l \Rightarrow \hat{\mathbf{x}}_{k+1} \in \mathcal{C}_{l-1}$ depends on the noise $\boldsymbol{\mu}_k, \boldsymbol{\mu}_{k+1}$. In the following, we construct a first step constraint from data and include robustness with respect to measurement noise. To that end, denote the set of feasible initial states by

$$\mathcal{C}_L = \{\hat{\mathbf{z}}_{0|k} \in \mathbb{R}^{n_x} \mid (\exists \boldsymbol{\alpha}_z \in \mathbb{R}^{N-L}) (33b)–(33f)\}. \quad (43)$$

Remark 8: In the model-based setting, \mathcal{C}_L can be computed via backward recursion [37] based on the nominal dynamics of the system and set algebra. In the data-driven setting, equivalently, \mathcal{C}_L may be computed by representing the system

matrices $(\mathbf{A}, \mathbf{B}, \mathbf{E})$ with data, based on a suitable extension of [3, Theorem 1] to include disturbance data. This implicit system identification step can be avoided by constructing the set of feasible $\boldsymbol{\alpha}$, $\mathcal{A}_z = \{\boldsymbol{\alpha}_z \in \mathbb{R}^{N-L} \mid (33b)–(33f)\}$, and projecting onto the first step as $\mathcal{C}_L = [\mathbf{H}_z]_{[1, n_x]} \mathcal{A}_z$.

Since recursive feasibility depends on the measured state $\hat{\mathbf{x}}_k$, see (42), we require robust positive control invariance of (a subset of) \mathcal{C}_L with respect to the evolution of $\hat{\mathbf{x}}_k$,

$$\hat{\mathbf{x}}_{k+1} = \mathbf{A}(\hat{\mathbf{x}}_k - \boldsymbol{\mu}_k) + \mathbf{B}\mathbf{u}_k + \mathbf{E}\mathbf{d}_k + \boldsymbol{\mu}_{k+1}. \quad (44)$$

That is, the control invariant set needs to be robust with respect to the extended disturbance term $\mathbf{w}_k := \mathbf{E}\mathbf{d}_k - \mathbf{A}\boldsymbol{\mu}_k + \boldsymbol{\mu}_{k+1}$, $\mathbf{w}_k \in \mathcal{W} = \mathbf{E}\mathcal{D} \oplus (-\mathbf{A}\mathcal{M}) \oplus \mathcal{M}$. The disturbance support set \mathcal{W} can be computed from data as

$$\mathcal{W} = \mathcal{E}_{d,1} \oplus (-\mathcal{E}_{\mu,1}) \oplus \mathcal{M}, \quad (45)$$

where $\mathcal{E}_{\mu,1}$ is as in (27) and $\mathcal{E}_{d,1}$, equal to $\mathbf{E}\mathcal{D}$, is obtained by defining the appropriate set of decision variables

$$\begin{aligned} \mathcal{A}_e &= \{\boldsymbol{\alpha} \in \mathbb{R}^{N-L} \mid [\mathbf{H}_u]_{[1, n_u]} \boldsymbol{\alpha} = \mathbf{0}, \\ &\mathbf{G}_d [\mathbf{H}_d]_{[1, n_o]} \boldsymbol{\alpha} \leq \mathbf{g}_d, [\mathbf{H}_x]_{[1, n_x]} \boldsymbol{\alpha} = \mathbf{0}\}, \end{aligned} \quad (46)$$

with zero input and initial state, and then projecting onto the first step state errors $\mathbf{e}_{1|k}$ via

$$\mathcal{E}_{d,1} = [\mathbf{H}_x]_{[n_x+1, 2n_x]} \mathcal{A}_e. \quad (47)$$

Based on \mathcal{W} (45), a robust control invariant subset of \mathcal{C}_L can be constructed as

$$\mathcal{C}_L^\infty = \bigcap_{i=0}^\infty \mathcal{C}_L^i, \quad (48)$$

where $\mathcal{C}_L^0 = \mathcal{C}_L$ and

$$\mathcal{C}_L^{i+1} = \left\{ \hat{\mathbf{x}} \in \mathcal{C}_L^i \left[\begin{array}{l} \exists \boldsymbol{\alpha} \in \mathbb{R}^{N-L}, \mathbf{u} \in \hat{\mathcal{U}}_0 : \\ \left[\begin{array}{c} \mathbf{u} \\ \mathbf{0} \\ \hat{\mathbf{x}} \\ \hat{\mathbf{x}}^+ \end{array} \right] = \left[\begin{array}{c} [\mathbf{H}_u]_{[1, n_u]} \\ [\mathbf{H}_d]_{[1, n_o]} \\ [\mathbf{H}_x]_{[1, n_x]} \\ [\mathbf{H}_x]_{[n_x+1, 2n_x]} \end{array} \right] \boldsymbol{\alpha} \\ \hat{\mathbf{x}} \in \mathcal{C}_L^i, \hat{\mathbf{x}}^+ \in \mathcal{C}_L^i \ominus \mathcal{W} \end{array} \right\}. \quad (49)$$

The idea in (49) is to tighten the set \mathcal{C}_L , until all contained states admit an input for which the successor state is in $\mathcal{C}_L^i \ominus \mathcal{W}$, far enough away from the boundary of the set to render it RPI with respect to all uncertainty in \mathcal{W} .

Remark 9: In the model-based setting, an equivalent computation (without measurement noise) has been used in [27] and the idea of the sequence is presented in [38, Chapter 5.3]. The set \mathcal{C}_L^∞ (48) is obtained by recursively computing \mathcal{C}_L^i (49), until $\mathcal{C}_L^i = \mathcal{C}_L^{i+1}$ for some $i \in \mathbb{N}$, which implies that $\mathcal{C}_L^\infty = \mathcal{C}_L^i$ is robust positive control invariant. In practice, the equality is relaxed to hold with a tolerance, see [37, Chapter 5.3] for details.

