

Conformal Prediction for Distribution-free Optimal Control of Linear Stochastic Systems

Eleftherios E. Vlahakis, Lars Lindemann, *Member, IEEE*, Pantelis Sopasakis, and Dimos V. Dimarogonas, *Fellow, IEEE*

Abstract—We address an optimal control problem for linear stochastic systems with unknown noise distributions and joint chance constraints using conformal prediction. Our approach involves designing a feedback controller to maintain an error system within a prediction region (PR). We define PRs as sublevel sets of a nonconformity score over error trajectories, enabling the handling of joint chance constraints. We propose two methods to design feedback control and PRs: one through direct optimization over error trajectory samples, and the other indirectly using the S -procedure with a disturbance ellipsoid obtained from data. By tightening constraints with PRs, we solve a relaxed problem to synthesize a feedback policy. Our method ensures reliable probabilistic guarantees based on marginal coverage, independent of data size.

Index Terms—conformal prediction, stochastic systems

I. INTRODUCTION

PROBABILISTIC guarantees play a crucial role in many applications that involve stochastic disturbances and safety. As chance-constrained problems are generally non-convex and computationally intractable, most approaches solve deterministic relaxations by applying constraint-tightening techniques based on available information about uncertainty [1]. Constraint tightening for probabilistic satisfaction has been studied when the underlying probability distribution is known [2], in cases of bounded disturbances [3], or using probabilistic reachable sets [4]–[6]. These approaches may be restricted to Gaussian settings, involve computationally costly operations, or often rely on the multivariate Chebyshev inequality, which can lead to conservative bounds. In contrast, sampling-based methods such as [7]–[9] can be significantly more flexible in relaxing chance-constrained problems by utilizing available disturbance samples, termed scenarios, and can provide formal guarantees based on scenario optimization (SO) [10], [11].

This work was supported by the Swedish Research Council (VR), the Knut & Alice Wallenberg Foundation (KAW), the Horizon Europe Grant SymAware and the ERC Consolidator Grant LEAFHOUND.

E.E. Vlahakis and D.V. Dimarogonas are with the Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, 10044, Stockholm, Sweden. Email: {vlahakis, dimos}@kth.se. L. Lindemann is with the Thomas Lord Department of Computer Science and the Ming Hsieh Department of Electrical and Computer Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, 90089, CA, USA. Email: llindema@usc.edu. P. Sopasakis is with the School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Northern Ireland, UK. Email: p.sopasakis@qub.ac.uk

An alternative distribution-free framework for streamlined uncertainty quantification, equipped with formal guarantees, is provided by conformal prediction (CP). Originally introduced by Vovk and Shafer [12], [13], CP uses a calibration dataset to infer prediction regions for a test dataset with a specified probability, while remaining distribution-agnostic. CP has been applied in machine learning [14], chance-constrained optimization [15], and control for safety verification and planning [16]–[18], see [19] for a recent survey.

Leveraging its distribution independence, we apply CP to tackle an optimal control problem for linear stochastic systems under joint chance constraints, assuming data-driven noise information. Our approach is closely related to [8], [9], which employ offline sampling for constraint tightening and probabilistic reachable set computation. Both methods assume a fixed feedback control policy, which preserves convexity in the underlying scenario-based relaxed programs, as required by SO in [10], [11] to enable solutions with a priori guarantees. Disturbance feedback parameterization [20] is employed in [7] for online SO-based stochastic optimal control, which maintains convexity and is well-suited for short control horizons. In our approach, we design a feedback controller using disturbance samples to ensure the error system remains within a prediction region (PR). Although state-feedback control introduces a nonlinear dependence of dynamics on the feedback gain, we can derive formal guarantees through a two-step training-calibration procedure provided by CP. We define PRs as sublevel sets of nonconformity scores, with thresholds based on empirical quantiles computed from data. To handle joint chance constraints, we introduce nonconformity scores capturing the maximum error across a trajectory. We propose two methods. In the first, we use a training dataset to optimize nonconformity score quantiles for controller design, followed by PR construction via CP on a calibration dataset. In the second method, we first obtain a disturbance PR via CP, and subsequently use the S -procedure to design a PR and a feedback gain for the error system. The PRs are used to tighten the constraints, enabling us to solve a deterministic relaxation problem that yields a feedback control policy with probabilistic guarantees. We demonstrate the performance of both methods through a numerical example and compare our approach with the SO-based randomized method in [7].

II. PRELIMINARIES AND PROBLEM SETUP

Notation: The set $\mathbb{N}_{[N_1, N_2]}$ collects all integers in $[N_1, N_2]$. $x(a : b) = (x(a), \dots, x(b))$ is an aggregate vector collecting

$x(t)$, $t \in \mathbb{N}_{[a,b]}$, representing a trajectory or a random process. The ceiling operator is $\lceil \cdot \rceil$. The probability of Y is $\Pr\{Y\}$. Quantile(\mathcal{D}, δ) is the δ -th quantile of a distribution \mathcal{D} , i.e., for $Z \sim \mathcal{D}$, $\text{Quantile}(\mathcal{D}, \delta) = \inf\{z : \Pr\{Z \leq z\} \geq \delta\}$. A being a (proper) subset of B is denoted as $(A \subsetneq B) A \subseteq B$. The Pontryagin set difference of $S_1, S_2 \subseteq \mathbb{R}^n$ is $S_1 \ominus S_2 = \{s_1 \in \mathbb{R}^n \mid s_1 + s_2 \in S_1, \forall s_2 \in S_2\}$. $\text{int}(X) = \{x \in X \mid \exists e > 0, \mathbb{B}(e) + x \in X\}$, where $\mathbb{B}(e) = \{x \mid \sqrt{x^\top x} \leq e\}$. The i th, the smallest, and the largest eigenvalues of a symmetric matrix M are $\lambda_i(M)$, $\lambda_{\min}(M)$, and $\lambda_{\max}(M)$, respectively. The determinant and trace of M is $\det M$ and $\text{trace } M$.

Conformal Prediction: Let $\mathcal{R}^{(0)}, \dots, \mathcal{R}^{(k)}$ be independent and identically distributed (i.i.d.) random variables. We will refer to $\mathcal{R}^{(j)}$ as a nonconformity score. Given a failure probability $\theta \in (0, 1)$, one wishes to obtain a bound $C \in \mathbb{R}$ for $\mathcal{R}^{(0)}$ so that

$$\Pr\{\mathcal{R}^{(0)} \leq C\} \geq 1 - \theta, \quad (1)$$

where C is based on the samples $\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}$, which we call calibration dataset. Specifically, C may be attained as $C = \text{Quantile}(\{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}, \infty\}, 1 - \theta)$, which is the $(1 - \theta)$ th quantile of the empirical distribution $\{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}, \infty\}$. Assuming $\mathcal{R}^{(1)} \leq \dots \leq \mathcal{R}^{(k)}$, one can pick $C = \mathcal{R}^{(p)}$, where $p = \lceil (k+1)(1 - \theta) \rceil$, which indicates the p th smallest nonconformity score. Note that C is finite with $p \in \mathbb{N}_{[1,k]}$ if $k \geq \lceil (k+1)(1 - \theta) \rceil$. This choice of C ensures that (1) holds since test point $\mathcal{R}^{(0)}$ and calibration data $\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}$ are i.i.d. [21]. This is summarized below.

