

Safe and Stable Control Synthesis for Uncertain System Models via Distributionally Robust Optimization

Kehan Long* Yinzhuang Yi* Jorge Cortés Nikolay Atanasov

Abstract—This paper considers enforcing safety and stability of dynamical systems in the presence of model uncertainty. Safety and stability constraints may be specified using a control barrier function (CBF) and a control Lyapunov function (CLF), respectively. To take model uncertainty into account, robust and chance formulations of the constraints are commonly considered. However, this requires known error bounds or a known distribution for the model uncertainty, and the resulting formulations may suffer from over-conservatism or over-confidence. In this paper, we assume that only a finite set of model parametric uncertainty samples is available and formulate a distributionally robust chance-constrained program (DRCCP) for control synthesis with CBF safety and CLF stability guarantees. To facilitate efficient computation of control inputs during online execution, we present a reformulation of the DRCCP as a second-order cone program (SOCP). Our formulation is evaluated in an adaptive cruise control example in comparison to 1) a baseline CLF-CBF quadratic programming approach, 2) a robust approach that assumes known error bounds of the system uncertainty, and 3) a chance-constrained approach that assumes a known Gaussian Process distribution of the uncertainty.

I. INTRODUCTION

With the increasing deployment of automatic control systems and robotic platforms in unstructured real-world environments, it is crucial to develop feedback controllers with safety and stability guarantees in the presence of model uncertainty. Enforcing safety by utilizing set invariance properties has become a mainstream approach for constrained control synthesis. Inspired by the property of control Lyapunov functions (CLFs) [1] to yield invariant level sets, control barrier functions (CBFs) [2] were introduced as a tool to verify that a desired safe subset of the state space is invariant. Stability and safety can be considered simultaneously by introducing CLF and CBF constraints on the control input in a quadratic program (QP) formulation for control synthesis [3], [4]. The reliability and efficiency of CLF-CBF-QP control synthesis has been evidenced in several robotic applications, including multi-agent systems [5], aerial robots [6], and walking robots [7].

The notion of safety in the presence of system model uncertainty has been mainly described in two ways: using robust constraints [8], [9] or chance constraints [10], [11]. Studies have also considered system uncertainty when pairing safety with stability in the CLF-CBF-QP formulation. Regarding robust formulations, Choi et al. [12] consider

model disturbances with a compact and convex support set and propose a robust control barrier value function to ensure safety. Similarly, [13] assumes bounded model uncertainty and reformulates the original safety and stability constraints as min-max constraints. Regarding probabilistic formulations, [14], [15] assume a Gaussian Process distribution for the model uncertainty and propose probabilistic versions of the CLF stability and CBF safety constraints. All these approaches require known error bounds or known distributions of the uncertainty. In addition, robust formulations may suffer from over-conservatism due to the worst-case error bounds, while chance-constrained formulations may suffer from over-confidence due to a distributional shift at deployment time.

To tackle such scenarios, we rely on a body of work from the literature on stochastic programming [16] that considers distributionally robust versions of stochastic optimization problems, see e.g. [17], [18]. In particular, distributionally robust chance-constrained programs (DRCCP) deal with uncertain variables in the constraints when only finitely many samples are available. The main idea is to construct an ambiguity ball centered at the empirical distribution obtained from the observed samples and with radius defined using a probability distance function, such as Kullback–Leibler divergence [19] or Wasserstein distance [20]–[24]. In DRCCP, the desired constraints must be satisfied with high probability for all distributions in the constructed ambiguity set. Given the ability to handle uncertainty with unknown or shifting distribution within the ambiguity set, distributionally robust formulations have been used to enforce constraints in reinforcement learning [25], [26] and Markov decision processes [27]–[29]. While these works are closely related, their focus is on discrete-time planning with robustness to uncertainty, while our work considers continuous-time control with safety and stability guarantees.

The contributions of this work are summarized as follows. First, we relax the assumption for safe and stable control synthesis that known error bounds or known distribution of model uncertainty are available by formulating distributionally robust safety and stability constraints using offline model uncertainty samples. Second, we show that the DRCCP control synthesis problem can be reformulated as a second-order cone program (SOCP) in two cases: when there is no restriction on the uncertainty support set and when the uncertainty support set is polyhedral. We demonstrate on an adaptive cruise control problem how our DRCCP SOCP guarantees safety in scenarios with incorrect model uncertainty error bounds or uncertainty distribution shift, in contrast with the vanilla CLF-CBF-QP approach, a robust

* These authors contributed equally.