The constraint on the first predicted state $\hat{\mathbf{z}}_{1|k}$,

$$\hat{\mathbf{z}}_{1|k} \in \hat{\mathcal{Z}}_f = \mathcal{C}_L^\infty \ominus \mathcal{W}, \quad (50)$$

guarantees that for all possible realizations of the disturbance \mathbf{d}_k and measurement noises $\boldsymbol{\mu}_k, \boldsymbol{\mu}_{k+1}$, the next measured state

$\hat{\mathbf{x}}_{k+1}$ is inside the robust positive control invariant subset \mathcal{C}_L^∞ of the L -step reachable initial state set \mathcal{C}_L .

Theorem 1 (Recursive feasibility of the proposed OCP): Let the first step constraint set of the receding horizon OCP (33) be given by (50). The receding horizon OCP is recursively feasible, and if $\mathbf{x}_0 \in \mathcal{C}_L^\infty \ominus \mathcal{M}$, the receding horizon OCP is feasible for all time.

Proof: If $\mathbf{x}_0 \in \mathcal{C}_L^\infty \ominus \mathcal{M}$, then $\hat{\mathbf{x}}_0 \in \mathcal{C}_L^\infty \subseteq \mathcal{C}_L$ and the OCP $\mathbb{P}(\hat{\mathbf{x}}_0)$ is feasible by construction of $\mathcal{C}_L = \{\hat{\mathbf{x}}_k \in \mathbb{R}^{n_x} \mid \mathcal{A}_F(\hat{\mathbf{x}}_k) \neq \emptyset\}$. For recursive feasibility, assume that the OCP is feasible at time step k with optimizer α_k^* such that $\alpha_k^* \in \mathcal{A}_F(\hat{\mathbf{x}}_k)$. By construction of the first-step constraint $\hat{\mathbf{z}}_{1|k}^* = [\mathbf{H}_x]_{[n_x+1, 2n_x]} \alpha_k^* \in \mathcal{C}_L^\infty \ominus \mathcal{W}$, the next measured state satisfies $\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{z}}_{1|L} + \mathbf{E}d_k - \mathbf{A}\mu_k + \mu_{k+1} \in \mathcal{C}_L^\infty$ for all possible realizations of μ_k, μ_{k+1}, d_k . Since $\mathcal{C}_L^\infty \subseteq \mathcal{C}_L$, $\mathcal{A}_F(\hat{\mathbf{x}}_{k+1}) \neq \emptyset$ and the OCP is recursively feasible by Proposition 3. ■

Theorem 2 (Closed loop chance constraint satisfaction): The proposed control law (Algorithm 1) leads to the satisfaction of chance constraints (3e) in closed loop for a risk parameter $p \in [p_{\min}, p_{\max}]$ with confidence $1 - \beta$.

Proof: From the constraints of the OCP, $\hat{\mathbf{z}}_{1|k}^* \in \hat{\mathcal{Z}}_1$. By Lemma 4, it follows that $\mathbf{z}_{1|k} \in \mathcal{Z}_1$, so that chance constraints are satisfied with \mathcal{Z}_1 defined as in (14), (15). Since the chosen solution method of the chance-constrained optimization problem (15) guarantees (15b) with confidence $1 - \beta$, satisfaction of chance constraints for \mathbf{x}_{k+1} is inherited with confidence $1 - \beta$ for a risk parameter p in the specified interval. The proof for the closed loop follows from recursive feasibility of the OCP, Theorem 1. ■

Remark 10: If the tightening parameter η is chosen such that (15b) holds with certainty, closed loop chance constraint satisfaction is guaranteed with certainty.

B. INPUT-TO-STATE STABILITY

In order to establish ISS according to Definition (8), we use a special ISS-Lyapunov function as proposed in [28], [39].

Definition 5 (ISS-Lyapunov function): A function $V : \mathcal{X}_0 \rightarrow \mathbb{R}_0^+$ is an ISS-Lyapunov function for system $\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{w}_k)$ if there exist functions $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ and $\gamma \in \mathcal{K}$ such that for all $\mathbf{x} \in \mathcal{X}_0, \mathbf{w} \in \mathcal{W}$

$$\alpha_1(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \alpha_2(\|\mathbf{x}\|), \quad (51a)$$

$$V(\mathbf{f}(\mathbf{x}, \mathbf{w})) - V(\mathbf{x}) \leq -\alpha_3(\|\mathbf{x}\|) + \gamma(\|\mathbf{w}\|). \quad (51b)$$

For more details on ISS and ISS-Lyapunov functions, see Section C of the supplementary material. In particular, we use the following corollary, which is a direct result of the Lipschitz continuity of linear functions.

Corollary 1: Consider a closed loop LTI system $\mathbf{f}(\mathbf{x}_k, \mathbf{w}_k) = \mathbf{A}\mathbf{x}_k + \mathbf{B}\kappa(\mathbf{x}_k) + \mathbf{w}_k$ with disturbance \mathbf{w}_k from a compact set \mathcal{W} . Let the set \mathcal{X}_0 contain the origin in its interior, be robust positive invariant for the system $\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{w}_k)$, and let $V : \mathcal{X}_0 \rightarrow \mathbb{R}_0^+$ be an ISS-Lyapunov function for the undisturbed system $\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{0})$. The closed loop

$\mathbf{f}(\mathbf{x}_k, \mathbf{w}_k)$ is ISS with respect to disturbances $\mathbf{w}_k \in \mathcal{W}$ if $\kappa(\mathbf{x}_k)$ is Lipschitz continuous on \mathcal{X}_0 .