Lemma 1 [21, Lemma 1] *If $\mathcal{R}^{(0)}, \dots, \mathcal{R}^{(k)}$ are i.i.d. random variables, then for any $\theta \in (0, 1)$, we have*

$$\Pr\{\mathcal{R}^{(0)} \leq \text{Quantile}(\{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}, \infty\}, 1 - \theta)\} \geq 1 - \theta. \quad (2)$$

Remark 1 *The coverage guarantees in (2) are marginal as the probability is defined over the randomness in the draw of test and calibration points $\mathcal{R}^{(0)}, \mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}$. Conditional coverage guarantees of the form $\Pr\{\mathcal{R}^{(0)} \leq C \mid \mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}\}$ are unfortunately not possible to obtain. However, one can show that the conditional probability is a random variable following a beta distribution centered at $1 - \theta$ regardless of k [14], [19]. Notably, probably approximately correct coverage guarantees $\Pr\{\Pr\{\mathcal{R}^{(0)} \leq \text{Quantile}(\{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}, \infty\}, 1 - \hat{\theta})\} \geq 1 - \theta\} \geq 1 - \beta$ can be obtained by setting $\hat{\theta} = \theta - \sqrt{\frac{\ln(1/\beta)}{2k}}$, with the “outer” probability taken with respect to the randomness in the draw of the calibration data $\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(k)}$, and $\beta \in (0, 1)$ [22].*

Problem Statement: We wish to solve the chance-constrained optimal control problem

$$\underset{\mathbf{u}(0:N-1), \mathbf{x}(1:N)}{\text{Minimize}} \quad \mathbb{E} \left[\sum_{t=0}^{N-1} \ell(x(t), u(t)) + V_f(x(N)) \right], \quad (3a)$$

$$\text{s.t. } x(t+1) = Ax(t) + Bu(t) + w(t), \quad t \in \mathbb{N}_{[0,N]}, \quad (3b)$$

$$\Pr\{x(t) \in \mathcal{X}_t, \forall t \in \mathbb{N}_{[1,N]}\} \geq 1 - \theta, \quad (3c)$$

$$\Pr\{u(t) \in \mathcal{U}, \forall t \in \mathbb{N}_{[0,N-1]}\} \geq 1 - \theta, \quad (3d)$$

$$x(0) = x_0, \quad (3e)$$

where $\ell : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ and $V_f : \mathbb{R}^n \rightarrow \mathbb{R}$ are cost functions penalizing the state and input trajectories $x(0 : N-1)$, $\mathbf{u}(0 : N-1)$, and the end point $x(N)$, respectively, of the discrete-time linear stochastic system with dynamics in (3b), and $\theta \in (0, 1)$ is a failure probability. In (3b), $x(t) \in \mathbb{R}^n$ is a state vector, $u(t) \in \mathbb{R}^m$ is an input vector, and $w(t)$ is a disturbance vector consisting of random variables drawn from arbitrary distributions. We assume that $w(t)$, $t \in \mathbb{N}_{[0,N-1]}$, are i.i.d., that is, $w(t) \sim \mathcal{D}$ for all $t \in \mathbb{N}_{[0,N-1]}$, with the underlying multivariate distribution \mathcal{D} being unknown. We assume that (A, B) is stabilizable with $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and the state and input constraint sets are ellipsoidal regions of \mathbb{R}^n and \mathbb{R}^m , respectively, defined as

$$\mathcal{X}_t = \{x \in \mathbb{R}^n \mid (x - p_t)^\top P_t (x - p_t) \leq 1\}, \quad (4a)$$

$$\mathcal{U} = \{u \in \mathbb{R}^m \mid u^\top Q u \leq 1\}, \quad (4b)$$

where $p_t \in \mathbb{R}^n$, $t \in \mathbb{N}_{[1,N]}$, and $P_t \in \mathbb{R}^{n \times n}$, $t \in \mathbb{N}_{[1,N]}$, and $Q \in \mathbb{R}^{m \times m}$ are symmetric positive definite matrices.

Solving (3) is challenging primarily due to the unknown distribution of $w(t)$. In the following, we assume that a disturbance dataset, $\mathcal{D}^w = \{\mathbf{w}^{(0)}, \dots, \mathbf{w}^{(k)}\}$, of $k+1$ samples of disturbance sequences is available, with $\mathbf{w}^{(j)} = (w^{(j)}(0), \dots, w^{(j)}(N-1))$, $j \in \mathbb{N}_{[0,k]}$, and $w^{(j)}(t) \sim \mathcal{D}$ for all $t \in \mathbb{N}_{[0,N-1]}$.

III. MAIN RESULTS

Due to linearity in (3b), the state $x(t)$ can be decomposed into a deterministic part, $z(t)$, and an error, $e(t)$, i.e., $x(t) = z(t) + e(t)$. We will tackle the stochastic control problem in (3)-(4) by the control law $u(t) = Ke(t) + v(t)$, where $K \in \mathbb{R}^{m \times n}$ is a state-feedback gain for the pair (A, B) , and $v(t)$ is a feedforward control action. Then, we may write

$$z(t+1) = Az(t) + Bv(t), \quad (5a)$$

$$e(t+1) = \bar{A}e(t) + w(t), \quad (5b)$$

where $z(0) = x(0)$, $e(0) = 0$, and $\bar{A} = A + BK$. This standard decomposition technique allows the two systems in (5) to be analyzed independently [3]. Unlike existing approaches that assume a fixed feedback controller [4], [9], we aim to design both feedback and feedforward control terms using the available disturbance dataset \mathcal{D}^w . We partition \mathcal{D}^w into training and calibration datasets, $\mathcal{D}_{\text{train}}^w = \{\mathbf{w}^{(k_1+1)}, \dots, \mathbf{w}^{(k)}\}$ and $\mathcal{D}_{\text{cal}}^w = \{\mathbf{w}^{(0)}, \dots, \mathbf{w}^{(k_1)}\}$, where $k_1 + 1 < k$. The training dataset is used to compute design parameters, such as K , while the calibration dataset is used to generate empirical distributions of nonconformity scores defined over error and input trajectories. Using Lemma 1, we then derive prediction regions for the error state $e(t)$ and the feedback term $Ke(t)$, with formal guarantees. Next, we define prediction regions for random vectors, and subsequently, show how to solve a relaxation of the problem in (3) optimizing over control actions, given a fixed feedback gain K .