The authors are with the Contextual Robotics Institute, University of California San Diego, La Jolla, CA 92093, USA (e-mails: {k3long, yiyi, cortes, natanasov}@ucsd.edu).

We gratefully acknowledge support from NSF RI IIS-2007141.

approach, and a chance-constrained approach.

II. PRELIMINARIES

This section reviews control Lyapunov and control barrier functions, distributionally robust modeling, and chance-constrained programming.

A. Optimization-based Control Synthesis

Consider a non-linear control-affine system¹:

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u} = [f(\mathbf{x}) \ g(\mathbf{x})] \cdot \begin{bmatrix} 1 \\ \mathbf{u} \end{bmatrix} \triangleq F(\mathbf{x})\mathbf{u}, \quad (1)$$

where $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state and $\mathbf{u} \in \mathcal{U} := \{1\} \times \mathbb{R}^m$ is the control input. Assume $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ and $g : \mathbb{R}^n \mapsto \mathbb{R}^{n \times m}$ are locally Lipschitz. We start by recalling the notions of CLF [1] and CBF [4], which play a key role in the synthesis of stable and safe controllers, respectively.

Definition II.1. A positive-definite continuously differentiable function $V : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$ is a *control Lyapunov function (CLF)* on \mathcal{X} for system (1) if there exists a class \mathcal{K} function α_V such that:

$$\inf_{\mathbf{u} \in \mathcal{U}} CLC(\mathbf{x}, \mathbf{u}) \leq 0, \quad \forall \mathbf{x} \in \mathcal{X} \setminus \{\mathbf{0}\}, \quad (2)$$

where the *control Lyapunov condition (CLC)* is:

$$CLC(\mathbf{x}, \mathbf{u}) := \mathcal{L}_f V(\mathbf{x}) + \mathcal{L}_g V(\mathbf{x})\mathbf{u} + \alpha_V(V(\mathbf{x})). \quad (3)$$

The existence of a CLF simplifies the stabilization problem considerably because a stabilizing feedback control law can be obtained in terms of the derivatives of the CLF [1].

In addition to stability, it is often necessary to ensure that the closed-loop system trajectories remain within a safe set $\mathcal{C} \subset \mathcal{X}$. To facilitate safe control synthesis, the safe set is specified as the zero superlevel set, $\mathcal{C} := \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) \geq 0\}$, of a function h .

Definition II.2. A continuously differentiable function $h : \mathbb{R}^n \mapsto \mathbb{R}$ is a *control barrier function (CBF)* on \mathcal{X} for system (1) if there exists an extended class \mathcal{K}_∞ function α_h with:

$$\sup_{\mathbf{u} \in \mathcal{U}} CBC(\mathbf{x}, \mathbf{u}) \geq 0, \quad \forall \mathbf{x} \in \mathcal{X}, \quad (4)$$

where the *control barrier condition (CBC)* is:

$$CBC(\mathbf{x}, \mathbf{u}) := \mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u} + \alpha_h(h(\mathbf{x})). \quad (5)$$

Noting that the CLF stability requirement in (2) and the CBF safety requirement in (4) are affine in \mathbf{u} , they can be

¹**Notation.** The sets of real, non-negative real, and natural numbers are denoted by \mathbb{R} , $\mathbb{R}_{\geq 0}$, and \mathbb{N} , respectively. For $N \in \mathbb{N}$, we let $[N] := \{1, 2, \dots, N\}$. We denote the distribution and expectation of a random variable Y by \mathbb{P} and $\mathbb{E}_{\mathbb{P}}(Y)$, respectively. We use $\mathbf{0}_n$ and $\mathbf{1}_n$ to denote the n -dimensional vector with all entries equal to 0 and 1, respectively. For scalar x , we define $(x)_+ := \max(x, 0)$. The L_2 norm for a vector \mathbf{x} is denoted by $\|\mathbf{x}\|$. We denote by $\mathbf{I}_m \in \mathbb{R}^{m \times m}$ the identity matrix and by \otimes the Kronecker product. We use $\text{vec}(\mathbf{X}) \in \mathbb{R}^{nm}$ to denote the vectorization of $\mathbf{X} \in \mathbb{R}^{n \times m}$, obtained by stacking its columns. The gradient of a differentiable function V is denoted by ∇V , while its Lie derivative along a vector field f by $\mathcal{L}_f V = \nabla V \cdot f$. A continuous function $\alpha : [0, a) \rightarrow [0, \infty)$ is of class \mathcal{K} if it is strictly increasing and $\alpha(0) = 0$. A continuous function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ is of extended class \mathcal{K}_∞ if it is of class \mathcal{K} and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$.