In the following, we use Corollary 1 to proof ISS of system (2a) under the proposed control law (36),

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\kappa(\mathbf{x}_k + \mu_k) + \mathbf{E}d_k. \quad (52)$$

Corollary 1 requires Lipschitz continuity of the control law. However, the measurement noise μ_k renders $\kappa(\mathbf{x}_k + \mu_k)$ discontinuous in \mathbf{x}_k . Consequently, direct stability proofs for the evolution of the state \mathbf{x}_k as in (52) are challenging. Therefore, as in [30], we first analyze stability for the equivalent evolution in terms of the state measurements (44),

$$\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{f}}(\hat{\mathbf{x}}_k, \mathbf{w}_k) = \mathbf{A}\hat{\mathbf{x}}_k + \mathbf{B}\kappa(\hat{\mathbf{x}}_k) + \mathbf{w}_k. \quad (53)$$

For the artificial noisy system (53), the control law $\kappa(\hat{\mathbf{x}}_k)$ is Lipschitz continuous in the state $\hat{\mathbf{x}}_k$ of the system by Proposition 1. By Corollary 1 then, system (53) is ISS with respect to \mathbf{w} if there is a Lipschitz continuous ISS-Lyapunov function for the uncertainty-free system

$$\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{f}}(\hat{\mathbf{x}}_k, \mathbf{0}) = \mathbf{A}\hat{\mathbf{x}}_k + \mathbf{B}\kappa(\hat{\mathbf{x}}_k). \quad (54)$$

System (54) coincides with the nominal dynamics (10a) assumed in the OCP. As a consequence, if the predictive control scheme were applied to system (54), the measured state at the next time step would be equal to the first predicted state (of the optimal state sequence),

$$\hat{\mathbf{z}}_{1|k}^* = \hat{\mathbf{x}}_{k+1} = \hat{\mathbf{z}}_{0|k+1}. \quad (55)$$

This equality allows us to express (cost) functions of the successor state $\hat{\mathbf{x}}_{k+1}$ in terms of the optimal solution of the OCP $\mathbb{P}(\hat{\mathbf{x}}_k)$, and thereby upper bound the descent of the cost function from time step k to $k + 1$. In the following, we show that the optimal cost J_L^* of the OCP,

$$J_L^*(\hat{\mathbf{x}}_k) = \sum_{l=0}^{L-1} J_s(\hat{\mathbf{z}}_{l|k}^*, \mathbf{u}_{l|k}^*) + J_f(\hat{\mathbf{z}}_{L|k}^*), \quad (56)$$

is an ISS-Lyapunov function for system (54) with proposed control law $\kappa(\hat{\mathbf{x}}_k) = \mathbf{K}\hat{\mathbf{x}}_k + \mathbf{v}_{0|k}^*$ (36), if the terminal cost function (35) is a Lyapunov function for system (54) under pure state feedback $\kappa(\hat{\mathbf{x}}_k) = \mathbf{K}\hat{\mathbf{x}}_k$.

Proposition 4 (Terminal cost function): Let \mathbf{P} be the optimal cost-to-go matrix of the LQR problem associated with the state feedback gain \mathbf{K} . The terminal cost function (35) is a Lyapunov function in the terminal set \mathcal{Z}_f (41) for the closed-loop system $\mathbf{x}_{k+1} = (\mathbf{A} + \mathbf{B}\mathbf{K})\mathbf{x}_k$, such that

$$J_f(\mathbf{x}_{k+1}) - J_f(\mathbf{x}_k) = -J_s(\mathbf{x}_k, \mathbf{K}\mathbf{x}_k). \quad (57)$$

The optimal cost-to-go matrix \mathbf{P} may be computed from data, see Section D of the supplementary material. The proof of Proposition (4) is standard in MPC and included in Section E of the supplementary material.

We proof the next result with the help of an explicit candidate solution. Given a feasible solution $\mathbf{U}_k^* = (\mathbf{u}_{0|k}^*, \mathbf{u}_{1|k}^*, \dots, \mathbf{u}_{L-1|k}^*)$, $\mathbf{u}_{i|k}^* = \mathbf{v}_{i|k}^* + \mathbf{K}\hat{\mathbf{z}}_{i|k}^*$, to the OCP (33) at time step k , a candidate solution for

the next time step $k + 1$ can be constructed by shifting the input trajectory and appending $\mathbf{K}\hat{\mathbf{z}}_{L|k}^*$. That is, a feasible, likely suboptimal, candidate solution for the next time step is $\tilde{\mathbf{U}}_{k+1} = (\mathbf{u}_{1|k}^*, \mathbf{u}_{2|k}^*, \dots, \mathbf{u}_{L-1|k}^*, \mathbf{K}\hat{\mathbf{z}}_{L|k}^*)$. Since predictions are exact (55) for the uncertainty-free system (54), we have $\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{z}}_{1|k}^*$ and the candidate state sequence is $\tilde{\mathbf{z}}_{k+1} = (\hat{\mathbf{z}}_{1|k}^*, \hat{\mathbf{z}}_{2|k}^*, \dots, \hat{\mathbf{z}}_{L|k}^*, \tilde{\mathbf{z}}_{L|k+1})$ with $\tilde{\mathbf{z}}_{L|k+1} = (\mathbf{A} + \mathbf{BK})\hat{\mathbf{z}}_{L|k}^* = [\mathbf{H}_x]_{[Ln_x+1, Ln_x+n_x]} \tilde{\boldsymbol{\alpha}}_{k+1}$ and

$$\tilde{\boldsymbol{\alpha}}_{k+1} = \begin{bmatrix} \mathbf{H}_u \\ [\mathbf{H}_x]_{[1, n_x]} \end{bmatrix}^\dagger \begin{bmatrix} \tilde{\mathbf{U}}_{k+1} \\ \hat{\mathbf{z}}_{1|k}^* \end{bmatrix}.$$