A. Open-loop control

Definition 1 *Let $\xi \in \mathbb{R}^n$ be a random vector. We call $\mathcal{E}_{1-\theta}(\xi) \subseteq \mathbb{R}^n$ a prediction region (PR) for ξ at a confidence level $1 - \theta$, if $\Pr\{\xi \in \mathcal{E}_{1-\theta}(\xi)\} \geq 1 - \theta$. We also*

denote $\mathcal{E}_{1-\theta}^{a:b}(\xi(t))$ a prediction region for the random process $(\xi(a), \dots, \xi(b))$, if $\Pr\{\xi(t) \in \mathcal{E}_{1-\theta}^{a:b}(\xi(t)) \forall t \in \mathbb{N}_{[a,b]}\} \geq 1 - \theta$.

Given PRs $\mathcal{E}_{1-\theta}^{1:N}(e(t))$, $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$, we can tighten the constraints in (4) to produce constraints for the deterministic variables $z(t)$, $v(t)$, in (5a). Conditions for tightening of the constraints in (4) are given next.

Lemma 2 Consider the ellipsoidal constraints in (4), and the error system in (5b). Let $\mathcal{E}_{1-\theta}^{1:N}(e(t))$ and $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$ be prediction regions for $(e(1), \dots, e(N))$ and $(Ke(0), \dots, Ke(N-1))$, respectively. If $\mathcal{E}_{1-\theta}^{1:N}(e(t)) \subseteq \mathbb{B}(C_e)$ and $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t)) \subseteq \mathbb{B}(C_u)$, with $C_e \in (0, 1/\sqrt{\lambda_{\max}(P_t)})$, $\forall t \in \mathbb{N}_{[1,N]}$, and $C_u \in (0, 1/\sqrt{\lambda_{\max}(Q)})$, respectively, then

$$\text{int}(\mathcal{X}_t \ominus \mathcal{E}_{1-\theta}^{1:N}(e(t))) \neq \emptyset \forall t \in \mathbb{N}_{[1,N]}, \quad (6a)$$

$$\text{int}(\mathcal{U} \ominus \mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))) \neq \emptyset. \quad (6b)$$

Proof: Define $c_{\min} = \min_t 1/\sqrt{\lambda_{\max}(P_t)}$. For the condition in (6a), it suffices to show that $\text{int}(\hat{\mathcal{X}}_t \ominus \mathbb{B}(C_e)) \neq \emptyset$, $\forall t \in \mathbb{N}_{[1,N]}$, where $\hat{\mathcal{X}}_t = \mathcal{X}_t - p_t$. Note that $\hat{\mathcal{X}}_t \ni 0$ is an ellipsoid, with the length of its smallest semi-axis equal to $1/\sqrt{\lambda_{\max}(P_t)}$. Thus, $\mathbb{B}(c_{\min}) \subseteq \hat{\mathcal{X}}_t$, $\forall t \in \mathbb{N}_{[1,N]}$. Moreover, since $C_e \in (0, c_{\min})$, we have $\mathbb{B}(c_{\min}) \ominus \mathbb{B}(C_e) = \mathbb{B}(c_{\min} - C_e) \subseteq \hat{\mathcal{X}}_t \ominus \mathbb{B}(C_e)$, $\forall t \in \mathbb{N}_{[1,N]}$. Taking the interior on both sides of this inclusion, the result follows, i.e., $\text{int}(\mathbb{B}(c_{\min} - C_e)) \subseteq \text{int}(\hat{\mathcal{X}}_t \ominus \mathbb{B}(C_e)) \neq \emptyset$, since $\text{int}(\mathbb{B}(c_{\min} - C_e)) \neq \emptyset$. The condition in (6b) is proved similarly. ■

We now present a relaxation of the problem in (3) with tighter state and input constraints based on PRs for a fixed K similar to the approaches in [3], [4], [8].

Theorem 1 Consider the optimization problem in (3)-(4), the system in (5) with PRs, $\mathcal{E}_{1-\theta}^{1:N}(e(t)) \subsetneq \mathbb{R}^n$, $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t)) \subsetneq \mathbb{R}^m$, for a fixed gain K , and

$$\text{Minimize}_{\mathbf{v}(0:N-1), \mathbf{z}(1:N)} \sum_{t=0}^{N-1} \ell(z(t), v(t)) + V_f(z(N)) \quad (7a)$$

$$\text{s.t. } z(t+1) = Az(t) + Bv(t), \quad t \in \mathbb{N}_{[0,N]}, \quad (7b)$$

$$z(t) \in \mathcal{Z}_t, \quad \forall t \in \mathbb{N}_{[1,N]}, \quad (7c)$$

$$v(t) \in \mathcal{V}, \quad \forall t \in \mathbb{N}_{[0,N-1]}, \quad (7d)$$

$$z(0) = x_0, \quad (7e)$$

where $\mathcal{Z}_t = \mathcal{X}_t \ominus \mathcal{E}_{1-\theta}^{1:N}(e(t))$, $t \in \mathbb{N}_{[1,N]}$, and $\mathcal{V} = \mathcal{U} \ominus \mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$. Assume that the conditions in (6) are true, and that (7) is feasible for $z(0) = x_0$, with an optimal solution $(\mathbf{v}^*(0:N-1), \mathbf{z}^*(1:N))$. Let $\mathbf{u}(0:N-1) = (u(0), \dots, u(N-1))$, with $u(t) = Ke^*(t) + v^*(t)$, where $e^*(t) = x(t) - z^*(t)$, $t \in \mathbb{N}_{[0,N-1]}$, and $x(t)$ is the state of the system (3b). Let the trajectory of (3b), $\mathbf{x}(1:N)$, originated at x_0 be driven by input and disturbance sequences, $\mathbf{u}(0:N-1)$, $\mathbf{w}(0:N-1)$, where individual disturbances follow the same distribution as the elements of $\mathcal{D}_{\text{cal}}^w$. Then, $(\mathbf{u}(0:N-1), \mathbf{x}(1:N))$ is a feasible solution to (3).

Proof: Since the conditions in (6) are true, it suffices to show that the probabilistic constraints in (3c)-(3d) are feasible.