enforced as constraints in an optimization problem. Given a baseline controller $\mathbf{k}(\mathbf{x})$, the following QP modifies the controller to guarantee safety and encourage stability:

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}_{\geq 0}} \quad & \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|^2 + \lambda \delta^2 \\ \text{s.t.} \quad & CLC(\mathbf{x}, \mathbf{u}) \leq \delta, CBC(\mathbf{x}, \mathbf{u}) \geq 0, \end{aligned} \quad (6)$$

where $\delta \in \mathbb{R}_{\geq 0}$ is a slack variable that relaxes the CLF constraints to ensure the feasibility of the QP, controlled by the scaling factor $\lambda > 0$.

We are interested in the control synthesis problem in (6) when the system dynamics in (1) are not perfectly known. Considering probabilistic uncertainty in the system model requires probabilistic versions of the safety and stability constraints in (6). We investigate how to handle model uncertainty using samples rather than a known distribution and whether the uncertainty-aware versions of the constraints in (6) remain convex and tractable.

B. Distributionally Robust Chance-constrained Program

To handle probabilistic constraints, we begin by reviewing chance-constrained programming. Throughout the paper we consider a complete separable metric space Ξ with metric d and associate with it a Borel σ -algebra \mathcal{F} and the set $\mathcal{P}(\Xi)$ of Borel probability measures on Ξ . A chance-constrained program (CCP) takes the form:

$$\begin{aligned} \min_{\mathbf{z} \in \mathcal{Z}} \quad & \mathbf{c}^\top \mathbf{z}, \\ \text{s.t.} \quad & \mathbb{P}(G(\mathbf{z}, \boldsymbol{\xi}) \leq 0) \geq 1 - \epsilon, \end{aligned} \quad (7)$$

with closed convex set $\mathcal{Z} \subseteq \mathbb{R}^n$ and uncertainty set $\Xi \subseteq \mathbb{R}^k$. The constraint function $G(\mathbf{z}, \boldsymbol{\xi}) \in \mathcal{Z} \times \Xi \mapsto \mathbb{R}$ depends both on the decision vector \mathbf{z} and an uncertainty vector $\boldsymbol{\xi}$, whose distribution \mathbb{P} is supported on Ξ , and $\epsilon \in (0, 1)$ is a user-specified risk tolerance. The feasible set defined by the chance constraint in (7) is not convex in general. Nemirovski and Shapiro [30] proposed a conservative convex approximation [30] of the feasible set in (7), which consists of replacing the chance constraint by a conditional value-at-risk (CVaR) constraint:

$$\begin{aligned} \min_{\mathbf{z} \in \mathcal{Z}} \quad & \mathbf{c}^\top \mathbf{z}, \\ \text{s.t.} \quad & CVaR_{1-\epsilon}^{\mathbb{P}}(G(\mathbf{z}, \boldsymbol{\xi})) \leq 0. \end{aligned} \quad (8)$$

The feasible set of (8) is a subset of the feasible set of (7). The following paragraph describes a way of defining CVaR.

Value-at-risk (VaR) at confidence level $1 - \epsilon$ for $\epsilon \in (0, 1)$ is defined as $VaR_{1-\epsilon}^{\mathbb{P}_q}(Q) := \inf_{t \in \mathbb{R}} \{t \mid \mathbb{P}_q(Q \leq t) \geq 1 - \epsilon\}$ for a random variable Q with distribution \mathbb{P}_q . VaR does not provide information about the right tail of the distribution, and optimization programs involving VaR variables are intractable in general [31]. To address this, Rockafellar and Uryasev [32] introduced conditional value-at-risk (CVaR), defined as $CVaR_{1-\epsilon}^{\mathbb{P}_q}(Q) = \mathbb{E}_{\mathbb{P}_q}[Q \mid Q \geq VaR_{1-\epsilon}^{\mathbb{P}_q}(Q)]$. CVaR can be also formulated as a convex program:

$$CVaR_{1-\epsilon}^{\mathbb{P}_q}(Q) := \inf_{t \in \mathbb{R}} [\epsilon^{-1} \mathbb{E}_{\mathbb{P}_q}[(Q + t)_+] - t]. \quad (9)$$