Based on Proposition 4, the difference between the terminal cost of the solution of the OCP at time step k and the terminal cost of the explicit candidate solution at the next time step $k + 1$ is upper bounded via

$$J_f(\tilde{\mathbf{z}}_{L|k+1}) - J_f(\hat{\mathbf{z}}_{L|k}^*) \leq -J_s(\hat{\mathbf{z}}_{L|k}^*, \mathbf{K}\hat{\mathbf{z}}_{L|k}^*). \quad (58)$$

Lemma 5 (ISS of the artificial system): The optimal cost J_L^* (56) is an ISS-Lyapunov function for system (54). Furthermore, J_L^* is Lipschitz continuous on \mathcal{C}_L^∞ . As a consequence, system (53) is ISS with region of attraction \mathcal{C}_L^∞ .

Proof: J_L^* is positive definite, defined on a domain that contains the origin with $J_L^*(\mathbf{0}) = 0$ and continuous by Proposition 1. By [40, Lemma 4.3], there exist functions $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ such that condition (51a) holds for $V = J_L^*$. Condition (51b) is $J_L^*(\hat{f}(\hat{\mathbf{x}}_k, \mathbf{0})) - J_L^*(\hat{\mathbf{x}}_k) \leq -\alpha(\|\hat{\mathbf{x}}_k\|)$ for some function $\alpha_3 \in \mathcal{K}_\infty$, which we show with the help of the suboptimal candidate solution for the next time step $k + 1$. The cost of the candidate solution, denoted by $\tilde{J}_L(\hat{\mathbf{x}}_{k+1})$, is

$$\sum_{l=1}^{L-1} J_s(\hat{\mathbf{z}}_{l|k}^*, \mathbf{u}_{l|k}^*) + J_s(\hat{\mathbf{z}}_{L|k}^*, \mathbf{K}\hat{\mathbf{z}}_{L|k}^*) + J_f(\tilde{\mathbf{z}}_{L|k+1}). \quad (59)$$

Since $J_L^*(\hat{\mathbf{x}}_{k+1}) \leq \tilde{J}_L(\hat{\mathbf{x}}_{k+1})$, the descent of the optimal cost function is upper bounded by $J_L^*(\hat{\mathbf{x}}_{k+1}) - J_L^*(\hat{\mathbf{x}}_k) \leq \tilde{J}_L(\hat{\mathbf{x}}_{k+1}) - J_L^*(\hat{\mathbf{x}}_k) \stackrel{(55)}{=} -J_s(\hat{\mathbf{x}}_k, \boldsymbol{\kappa}(\hat{\mathbf{x}}_k)) + J_s(\hat{\mathbf{z}}_{L|k}^*, \mathbf{K}\hat{\mathbf{z}}_{L|k}^*) +$

$J_f(\tilde{\mathbf{z}}_{L|k+1}) - J_f(\hat{\mathbf{z}}_{L|k}^*) \stackrel{(58)}{\leq} -J_s(\hat{\mathbf{x}}_k, \boldsymbol{\kappa}(\hat{\mathbf{x}}_k))$. Since $\boldsymbol{\kappa}(\hat{\mathbf{x}}_k)$ is continuous in $\hat{\mathbf{x}}_k$ by Proposition 1, and $J_s(\hat{\mathbf{x}}_k, \mathbf{u}_k)$ (34) is continuous in $\hat{\mathbf{x}}_k$ and \mathbf{u}_k , the composition $J_s(\hat{\mathbf{x}}_k, \boldsymbol{\kappa}(\hat{\mathbf{x}}_k))$ is continuous in $\hat{\mathbf{x}}_k$. Hence, there exists $\alpha_3 \in \mathcal{K}_\infty$ such that $\alpha_3(\|\hat{\mathbf{x}}_k\|) \leq J_s(\hat{\mathbf{x}}_k, \boldsymbol{\kappa}(\hat{\mathbf{x}}_k))$ by [40, Lemma 4.3]. As a consequence, $J_L^*(\hat{\mathbf{x}}_{k+1}) - J_L^*(\hat{\mathbf{x}}_k) \leq -\alpha_3(\|\hat{\mathbf{x}}_k\|)$. With condition (51a) and (51b) satisfied for $\mathbf{w} = 0$, J_L^* is an ISS-Lyapunov function for system (54).

By construction, \mathcal{C}_L^∞ (48) is compact, contains the origin, and is robust positive invariant for system (53) with respect to all possible disturbances $\mathbf{w}_k \in \mathcal{W}$. Since the quadratic optimal cost $J_L^*(\cdot)$ is Lipschitz continuous on the compact region of attraction \mathcal{C}_L^∞ [28, Proposition 17], system (53) is ISS by Corollary 1. ■

Lemma 5 shows that the artificial system (53) is input-to-state stable under the proposed control law (36). ISS for the actual closed loop system (52) follows from the measurement equation $\hat{\mathbf{x}}_k = \mathbf{x}_k + \boldsymbol{\mu}_k$ and the definition of input-to-state

stability, as in [30, Theorem 3.3], based on properties of comparison functions.

Theorem 3 (ISS of the closed loop system): System (52) is ISS under the proposed control law (36), with region of attraction $\mathcal{X}_0 = \mathcal{C}_L^\infty \ominus \mathcal{M}$.