Define events $X := x(t) \in \mathcal{X}_t \forall t \in \mathbb{N}_{[1,N]}$, $E := e(t) \in \mathcal{E}_{1-\theta}^{1:N}(e(t)) \forall t \in \mathbb{N}_{[1,N]}$, and $Z := z(t) \in \mathcal{Z}_t \forall t \in \mathbb{N}_{[1,N]}$, and let \hat{E} be the complement of E . Since (7) is feasible by assumption, $\Pr\{Z\} = 1$, implying that $\Pr\{X | E\} = 1$, by definition of $\mathcal{Z}_t = \mathcal{X}_t \ominus \mathcal{E}_{1-\theta}^{1:N}(e(t))$. From the law of total probability, we have $\Pr\{X\} = \Pr\{X | E\} \Pr\{E\} + \Pr\{X | \hat{E}\} \Pr\{\hat{E}\} \geq 1 - \theta$, which follows from the fact that $\Pr\{X | E\} = 1$, $\Pr\{E\} \geq 1 - \theta$, and $\Pr\{X | \hat{E}\} \Pr\{\hat{E}\} \geq 0$. The constraint in (3d) can be shown by similar reasoning. ■

We note that the conditions in (6) of Lemma 2 ensure that (7) is well defined. Theorem 1 can be stated independently. We also remark that the input $u(t) = Ke(t) + v(t)$ may be unbounded as $e(t)$ is influenced by a stochastic, potentially unbounded, disturbance. The chance constraint in (3d) allows the input constraint to be violated for some disturbance realizations, with the parameter θ specifying the extent of such violations. Next, we design PRs that bound the feedback term $Ke(t)$ with a probability of at least $1 - \theta$.

B. Prediction regions and feedback control

Assuming that PRs for the error system in (5b) with a fixed state-feedback gain K are available, Theorem 1 provides a method to construct a feasible solution for the problem in (3)-(4) by solving a tractable optimization problem. Here, we outline two methods for designing K and approximate PRs for the error system. The first method formulates an optimization problem over a training dataset of error trajectories, where each trajectory sample depends nonlinearly on K . After solving this, conformal prediction is applied on a calibration dataset to compute PRs, referred to as the *direct method*, as CP is applied directly to error trajectories. The second method uses conformal prediction on disturbance samples to identify an ellipsoidal set that bounds disturbance sequences with the specified probability. The *S-procedure* is then used to compute an admissible state-feedback gain and an ellipsoidal PR. This is called the *indirect method*, as CP is applied to disturbance samples to indirectly infer coverage for the error system's PRs.

1) Direct method: From the available disturbance dataset \mathcal{D}^w , we define a trajectory dataset as

$$\mathcal{D}^e = \{e^{(0)}, \dots, e^{(k)}\}, \quad (8a)$$

$$e^{(j)} = (e^{(j)}(1), \dots, e^{(j)}(N)), \quad j \in \mathbb{N}_{[0,k]} \quad (8b)$$

$$e^{(j)}(t) = \sum_{i=1}^t (A + BK)^{i-1} w^{(j)}(t-i), \quad t \in \mathbb{N}_{[1,N]}. \quad (8c)$$

In the following, we partition \mathcal{D}^e into $\mathcal{D}_{\text{train}}^e$ and $\mathcal{D}_{\text{cal}}^e$, where trajectory samples in $\mathcal{D}_{\text{train}}^e$ are constructed by disturbance samples in $\mathcal{D}_{\text{train}}^w$ and are parameterized by the feedback gain K , while trajectory samples in $\mathcal{D}_{\text{cal}}^e$ are constructed by disturbance samples in $\mathcal{D}_{\text{cal}}^w$ for a fixed K . We also introduce the nonconformity scores

$$\mathcal{R}_e^{(j)} = \max(\|e^{(j)}(1)\|, \dots, \|e^{(j)}(N)\|), \quad (9a)$$

$$\mathcal{R}_u^{(j)} = \max(\|Ke^{(j)}(0)\|, \dots, \|Ke^{(j)}(N-1)\|), \quad (9b)$$

which will help us identify PRs for the random processes $(e(1), \dots, e(N))$ and $(Ke(0), \dots, Ke(N-1))$, as Euclidean balls bounding $\|e(t)\|$ and $\|Ke(t)\|$ uniformly for all $t \in$

$\mathbb{N}_{[0,N]}$, enabling efficient tightening of the ellipsoids in (4). Consider now the optimization problem

$$\text{Minimize}_{K, \eta_e, \eta_u} \eta_e + \gamma \eta_u \quad (10a)$$

$$\text{s.t. } \Pr\{\mathcal{R}_e^{(0)} \leq \eta_e\} \geq 1 - \theta, \quad (10b)$$

$$\Pr\{\mathcal{R}_u^{(0)} \leq \eta_u\} \geq 1 - \theta, \quad (10c)$$

where the objective is to obtain a state-feedback gain K and minimal PRs $\mathcal{E}_{1-\theta}^{1:N}(e(t))$, $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$ induced by the constraints in (10b)-(10c), and the nonconformity scores $\mathcal{R}_e^{(j)}$, $\mathcal{R}_u^{(j)}$ in (9). The objective in (10a) balances the trade-off between minimizing the error cost η_e and the feedback control cost η_u , as indicated by the nonconformity scores $\mathcal{R}_e^{(0)}$ and $\mathcal{R}_u^{(0)}$ in (9). A specific value for the tuning parameter γ , which leads to a cost function that supports the desired performance while adhering to state and input constraints, is recommended later. At first glance, one might consider using Lemma 1 to replace the probabilistic constraints in (10b)-(10c) with constraints as in (2), that is, $\text{Quantile}(\{\mathcal{R}_e^{(1)}, \dots, \mathcal{R}_e^{(k_1)}, \infty\}, 1 - \theta) \leq \eta_e$, and $\text{Quantile}(\{\mathcal{R}_u^{(1)}, \dots, \mathcal{R}_u^{(k_1)}, \infty\}, 1 - \theta) \leq \eta_u$, by constructing empirical nonconformity scores $\mathcal{R}_e^{(j)}$ and $\mathcal{R}_u^{(j)}$ using the available calibration dataset $\mathcal{D}_{\text{cal}}^e$. However, this is not straightforward because the random variables $\mathcal{R}_e^{(0)}, \dots, \mathcal{R}_e^{(k_1)}$ and $\mathcal{R}_u^{(0)}, \dots, \mathcal{R}_u^{(k_1)}$ do not retain the i.i.d property, as they depend on the decision variable K . Instead, we can first obtain a state-feedback gain K using a separate training dataset $\mathcal{D}_{\text{train}}^e$ and subsequently apply conformal prediction to attain coverage guarantees for the constraints in (10b)-(10c) by computing a calibration dataset $\mathcal{D}_{\text{cal}}^e$ for the obtained gain K . This approach is shown next.