Both the formulations in (7) and (8) assume that \mathbb{P} , the true distribution of $\boldsymbol{\xi}$, is known. When this is not the case, one can instead resort to distributionally robust formulations [20], [22]. Assume we only have access to samples $\{\boldsymbol{\xi}_i\}_{i \in [N]}$ from the true distribution of $\boldsymbol{\xi}$. We describe a way of constructing an ambiguity set of distributions that could have potentially generated such samples. Let $\mathcal{P}_p(\Xi) \subseteq \mathcal{P}(\Xi)$ be the set of Borel probability measures with finite p -th moment for $p \geq 1$. The p -Wasserstein distance [23] between two probability measures μ, ν in $\mathcal{P}_p(\Xi)$ is:

$$W_p(\mu, \nu) := \left(\inf_{\gamma \in \mathcal{Q}(\mu, \nu)} \left[\int_{\Xi \times \Xi} d(x, y)^p d\gamma(x, y) \right] \right)^{\frac{1}{p}}, \quad (10)$$

where $\mathcal{Q}(\mu, \nu)$ denotes the measures on $\Xi \times \Xi$ with marginals μ and ν on the first and second factors, and d denotes the metric in the space Ξ .

Let $\hat{\mathbb{P}}_N := \frac{1}{N} \sum_{i=1}^N \delta_{\boldsymbol{\xi}_i}$ denote the discrete empirical distribution constructed from the observed samples $\{\boldsymbol{\xi}_i\}_{i=1}^N$. Using the Wasserstein distance (10), one can define a Wasserstein ambiguity set of radius r centered at $\hat{\mathbb{P}}_N$:

$$\mathcal{M}_N^r := \{\mu \in \mathcal{P}_p(\Xi) \mid W_p(\mu, \hat{\mathbb{P}}_N) \leq r\}, \quad (11)$$

and, in turn, a distributionally robust chance-constrained program (DRCCP):

$$\begin{aligned} & \min_{\mathbf{z} \in \mathcal{Z}} \mathbf{c}^\top \mathbf{z}, \\ \text{s.t.} & \inf_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(G(\mathbf{z}, \boldsymbol{\xi}) \leq 0) \geq 1 - \epsilon. \end{aligned} \quad (12)$$

The constraint in (12) is equivalent to $\sup_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(G(\mathbf{z}, \boldsymbol{\xi}) \geq 0) \leq \epsilon$. Thus, mimicking the convexification for CCP in (8), one can use CVaR to obtain a convex approximation of (12):

$$\begin{aligned} & \min_{\mathbf{z} \in \mathcal{Z}} \mathbf{c}^\top \mathbf{z}, \\ \text{s.t.} & \sup_{\mathbb{P} \in \mathcal{M}_N^r} \text{CVaR}_{1-\epsilon}^{\mathbb{P}}(G(\mathbf{z}, \boldsymbol{\xi})) \leq 0. \end{aligned} \quad (13)$$

III. PROBLEM FORMULATION

We study the problem of enforcing safety and stability of control-affine dynamical systems with model uncertainty. Critically, we do not assume that the probability distribution or error bounds for the model uncertainty are known. We model the uncertainty in the system in (1) using a nominal model $\tilde{F}(\mathbf{x})$ and a linear combination of k perturbations:

$$\dot{\mathbf{x}} = F(\mathbf{x})\mathbf{u} = (\tilde{F}(\mathbf{x}) + \sum_{j=1}^k W_j(\mathbf{x})\boldsymbol{\xi}_j)\mathbf{u}. \quad (14)$$

For $1 \leq j \leq k$, we use $W_j(\mathbf{x}) \in \mathbb{R}^{n \times (m+1)}$ to denote the possible model perturbations, and $\boldsymbol{\xi} \in \mathbb{R}^k$ with elements $\xi_j \in \mathbb{R}$ to denote the corresponding unknown weights. We assume the perturbations $W_j(\mathbf{x})$ are known, while the weights $\boldsymbol{\xi}$ are stochastic. We require a set of historical realizations $\{\boldsymbol{\xi}_i\}_{i=1}^N$ as training data, which can be obtained from past state-control system trajectories using system identification techniques, e.g., based on neural ordinary differential equations [33] or Koopman operator theory [34].

Many control applications require safety and stability guarantees for an uncertain system under online error realizations. This motivates us to consider a distributionally robust formulation for online control synthesis.