Proof: The proof follows from [30, Theorem 3.3], since both necessary Assumptions [30, Assumptions 3.1, 3.2] hold with $\mathbf{Ax} + \mathbf{Bu}$ Lipschitz continuous in \mathbf{x} , \mathcal{W} compact and therefore contained in a closed ball of some radius λ , the artificial system (53) ISS on \mathcal{X}_0 w.r.t. the extended additive disturbance $\mathbf{w} \in \mathcal{W}$, and \mathcal{X}_0 robust positive invariant w.r.t. the extended additive disturbance by construction. For completeness, a self-contained proof is included in Section F of the supplementary material. ■

In summary, we have shown that the proposed receding horizon OCP is recursively feasible, and the resulting data-driven predictive controller (36) renders the closed loop system ISS with respect to both additive disturbances and measurement noise.

VI. ON THE CASE OF INEXACT DATA

The availability of exact trajectory data is a standing assumption in this work (Assumption 1). In essence, we take a *certainty-equivalence* approach, i.e., we assume that the data-driven representation captures the *true* system, comparable to the case where system matrices $(\mathbf{A}, \mathbf{B}, \mathbf{E})$ are known. The proposed control scheme and the proof of its properties were nontrivial since we investigated a realistic setting *during* the control phase, with the system state influenced by unknown disturbances and only accessible via noisy measurements. Even though offline data can be averaged, filtered, or similarly de-noised, in practice, data are still likely inexact, and Assumption 1 is likely violated. As a consequence, the hypothesis of Lemma 2 no longer holds.

In that case, there are four consequences for the presented control scheme: 1) predictions within the OCP are inexact even for the nominal system; 2) the computed constraint sets are not as specified and may lose their declared properties; 3) the computed LQR gain \mathbf{K} is not the actual LQR gain associated with specified \mathbf{Q} and \mathbf{R} , and may even be non-stabilizing; 4) the computed solution of the algebraic Riccati equation \mathbf{P} is (likely) no longer associated with \mathbf{K} via the Lyapunov equation. Problem 1) may be alleviated by a bound on the prediction error, which could be accounted for in a further constraint tightening. Specifying such a bound is a pressing open challenge in the data-driven control literature. Problem 2) may be alleviated at the cost of conservatism by guaranteeing either under- or overapproximation of the original sets, whichever is appropriate for the resulting properties. Problem 3) is not as crucial since only optimality is lost as long as \mathbf{K} is still stabilizing. In [41], Dörfler et al. present a regularized data-driven LQR which leads to stabilizing controllers when the signal-to-noise ratio is large. Problem 4) has implications related to stability: the solution of the algebraic Riccati equation \mathbf{P} is used to guarantee that the terminal cost function is




FIGURE 2. Double-mass-spring-damper model used for the simulation. The states are angles and angular velocities of the masses, the input is the torque on mass 1, and the disturbance is the torque on mass 2.

a Lyapunov function for the nominal system under control $u = Kx$. This in turn lets us bound the descent of the stage cost. A similar bound may still be obtained without P being accurate.

In practice, deviations from the case of exact data can be reduced by implementing minor regularizing adjustments, the benefits of which have been commonly observed in the literature on data-driven control with inexact data [19], [41], [42]. In particular, it is advantageous to penalize large norms of the decision variable α , because any noise inside the Hankel matrix (33c) is amplified directly by α . In line with [41], we propose to add the term $\lambda_\alpha \|\Pi\alpha_k\|$ to the cost function (33a), where $\lambda_\alpha > 0$ is a regularization parameter and Π is a matrix defined as

$$\Pi := I - H_{\text{udx}}^\dagger H_{\text{udx}}, \quad H_{\text{udx}} := \begin{bmatrix} H_u - \tilde{K}H_x \\ H_d \\ [H_x]_{[1, n]} \end{bmatrix}. \quad (60)$$

With the above definition, $\|\Pi\alpha\|$ is a measure for the distance of α to the image of H_{udx}^\dagger . For λ_α sufficiently large, the equality $\Pi\alpha_k = \mathbf{0}$ holds, which means that $\alpha_k \in \text{image } H_{\text{udx}}^\dagger$ [41]. The presented sets of admissible decision variables α , such as \mathcal{A}_μ (27) of Lemma 4 and \mathcal{A}_e in (46) may be (partially) de-noised by enforcing $\Pi\alpha = \mathbf{0}$ in the definition of each set.

Similarly, instead of directly computing state errors from disturbance samples with the pseudo-inverse in (21), α_e may first be obtained by employing a convex cost function $l_\alpha(\cdot) : \mathbb{R}^{N-L} \rightarrow \mathbb{R}_{\geq 0}$ other than $\|\alpha_e\|_2$ and solving

$$\alpha_e = \arg \min_{\tilde{\alpha}_e} l_\alpha(\tilde{\alpha}_e) \quad (61)$$

$$\text{s.t.} \quad \begin{bmatrix} \mathbf{0} \\ \underline{d}_{[0, L]} \\ e_0 \end{bmatrix} = \begin{bmatrix} H_u - \tilde{K}H_x \\ H_d \\ [H_x]_{[1, n]} \end{bmatrix} \tilde{\alpha}_e. \quad (62)$$

The computation of the LQR state feedback gain may be regularized in similar fashion, see Section B of the supplementary material.