Let $\mathcal{D}_{\text{train}}^e = \{e^{(k_1+1)}, \dots, e^{(k)}\}$, where $e^{(j)}$ is a trajectory sample parameterized over state-feedback gains $K \in \mathbb{R}^{m \times n}$ and driven by a training disturbance sequence $w^{(j)} \in \mathcal{D}_{\text{train}}^w$, with $j \in \mathbb{N}_{[k_1+1, k]}$. Then, we formulate the problem of finding a state-feedback gain K by solving

$$\text{Minimize}_{K, \eta_e < \eta_e^{\max}, \eta_u < \eta_u^{\max}} \eta_e + \gamma \eta_u \quad (11a)$$

$$\text{s.t. } \text{Quantile}(\{\mathcal{R}_e^{(k_1+1)}(K), \dots, \mathcal{R}_e^{(k)}(K)\}, \hat{\theta}) \leq \eta_e, \quad (11b)$$

$$\text{Quantile}(\{\mathcal{R}_u^{(k_1+1)}(K), \dots, \mathcal{R}_u^{(k)}(K)\}, \hat{\theta}) \leq \eta_u, \quad (11c)$$

where $\eta_e^{\max} = \min_{t \in \mathbb{N}_{[1, N]}} (1/\sqrt{\lambda_{\max}(P_t)})$, $\eta_u^{\max} = 1/\sqrt{\lambda_{\max}(Q)}$, $\mathcal{R}_e^{(j)}(K)$ ($\mathcal{R}_u^{(j)}(K)$) denotes the parameterization over the feedback gain K , and we choose $\gamma = \frac{\eta_e^{\max}}{\eta_u^{\max}}$, and $\hat{\theta} = (1 + \frac{1}{k-k_1-1})(1 - \theta)$. Note that ∞ is omitted from the empirical distributions in (11b)-(11c) and $1 - \theta$ is replaced by $(1 + \frac{1}{k-k_1-1})(1 - \theta)$ (see [19, Sec. 2.1] for details). Solving (11) is indeed difficult due to the nonlinear dependence of the constraints in (11b)-(11c) on the decision variable K . One can employ a nonlinear solver to obtain a feasible solution to (11). See, e.g., Sec. IV, where we employ a genetic algorithm to solve (11) for an academic example.

Next, given a state-feedback gain K obtained from (11), we can provide PRs, $\mathcal{E}_{1-\theta}^{1:N}(e(t))$, $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$, via conformal prediction using empirical distributions of the nonconformity scores in (9) generated by constructing the calibration dataset $\mathcal{D}_{\text{cal}}^e$. This is formally stated next.

Lemma 3 Construct the calibration trajectory dataset $\mathcal{D}_{\text{cal}}^e$ as in (8), using the calibration disturbance set $\mathcal{D}_{\text{cal}}^w$ and the solution K from (11). Let $\{\mathcal{R}_e^{(0)}, \dots, \mathcal{R}_e^{(k_1)}\}$, $\{\mathcal{R}_u^{(0)}, \dots, \mathcal{R}_u^{(k_1)}\}$ be the empirical distributions in (9a), (9b), respectively, computed by the samples in $\mathcal{D}_{\text{cal}}^e$ and the gain K . Compute $C_e = \text{Quantile}(\{\mathcal{R}_e^{(1)}, \dots, \mathcal{R}_e^{(k_1)}, \infty\}, 1 - \theta)$ and $C_{Ke} = \text{Quantile}(\{\mathcal{R}_u^{(1)}, \dots, \mathcal{R}_u^{(k_1)}, \infty\}, 1 - \theta)$. Then, if $\mathcal{E}_{1-\theta}^{1:N}(e(t)) := \mathbb{B}(C_e)$, $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t)) := \mathbb{B}(C_{Ke})$, we have

$$\Pr\{e^{(0)}(t) \in \mathcal{E}_{1-\theta}^{1:N}(e(t)), t \in \mathbb{N}_{[1, N]}\} \geq 1 - \theta, \quad (12a)$$

$$\Pr\{Ke^{(0)}(t) \in \mathcal{E}_{1-\theta}^{0:N-1}(Ke(t)), t \in \mathbb{N}_{[0, N]}\} \geq 1 - \theta. \quad (12b)$$

Proof: First we prove that $\mathcal{D}_{\text{cal}}^e$ consists of $k_1 + 1$ i.i.d. random trajectories of (5b). For each fixed $t \in \mathbb{N}_{[1, N]}$, the random vector $e^{(j)}(t)$ in (8c) is identically distributed across all $j \in \mathbb{N}_{[0, k_1]}$. This means that for any fixed time index t , the distribution of $e^{(j)}(t)$ does not depend on the index j for a given K . Therefore, since the distribution of each individual $e^{(j)}(t)$ is the same for all j , it follows that the entire random process $e^{(j)}(1 : N)$ in (8b) is identically distributed across $j \in \mathbb{N}_{[0, k_1]}$. By the independence of $w^{(j)}$ across all $j \in \mathbb{N}_{[0, k_1]}$, $e^{(j)}(1 : N)$, $j \in \mathbb{N}_{[0, k_1]}$, are i.i.d. random processes, implying that $\mathcal{R}_e^{(0)}, \dots, \mathcal{R}_e^{(k_1)}$ are $k_1 + 1$ i.i.d random variables. Thus, by Lemma 1, we obtain the coverage $\Pr\{\mathcal{R}_e^{(0)} \leq C_e\} \geq 1 - \theta$, which can be written as $\Pr\{\max(\|e^{(0)}(1)\|, \dots, \|e^{(0)}(N)\|) \leq C_e\} \geq 1 - \theta$ by (9a), or as $\Pr\{e^{(0)}(1) \in \mathbb{B}(C_e) \wedge \dots \wedge e^{(0)}(N) \in \mathbb{B}(C_e)\} \geq 1 - \theta$, leading to (12a). The proof of (12b) follows identical lines. ■

We summarize important remarks for selecting the nonconformity scores in (9) as follows: 1) The PRs obtained in Lemma 3 give marginal guarantees, that is, the probabilities in (12) are averaged over the randomness of both the training and calibration data. 2) The ball $\mathbb{B}(\|K\|C_e)$ is a conservative PR, i.e., $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$, motivating the use of two nonconformity scores. 3) We can select ∞ -norm-based nonconformity scores to infer box-shaped PRs when more appropriate, e.g., for polyhedral constraints.

2) Indirect method: First, we compute an ellipsoid containing all the training disturbance samples in $\mathcal{D}_{\text{train}}^w$ by solving the optimization

$$\text{Minimize}_{\hat{Y} \succ 0} \log \det \hat{Y}^{-1} \quad (13a)$$

$$\text{s.t. } \|\hat{Y} w^{(j)}(t)\| \leq 1, t \in \mathbb{N}_{[0, N-1]}, j \in \mathbb{N}_{[k_1+1, k]}. \quad (13b)$$

Note that (13) is a convex problem [23, Sec. 8.4], with \hat{Y} being a symmetric positive definite matrix. Also, the volume of $\mathcal{W} = \{w \mid \|\hat{Y} w\| \leq 1\}$ is minimum since it is proportional to $\det \hat{Y}$ by $\hat{Y} \succ 0$. Next, we obtain an ellipsoidal prediction region $\mathcal{E}_{1-\theta}^{0:N-1}(w(t))$. Consider the nonconformity score

$$\mathcal{R}_w^{(j)} = \max(\|\hat{Y} w^{(j)}(0)\|, \dots, \|\hat{Y} w^{(j)}(N-1)\|), \quad (14)$$

and compute

$$C_w = \text{Quantile}(\{\mathcal{R}_w^{(1)}, \dots, \mathcal{R}_w^{(k_1)}, \infty\}, 1 - \theta), \quad (15)$$

over the calibration dataset $\mathcal{D}_{\text{cal}}^w$. Then, by Lemma 1, we obtain $\Pr\{\mathcal{R}_w^{(0)} \leq C_w\} \geq 1 - \theta$, which implies that

$\Pr\{\|\hat{Y}w^{(0)}(t)\| \leq C_w \ \forall t \in \mathbb{N}_{[0, N-1]}\} \geq 1 - \theta$, i.e., the ellipsoid

$$\mathcal{W} = \{w \mid w^\top Y w \leq 1\}, \quad (16)$$

is a PR, $\mathcal{E}_{1-\theta}^{0:N-1}(w(t))$, where $Y = \hat{Y}^\top \hat{Y} / C_w^2$. Thus, the feedback design problem can be reduced to the synthesis of an ellipsoidal region for the error system (5b) that is robustly controlled invariant [24, Def. 2.3] by a feedback controller K for all disturbances lying in \mathcal{W} in (16). In other words, we seek

$$\mathcal{E} = \{e \mid e^\top \Phi e \leq 1\}, \quad (17)$$

where $\Phi \succ 0$ is a symmetric positive definite matrix, and a state-feedback gain K , such that

$$[(A + BK)e + w]^\top \Phi [(A + BK)e + w] \leq 1, \ \forall w \in \mathcal{W}, \quad (18)$$

which ensures that \mathcal{E} in (17) is invariant for all disturbance sequences in \mathcal{W} , thereby serving as a PR. By S -procedure [23, Appendix B.2], and the fact that $\Phi \succ 0$ and $Y \succ 0$ [25, Theorem 4.2], the invariance condition in (18) can be translated into finding a feasible solution to the BMI:

$$\begin{bmatrix} \lambda_0 \Phi - \bar{A}^\top \Phi \bar{A} & -\bar{A}^\top \Phi \\ -\Phi \bar{A} & \lambda_1 Y - \Phi \end{bmatrix} \succeq 0, \quad (19a)$$

$$\lambda_0, \lambda_1 \geq 0, \ 1 - \lambda_0 - \lambda_1 \geq 0, \quad (19b)$$

where $\bar{A} = A + BK$. By applying Schur complement twice and noting that $\lambda_1 Y - \Phi \succ 0$ is equivalent to $\Phi^{-1} - 1/\lambda_1 Y^{-1} \succ 0$, (19a) can be enforced by requiring

$$\begin{bmatrix} \hat{\Phi} - 1/\lambda_1 Y^{-1} & A\hat{\Phi} + B\Psi \\ (A\hat{\Phi} + B\Psi)^\top & \lambda_0 \hat{\Phi} \end{bmatrix} \succeq 0, \quad (20)$$

which is linear in $\hat{\Phi}$ and Ψ , where $\hat{\Phi} = \Phi^{-1}$ and $\Psi = K\hat{\Phi}$. To ensure that $K = \Psi\hat{\Phi}^{-1}$ is an admissible linear controller, we also require that $\max_{e^\top \Phi e \leq 1} \|Q^{1/2} K e\| < 1$, which, by using the state-space transformation $\hat{e} = \Phi^{1/2} e$, is written as

$$\begin{bmatrix} \hat{\Phi} & \Psi^\top Q^{1/2} \\ Q^{1/2} \Psi & I \end{bmatrix} \succ 0. \quad (21)$$

We are now ready to state the following result.

Theorem 2 Consider the ellipsoid \mathcal{W} as defined in (16) such that $\Pr\{w^{(0)}(t) \in \mathcal{W} \ \forall t \in \mathbb{N}_{[0, N-1]}\} \geq 1 - \theta$. Let $\hat{\Phi}^*$, Ψ^* , λ_0^* , λ_1^* be a feasible solution to the following program:

$$\begin{aligned} & \text{Minimize} \quad \text{trace } \hat{\Phi} \\ & \hat{\Phi} \succ 0, \ \Psi, \ \lambda_0, \ \lambda_1 \end{aligned} \quad (22a)$$

$$\text{s.t. (19b), (20), (21), } \hat{\Phi} \prec (P_t)^{-1}, \ \forall t \in \mathbb{N}_{[1, N]} \quad (22b)$$

and define $\Phi = (\hat{\Phi}^*)^{-1}$. Consider the error system in (5b), with $K = \Psi^* \Phi$, and the ellipsoidal regions in (4), and define $\mathcal{E} = \{e \mid e^\top \Phi e \leq 1\}$. Then, i) $\Pr\{e^{(0)}(t) \in \mathcal{E} \ \forall t \in \mathbb{N}_{[1, N]}\} \geq 1 - \theta$, ii) $\text{int}(\mathcal{X}_t \ominus \mathcal{E}) \neq \emptyset \ \forall t \in \mathbb{N}_{[1, N]}$, iii) $\text{int}(\mathcal{U} \ominus \mathcal{E}_u) \neq \emptyset$, where $\mathcal{E}_u = \{u = Ke \mid e \in \mathcal{E}\}$.