VII. SIMULATION EXAMPLE

In this section, we present a simulation example and demonstrate the effectiveness of the proposed control algorithm. Consider the double-mass-spring-damper system, shown in

Fig. 2. The discrete-time dynamics of the system read

$$x_{k+1} = \begin{bmatrix} 0.952 & 0.048 & 0.094 & 0.002 \\ 0.048 & 0.952 & 0.002 & 0.094 \\ -0.920 & 0.920 & 0.859 & 0.046 \\ 0.920 & -0.920 & 0.046 & 0.858 \end{bmatrix} x_k + \begin{bmatrix} 0.048 \\ 0.001 \\ 0.936 \\ 0.016 \end{bmatrix} u_k + \begin{bmatrix} 0.001 \\ 0.048 \\ 0.016 \\ 0.94 \end{bmatrix} d_k, \quad (63)$$

following discretization of the continuous-time dynamics (see Section G of the supplementary material) with sampling time $\Delta t = 0.1$ s. The state vector $x = [\theta_1 \ \theta_2 \ \omega_1 \ \omega_2]^\top$ consists of the respective angles and angular velocities of the two masses. The input $u = M_M$ is the torque on mass 1, the disturbance $d = M_d$ is the torque on mass 2. The goal is to stabilize the origin while minimizing cumulative stage costs $\sum_k x_k^\top Q x_k + u_k R u_k$ with $Q = \text{diag}(1, 100, 1, 1)$, $R = 1$, and respecting input constraints $|u| \leq 5$ Nm and component-wise state constraints $-\mathbf{x}_{\max} \leq x \leq \mathbf{x}_{\max}$, with $\mathbf{x}_{\max}^\top = [2\pi \ 2\pi \ \pi \cdot 0.5 \ 0.5\pi]^\top$. The disturbance d_k acting on mass 2 is sampled randomly from a zero-mean normal distribution with variance $\bar{d} = (0.2 \text{ Nm})^2$, truncated at the bounds of the interval $[-\bar{d}, \bar{d}]$. Online state measurements are corrupted by measurement noise μ_k , which is sampled from a zero-mean normal distribution with covariance matrix $\bar{\mu}I_4$, $\bar{\mu} = 0.01^2$, truncated such that $\|\mu\|_\infty \leq \bar{\mu}$.

A. EXACT OFFLINE DATA

For the proposed control algorithm, the model (63) is assumed to be unknown. To retrieve trajectory data as in Assumption 1, we simulate the system in open-loop for 50 time-steps, with admissible inputs chosen at random. With the trajectory data, we generate a persistently excited nominal state trajectory via (20a), which lets us compute the LQR feedback gain K based on [3, Theorem 4] and the specified Q, R . By Proposition 1 of the supplementary material, we retrieve the associated matrix P for the definition of the terminal cost function. For the probabilistic constraint tightening, we use 2924 disturbance measurements (Assumption 2 a) to solve the probabilistic optimization problems (15), (31), with confidence $1 - \beta = 0.99$ for a risk parameter $p \in [0.88, 0.92]$. The resulting constraint sets are further tightened to account for the bounded measurement noise (see Section IV.IV-B), for which we assume the bound $\bar{\mu}$ to be known (Assumption 2 b). We employ a Monte-Carlo simulation of 100 runs, each with different realizations of the online measurement noise and additive disturbance, a length of 50 time steps, and fixed initial state $x_0 = [1.5 \ 2 \ 0 \ 0]^\top$.

All simulations are carried out in MATLAB, with polyhedral constraints specified based on MPT3 [43], linear matrix inequalities solved with CVX [44], and the OCP (33) solved by MATLAB's `quadprog`. With a prediction horizon set to $L = 10$, the mean computation time for the solution of




FIGURE 3. Trajectories of 100 simulations for different realizations of online disturbance and measurement noise. Dotted lines are constraints. The average of all trajectories is shown highlighted.

the OCP was 1.650.ms on an Intel I5-13600K. The resulting trajectories are shown in Fig. 3. In all scenarios, the OCP remained feasible for all time steps and the proposed controller stabilized the system around the origin while respecting input constraints. Chance constraints were met, with rare constraint violations in the interval [0.3s, 1.3s]. This conservatism is introduced by the robust treatment of the measurement noise. More frequent constraint violations occur if the bound on the measurement noise is reduced, see the additional simulation examples in Section H of the supplementary material.

B. INEXACT OFFLINE DATA

Although inexact data is not explicitly considered in the control design beyond the regularization techniques presented in Section VI, we demonstrate the applicability of our proposed control algorithm in a more realistic setting. To that end, we perturb all measured states and disturbances (including those used for the probabilistic constraint tightening) by measurement noise μ as defined above, with $\bar{\mu} = 0.01^2$ unchanged. In other words, we now consider the same measurement noise to be present both online and offline, and also similarly perturb disturbance samples, emulating a prior disturbance estimation procedure. In order to investigate the effect of noise inside the offline data, we sample 5 different realizations of offline noise and use it to perturb the same state and disturbance measurements. For each of the 5 sets of data, we compute tightened constraint sets, controller gain K , and cost-to-go P , but employ simple robustifications as discussed in Section VI. In the cost function of all 5 resulting OCPs, we add the term $5\|\Pi\alpha\|_2^2$.

For all 5 resulting data-driven predictive controllers, we again simulate 100 different realizations of online measurement noise and disturbance each. In the case of $\bar{d} = (0.2\text{Nm})^2$, we observed that about one out of five controllers




FIGURE 4. Median cost trajectories for 100 runs based on offline state and disturbance data perturbed by measurement noise, each corresponding to different realizations of offline noise. For comparison, the trajectory with exact offline data is shown dashed.

TABLE 1 Performance comparison of exact and noisy (inexact) state and disturbance data.

\bar{d} in $(\text{Nm})^2$	Data	OCP infeasible	Mean Cost	STD Cost
0.1	Exact	0	2974.9	50.6
0.1	Inexact	0	3757.5	2098.6
0.2	Exact	0	2979.3	117.1
0.2	Inexact	0	5198.9	4603.5

did not steer the state to the origin (see Section I of the supplementary material). We lowered the disturbance level to $\bar{d} = (0.1\text{Nm})^2$ and repeated the simulations. Median cost trajectories for the 5 different resulting predictive controllers are shown in Fig. 4. Whereas the control performance varies significantly, all predictive controllers stabilized the system around the origin. The resulting mean and standard deviation of the total costs can be seen in Table 1, which also includes results for the predictive controller based on exact (unperturbed, noise-free) offline data. For the given parameters, even with noisy data, the OCP was feasible for all time steps. For the setting with lower disturbance levels, mean costs increased by 26.3%. In contrast, median costs increased by only 6.7% from 2749.5 to 2934.3. Overall, we observed that in the present setting, the effect of noisy data is often small, and rarely very large. In hindsight, this further motivates a stochastic, as opposed to robust, future road to data-driven control.