Proof: i) By the feasibility of (22), the ellipsoid \mathcal{E} is robustly invariant for (5b) for all $w(t) \in \mathcal{W}$ when $K = \Psi^* \Phi$, that is $e(t) \in \mathcal{E}$ for all $w(t) \in \mathcal{W}$. Since $\Pr\{w^{(0)}(t) \in \mathcal{W} \ \forall t \in \mathbb{N}_{[0, N-1]}\} \geq 1 - \theta$, one can deduce that $\Pr\{e^{(0)}(t) \in \mathcal{E} \ \forall t \in \mathbb{N}_{[1, N]}\} \geq 1 - \theta$. ii) Without loss of generality, assume that \mathcal{X}_t

is centered at the origin. By the feasibility of the constraints in (22b), we have that $\Phi \succ P_t$, for $t \in \mathbb{N}_{[1, N]}$, which implies that $\mathcal{E} \subsetneq \text{int}(\mathcal{X}_t)$, $\forall t \in \mathbb{N}_{[1, N]}$. Then, it follows that $\text{int}(\mathcal{X}_t \ominus \mathcal{E}) \neq \emptyset$. iii) Similarly, by the feasibility of the LMI in (21), it follows that $\mathcal{E}_u \subsetneq \text{int}(\mathcal{U})$, which implies that $\text{int}(\mathcal{U} \ominus \mathcal{E}_u) \neq \emptyset$. ■

Note that the optimization in (22) is convex for any fixed pair (λ_0, λ_1) and thus can be solved using a grid search over (λ_0, λ_1) . It also minimizes an upper bound on $\det \hat{\Phi}$, as shown by $\det \hat{\Phi} \leq \left(\frac{1}{n} \text{trace } \hat{\Phi}\right)^n$.

IV. EXAMPLE

We consider the problem in (3)-(4) for a discrete-time double integrator system, with a horizon of $N = 100$ time steps, cost functions $\ell(x(t), u(t)) = (u(t))^2$, $V_f(x(100)) = 100x(100)^\top x(100)$, failure probability $\theta = 0.05$, $A = \begin{bmatrix} 1 & 0.5 \\ 0 & 1 \end{bmatrix}$, $B = \begin{bmatrix} 0 \\ 0.5 \end{bmatrix}$, and initial condition $x(0) = (2, -1)$. For the constraints in (3c)-(3d), we select $P_t = I_2^{1/10}$, $p_t = 0$, for $t \in \mathbb{N}_{[1, N]}$, and $Q = 1$. For the unknown vector $w(t) = (w_1(t), w_2(t))$, we generate calibration and training disturbance datasets, $\mathcal{D}_{\text{train}}^w$ and $\mathcal{D}_{\text{cal}}^w$, each with 100 sequence samples. We sample $w_1(t)$ from $\mathcal{N}(-0.01, 0.005)$, and $w_2(t)$ from a gamma distribution with shape 5.5 and scale 0.005, each multiplied equally likely by 1 or -1 .

Direct method: We formulate the problem in (11), where $\gamma = \eta_e^{\max} / \eta_u^{\max}$, with $\eta_e^{\max} = \sqrt{10}$, $\eta_u^{\max} = 1$, and the constraints in (11b)-(11c) are constructed using the training dataset $\mathcal{D}_{\text{train}}^w$. To solve this problem we employ a genetic algorithm (GA), with a population size of 150 candidates for 50 generations. The GA returns $K_{\text{dir}} = \begin{bmatrix} -0.241 & -0.787 \end{bmatrix}$ within 5 minutes. Next, we obtain PRs, $\mathcal{E}_{0.95}^{1:N}(e(t))$, $\mathcal{E}_{0.95}^{0:N-1}(Ke(t))$, by constructing a calibration trajectory dataset $\mathcal{D}_{\text{cal}}^e$ by the disturbance dataset $\mathcal{D}_{\text{cal}}^w$ as in (8) for the obtained K_{dir} . Following Lemma 3, we obtain $\mathcal{E}_{1-\theta}^{1:N}(e(t)) = \mathbb{B}(0.5785)$ (see blue dashed circle in Fig. 1 (left)) and $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t)) = \mathbb{B}(0.1271)$, which verify the conditions in (6) of Lemma 2.

Indirect method: We construct an ellipsoid $\hat{\mathcal{W}}$ that confines all training disturbance samples in $\mathcal{D}_{\text{train}}^w$ by solving the problem in (13). Then, using the nonconformity score from (14) and computing the empirical quantile from (15) over the calibration dataset $\mathcal{D}_{\text{cal}}^w$ we obtain a tighter ellipsoid $\mathcal{W} \subseteq \hat{\mathcal{W}}$, defined in (16) with $Y = \begin{bmatrix} 12.6733 & -1.0720 \\ -1.0720 & 114.7949 \end{bmatrix}$, which is a PR, $\mathcal{E}_{1-\theta}^{0:N-1}(w(t))$, by Lemma 3. For the obtained PR, \mathcal{W} , we construct a feedback gain $K_{\text{ind}} = \begin{bmatrix} -1.4140 & -2.3412 \end{bmatrix}$, and the ellipsoid $\mathcal{E} = \{e \mid e^\top \Phi e \leq 1\}$, with $\Phi = \begin{bmatrix} 3.4644 & 3.8069 \\ 3.8069 & 5.6494 \end{bmatrix}$, by solving (22). The problem was formulated in YALMIP and solved using the SDP solver SEDUMI, with a grid search over 210 pairs (λ_0, λ_1) , within 37 seconds. Note that \mathcal{E} is a PR, $\mathcal{E}_{1-\theta}^{1:N}(e(t))$, by Theorem 2. For K_{ind} , we obtain a PR, $\mathcal{E}_{0.95}^{0:N-1}(Ke(t)) = \mathbb{B}(0.4408)$, by Lemma 3.

We use the PRs, $\mathcal{E}_{1-\theta}^{1:N}(e(t))$ and $\mathcal{E}_{1-\theta}^{0:N-1}(Ke(t))$, from the direct and indirect methods to tighten the constraints in (3c)-(3d) and solve the deterministic problem in (7). The average solve time for the problem in (7) on a laptop with an Intel i7-1185G7 processor and 32 GB of RAM is less than 0.02 seconds using the MOSEK solver in YALMIP. Using the resulting solutions, we test constraint satisfaction for 10^4 new disturbance realizations. For the direct method, state and input

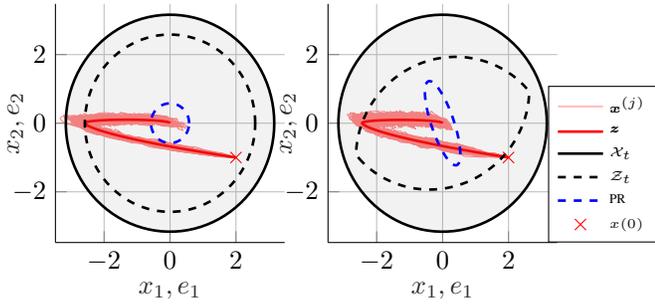


Fig. 1: Direct method (left), Indirect method (right): Representation of the original state constraint \mathcal{X}_t (solid black), the tighter state constraint $\mathcal{Z}_t = \mathcal{X}_t \ominus \mathcal{E}_{0.95}^{1:N}(e(t))$ (dashed black), the PR $\mathcal{E}_{0.95}^{1:N}(e(t))$ (dashed blue), 100 sample trajectories $\mathbf{x}^{(j)}(0 : N)$ (light red), and the deterministic trajectory $\mathbf{z}(0 : N)$ (red), with initial condition $\mathbf{x}(0)$ (red cross).

constraints are satisfied with frequencies of 99.11% and 100%, meeting the specified probability. For the indirect method, state and input constraints are satisfied for all sampled disturbances. Noisy trajectories are illustrated for both methods in Fig. 1.