VIII. CONCLUSION

This work presented a novel data-driven stochastic predictive control scheme for the control of constrained LTI systems subject to stochastic disturbances and measurement noise. The goal was to develop a lightweight and efficient predictive control strategy that leverages the deterministic fundamental lemma while also considering the probabilistic information of disturbances for performance-oriented control. The key idea

of the paper was to leverage a data-driven formulation of robust and stochastic tubes that leads to a recursively feasible receding horizon OCP and a predictive controller that renders the system input-to-state stable with respect to both disturbance and measurement noise. On the one hand, we provided a lightweight deterministic receding horizon OCP that allows for the data-driven control of stochastic systems without the need of a model. On the other hand, we provided a data-driven tube-based formulation that allows for the translation of state-of-the-art stochastic and robust tube-based MPC results into the data-driven domain.

REFERENCES

- [1] J. C. Willems, P. Rapisarda, I. Markovskiy, and L. M. B. De Moor, "A note on persistency of excitation," *Syst. Control Lett.*, vol. 54, no. 4, pp. 325–329, 2005.
- [2] I. Markovskiy and P. Rapisarda, "Data-driven simulation and control," *Int. J. Control*, vol. 81, no. 12, pp. 1946–1959, 2008.
- [3] C. De Persis and P. Tesi, "Formulas for data-driven control: Stabilization, optimality, and robustness," *IEEE Trans. Autom. Control*, vol. 65, no. 3, pp. 909–924, Mar. 2020.
- [4] H. J. van Waarde, J. Eising, and H. L. Trentelman, "Data informativity: A new perspective on data-driven analysis and control," *IEEE Trans. Autom. Control*, vol. 65, no. 11, pp. 4753–4768, Nov. 2020.
- [5] J. Berberich and F. Allgöwer, "A trajectory-based framework for data-driven system analysis and control," in *Proc. IEEE Eur. Control Conf.*, 2020, pp. 1365–1370.
- [6] M. Yin, A. Iannelli, and R. S. Smith, "Maximum likelihood estimation in data-driven modeling and control," *IEEE Trans. Autom. Control*, vol. 68, no. 1, pp. 317–328, Jan. 2023.
- [7] I. Markovskiy and F. Dörfler, "Behavioral systems theory in data-driven analysis, signal processing, and control," *Annu. Rev. Control*, vol. 52, pp. 42–64, 2021.
- [8] H. Yang and S. Li, "A data-driven predictive controller design based on reduced Hankel matrix," in *Proc. IEEE 10th Asian Control Conf.*, 2015, pp. 1–7.
- [9] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," in *Proc. IEEE 18th Eur. Control Conf.*, 2019, pp. 307–312.
- [10] 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," *Annu. Rev. Control*, vol. 55, pp. 92–117, 2023.
- [11] L. Huang, J. Lygeros, and F. Dörfler, "Robust and Kernelized data-enabled predictive control for nonlinear systems," 2022, [arXiv:2206.01866](https://arxiv.org/abs/2206.01866).
- [12] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Linear tracking MPC for nonlinear systems – Part II: The data-driven case," *IEEE Trans. Autom. Control*, vol. 67, no. 9, pp. 4406–4421, Sep. 2022.
- [13] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE Trans. Autom. Control*, vol. 66, no. 4, pp. 1702–1717, Apr. 2021.
- [14] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven tracking MPC for changing setpoints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6923–6930, 2020.
- [15] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Robust constraint satisfaction in data-driven MPC," in *Proc. IEEE 59th Conf. Decis. Control*, 2020, pp. 1260–1267.
- [16] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "On the design of terminal ingredients for data-driven MPC," in *IFAC-PapersOnLine*, vol. 54, no. 6, pp. 257–263, 2021.
- [17] J. Bongard, J. Berberich, J. Köhler, and F. Allgöwer, "Robust stability analysis of a simple data-driven model predictive control approach," *IEEE Trans. Autom. Control*, vol. 68, no. 5, pp. 2625–2637, May 2023.
- [18] L. Huang, J. Zhen, J. Lygeros, and F. Dörfler, "Robust data-enabled predictive control: Tractable formulations and performance guarantees," *IEEE Trans. Autom. Control*, vol. 68, no. 5, pp. 3163–3170, May 2023.
- [19] J. Coulson, J. Lygeros, and F. Dörfler, "Regularized and distributionally robust data-enabled predictive control," in *Proc. IEEE 58th Conf. Decis. Control*, 2019, pp. 2696–2701.
- [20] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," *IEEE Trans. Autom. Control*, vol. 67, no. 7, pp. 3289–3304, Jul. 2022.
- [21] Y. Wang, C. Shang, and D. Huang, "Data-driven control of stochastic systems: An innovation estimation approach," 2022, [arXiv:2209.08995](https://arxiv.org/abs/2209.08995).
- [22] G. Pan, R. Ou, and T. Faulwasser, "On a stochastic fundamental lemma and its use for data-driven optimal control," *IEEE Trans. Autom. Control*, early access, Dec. 27, 2022.
- [23] G. Pan, R. Ou, and T. Faulwasser, "Towards data-driven stochastic predictive control," 2022, [arXiv:2212.10663](https://arxiv.org/abs/2212.10663).
- [24] G. Pan, R. Ou, and T. Faulwasser, "Data-driven stochastic output-feedback predictive control: Recursive feasibility through interpolated initial conditions," in *Proc. 5th Annu. Learn. Dyn. Control Conf.*, vol. 211, Jun. 2023, pp. 980–992.
- [25] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Stochastic MPC with offline uncertainty sampling," *Automatica*, vol. 81, pp. 176–183, 2017.
- [26] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Syst. Mag.*, vol. 36, no. 6, pp. 30–44, Dec. 2016.
- [27] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Constraint-tightening and stability in stochastic model predictive control," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3165–3177, Jul. 2017.
- [28] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.
- [29] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design.*, 2nd ed., Madison, Wisconsin, USA: Nob Hill Publishing, 2017.
- [30] B. J. P. Roset, W. P. M. H. Heemels, M. Lazar, and H. Nijmeijer, "On robustness of constrained discrete-time systems to state measurement errors," *Automatica*, vol. 44, no. 4, pp. 1161–1165, 2008.
- [31] S. Küçükyavuz and R. Jiang, "Chance-constrained optimization under limited distributional information: A review of reformulations based on sampling and distributional robustness," *EURO J. Comput. Optim.*, vol. 10, 2022, Art. no. 100030.
- [32] M. C. Campi and S. Garatti, "A sampling-and-discarding approach to chance-constrained optimization: Feasibility and optimality," *J. Optim. Theory Appl.*, vol. 148, no. 2, pp. 257–280, 2011.
- [33] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints—A review," *J. Process Control*, vol. 44, pp. 53–67, 2016.
- [34] 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, Jul. 2020.
- [35] L. Grüne and J. Pannek, "Nonlinear model predictive control," in *Nonlinear Model Predictive Control*. Berlin, Germany: Springer, 2017, pp. 45–69.
- [36] B. Kouvaritakis, M. Cannon, S. V. Rakovic, and Q. Cheng, "Explicit use of probabilistic distributions in linear predictive control," *Automatica*, vol. 46, no. 10, pp. 1719–1724, 2010.
- [37] P.-O. Gutman and M. Cwikel, "An algorithm to find maximal state constraint sets for discrete-time linear dynamical systems with bounded controls and states," *IEEE Trans. Autom. Control*, vol. 32, no. 3, pp. 251–254, Mar. 1987.
- [38] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, vol. 78, 2nd ed. Berlin, Germany: Springer, 2015.
- [39] Z.-P. Jiang and Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," *Automatica*, vol. 37, no. 6, pp. 857–869, 2001.
- [40] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall and London, Pearson Education, 2002.
- [41] F. Dörfler, J. Coulson, and I. Markovskiy, "Bridging direct and indirect data-driven control formulations via regularizations and relaxations," *IEEE Trans. Autom. Control*, vol. 68, no. 2, pp. 883–897, Feb. 2023.
- [42] F. Dörfler, P. Tesi, and C. de Persis, "On the role of regularization in direct data-driven LQR control," in *Proc. IEEE 61st Conf. Decis. Control*, 2022, pp. 1091–1098.
- [43] M. Herceg, M. Kvasnica, C. N. Jones, and M. Morari, "Multi-parametric toolbox 3.0," in *Proc. Eur. Control Conf.*, Zürich, Switzerland, 2013, pp. 502–510. [Online]. Available: <http://control.ee.ethz.ch/mpt>