The conservatism of the indirect method is expected, as the dashed blue ellipsoid in Fig. 1 (right) is a PR for the error state $e^{(0)}(t)$ for all $t > 0$ and any disturbance within $\mathcal{E}_{0.95}(w(t))$. In contrast, the direct method synthesizes a gain K_{dir} from available training disturbance sequences. The merit of our approach is that the two methods can be used independently or in combination. With long horizons, it may be advantageous to collect disturbance samples and solve the problem in (22), followed by determining PRs using Lemma 3.

Comparison with SO: We evaluate our method’s efficiency against the randomized approach from [7], which uses disturbance feedback parameterization [20] to maintain a convex scenario-based relaxation. By enforcing a time-invariant feedback structure (see [7, Sec. III]), we generate scenarios for the constraints in (3b)–(3d) and solve the scenario optimization with 100 scenarios within 1.5 minutes using the MOSEK solver. However, achieving the specified probabilistic constraints, e.g., with confidence level $\beta = 10^{-3}$ requires at least 5,739 scenarios by [7, Theorem 1] and [9, Theorem 1], which demands computational time exceeding one hour using our software. In contrast, our synthesis approach in (11) or (22) does not impose a minimum number of training scenarios. The desired confidence can be achieved offline through calibration by Lemma 3 for any number of scenarios. Notably, calibration for 5,739 scenarios can be completed within 4 seconds on our computer. The software used in this section is on [26].

V. CONCLUSION

We have addressed an optimal control problem for linear stochastic systems with unknown distributions using conformal prediction. We propose two methods to design a feedback controller from data that keeps the error state within a prediction region. We provide guarantees independent of data size through a two-step training-calibration process. In future work, we will extend this approach to a stochastic model predictive control setting, exploring closed-loop properties.

REFERENCES

- [1] A. Mesbah, “Stochastic model predictive control: An overview and perspectives for future research,” *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.
- [2] B. Kouvaritakis, M. Cannon, S. V. Raković, and Q. Cheng, “Explicit use of probabilistic distributions in linear predictive control,” *Automatica*, vol. 46, no. 10, pp. 1719–1724, 2010.
- [3] M. Cannon, B. Kouvaritakis, S. V. Raković, and Q. Cheng, “Stochastic tubes in model predictive control with probabilistic constraints,” *IEEE Transactions on Automatic Control*, vol. 56, no. 1, pp. 194–200, 2011.
- [4] L. Hewing and M. N. Zeilinger, “Stochastic model predictive control for linear systems using probabilistic reachable sets,” in *2018 IEEE Conf. on Decision and Control (CDC)*, 2018, pp. 5182–5188.
- [5] J. Köhler and M. N. Zeilinger, “Recursively feasible stochastic predictive control using an interpolating initial state constraint,” *IEEE Control Systems Letters*, vol. 6, pp. 2743–2748, 2022.
- [6] E. E. Vlahakis, L. Lindemann, P. Sotasakis, and D. V. Dimarogonas, “Probabilistic tube-based control synthesis of stochastic multi-agent systems under signal temporal logic,” 2024, [accepted to CDC24]. [Online]. Available: <https://arxiv.org/abs/2405.02827>
- [7] M. Prandini, S. Garatti, and J. Lygeros, “A randomized approach to stochastic model predictive control,” in *2012 IEEE Conf. on Decision and Control (CDC)*, 2012, pp. 7315–7320.
- [8] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, “Constraint-tightening and stability in stochastic model predictive control,” *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3165–3177, 2017.
- [9] L. Hewing and M. N. Zeilinger, “Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control,” *IEEE Control Systems Letters*, vol. 4, no. 2, pp. 450–455, 2020.
- [10] G. Calafiore and M. C. Campi, “Uncertain convex programs: randomized solutions and confidence levels,” *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.
- [11] M. C. Campi and S. Garatti, “The exact feasibility of randomized solutions of uncertain convex programs,” *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.
- [12] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. Springer Science & Business Media, 2005.
- [13] G. Shafer and V. Vovk, “A tutorial on conformal prediction,” *Journal of Machine Learning Research*, vol. 9, no. 3, pp. 371–421, 2008.
- [14] A. N. Angelopoulos and S. Bates, “A gentle introduction to conformal prediction and distribution-free uncertainty quantification,” 2022. [Online]. Available: <https://arxiv.org/abs/2107.07511>
- [15] Y. Zhao, X. Yu, J. V. Deshmukh, and L. Lindemann, “Conformal predictive programming for chance constrained optimization,” 2024. [Online]. Available: <https://arxiv.org/abs/2402.07407>
- [16] A. Muthali, H. Shen, S. Deglurkar, M. H. Lim, R. Roelofs, A. Faust, and C. Tomlin, “Multi-agent reachability calibration with conformal prediction,” in *62nd IEEE Conf. on Decision and Control (CDC)*, 2023, pp. 6596–6603.
- [17] C. Stamouli, L. Lindemann, and G. Pappas, “Recursively feasible shrinking-horizon MPC in dynamic environments with conformal prediction guarantees,” in *Proc. of the 6th Annual Learning for Dynamics & Control Conf.*, vol. 242. PMLR, 2024, pp. 1330–1342.
- [18] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, “Safe planning in dynamic environments using conformal prediction,” *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 5116–5123, 2023.
- [19] L. Lindemann, Y. Zhao, X. Yu, G. J. Pappas, and J. V. Deshmukh, “Formal verification and control with conformal prediction,” 2024. [Online]. Available: <https://arxiv.org/abs/2409.00536>
- [20] 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.
- [21] R. J. Tibshirani, R. F. Barber, E. J. Candès, and A. Ramdas, “Conformal prediction under covariate shift,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [22] V. Vovk, “Conditional validity of inductive conformal predictors,” in *Proc. of the Asian Conf. on Machine Learning*, vol. 25, 2012, pp. 475–490.
- [23] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [24] F. Blanchini, “Set invariance in control,” *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [25] I. Pólik and T. Terlaky, “A survey of the S-lemma,” *SIAM Review*, vol. 49, no. 3, pp. 371–418, 2007.
- [26] E. E. Vlahakis, https://github.com/leftervisr183/Vlahakis_LCSS2024.