[44] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming,” version 2.1, Mar. 2014. [Online]. Available: <http://cvxr.com/cvx>



SEBASTIAN KERZ (Graduate Student Member, IEEE) received the M.Sc. degree in electrical engineering and information technology from Technical University Darmstadt, Darmstadt, Germany, in 2019. During his studies, he was a Visiting Student with the University of Toronto, Toronto, ON, Canada, in 2017. He is currently a Research Associate and Ph.D. student with the Chair of Automatic Control Engineering, the Technical University of Munich, Munich, Germany. His research interests include data-driven control, optimal control, and the intersection between control theory and machine learning.



JOHANNES TEUTSCH received the B.Sc. and M.S. degrees in electrical engineering and information technology from the Technical University of Munich, Munich, (TUM), Germany, in 2019 and 2021, respectively. He is currently a Research Associate and Ph.D. student with the Chair of Automatic Control Engineering, TUM. His research interests include (model) predictive control, with a focus on data-driven predictive control in uncertain environments.



TIM BRÜDIGAM (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Technical University of Munich (TUM), Munich, Germany, in 2014 and 2017, respectively. In 2022, he received the Ph.D. degree from TUM. During his studies, he was a Research Scholar with the California Institute of Technology, Pasadena, CA, USA, in 2016. He joined Chair of Automatic Control Engineering, TUM, as a Research Associate, in 2017, and was a Visiting Researcher with the University of California, Berkeley, CA, in 2021. His research interests include advancing Model Predictive Control (MPC), especially stochastic MPC, with possible application in automated driving.



MARION (NEE SOBOTKA) LEIBOLD received the Diploma in applied mathematics from Technical University of Munich, Munich, Germany, in 2002. She received the Ph.D. and Habilitation degrees in control theory from Technical University of Munich, in 2007 and 2020, respectively. She is currently a Senior Researcher with the Faculty of Electrical Engineering and Information Technology, Institute of Automatic Control Engineering, Technical University of Munich. Her research interests include optimal control and non-linear control theory and the applications to robotics and automated driving.



DIRK WOLLHERR (Senior Member, IEEE) received the Dipl.-Ing., Dr.- Ing., and Habilitation degrees in electrical engineering from Technical University of Munich, Germany, in 2000, 2005, and, 2013, respectively. He is a Senior Researcher in robotics and control and cognitive systems. His research interests include automatic control, robotics autonomous mobile robots, human-robot interaction, and socially aware collaboration and joint action. He is with the management boards of several conferences and Associate Editor for IEEE TRANSACTIONS ON MECHATRONICS.