Problem 1 (Distributionally Robust Safety and Stability for Uncertain Systems). Consider a nominal model $\tilde{F}(\mathbf{x})$ and perturbation matrices $W_j(\mathbf{x})$, $j \in [k]$ for the system dynamics in (1). Given observations $\{\boldsymbol{\xi}_i\}_{i=1}^N$ of the model uncertainty $\boldsymbol{\xi}$ with support set Ξ , design a feedback controller $\mathbf{k}^* : \mathbb{R}^n \mapsto \mathcal{U}$ with a risk-tolerance parameter $\epsilon \in (0, 1)$ such that, for each $\mathbf{x} \in \mathcal{X}$:

$$\begin{aligned} & \inf_{\mathbb{P} \in \mathcal{M}_N^{r_1}} \mathbb{P}(\text{CLC}(\mathbf{x}, \mathbf{k}^*(\mathbf{x}), \boldsymbol{\xi})) \leq \delta) \geq 1 - \epsilon, \\ & \inf_{\mathbb{P} \in \mathcal{M}_N^{r_2}} \mathbb{P}(\text{CBC}(\mathbf{x}, \mathbf{k}^*(\mathbf{x}), \boldsymbol{\xi})) \geq 0) \geq 1 - \epsilon, \end{aligned} \quad (15)$$

where $\mathcal{M}_N^{r_1}$, $\mathcal{M}_N^{r_2}$ are Wasserstein ambiguity sets with user-specified radii r_1 and r_2 .

While we do not assume a particular distribution for $\boldsymbol{\xi}$, the Wasserstein ball radii r_1, r_2 specify the maximal shift of the true distribution of $\boldsymbol{\xi}$ from the empirical distribution $\hat{\mathbb{P}}_N$ of the historical samples that our method can handle.

We consider two cases based on the information available about the support set Ξ . In the first case, we consider an unbounded support set $\Xi = \mathbb{R}^k$; in the second case, we assume a compact polyhedron set $\Xi = \{\boldsymbol{\xi} \in \mathbb{R}^k \mid \mathbf{C}\boldsymbol{\xi} \leq \mathbf{d}\}$. Inspired by the CLF-CBF-QP in (6), we consider the following DRCCP formulation to enforce safety and stability with high probability and out-of-sample errors by leveraging the CVaR approximations (13) and the CVaR definition (9),

$$\begin{aligned} & \min_{\mathbf{u} \in \mathcal{U}, \delta \in \mathbb{R}} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|^2 + \lambda \delta^2 \\ \text{s.t.} & \sup_{\mathbb{P} \in \mathcal{M}_N^{r_1}} \inf_{t \in \mathbb{R}} [\epsilon^{-1} \mathbb{E}_{\mathbb{P}}[(\text{CLC}(\mathbf{x}, \mathbf{u}, \boldsymbol{\xi}) + t - \delta)_+] - t] \leq 0, \\ & \sup_{\mathbb{P} \in \mathcal{M}_N^{r_2}} \inf_{t \in \mathbb{R}} [\epsilon^{-1} \mathbb{E}_{\mathbb{P}}[(-\text{CBC}(\mathbf{x}, \mathbf{u}, \boldsymbol{\xi}) + t)_+] - t] \leq 0. \end{aligned} \quad (16)$$

Although the constraints in (16) are convex, the program is intractable [20], [23] due to the search of suprema over the Wasserstein ambiguity set. In the following sections, we discuss our approach to identify tractable reformulations of (16) and enable online stable and safe control synthesis.

IV. TRACTABLE REFORMULATION OF CONTROL SYNTHESIS WITH MODEL UNCERTAINTY

This section presents our approach for solving (16). To simplify the notation, we use the vectorization of $F(\mathbf{x})$,

$$\text{vec}(F(\mathbf{x})) = \text{vec}(\tilde{F}(\mathbf{x})) + \mathbf{W}(\mathbf{x})\boldsymbol{\xi}, \quad (17)$$

where

$$\mathbf{W}(\mathbf{x}) = [\text{vec}(W_1(\mathbf{x})) \quad \cdots \quad \text{vec}(W_k(\mathbf{x}))] \in \mathbb{R}^{n(m+1) \times k}.$$

Observe that the CBC expression in (5) is affine in both \mathbf{u} and $\boldsymbol{\xi}$. Using the Kronecker product property $\text{vec}(\mathbf{ABC}) = (\mathbf{C}^\top \otimes \mathbf{A})\text{vec}(\mathbf{B})$ and $\text{vec}(F(\mathbf{x}))$ in (17), we have:

$$\text{CBC}(\mathbf{x}, \mathbf{u}, \boldsymbol{\xi}) = [\nabla_{\mathbf{x}} h(\mathbf{x})]^\top F(\mathbf{x})\mathbf{u} + \alpha_h(h(\mathbf{x}))$$

$$\begin{aligned}
&= \underline{\mathbf{u}}^\top \text{vec}([\nabla_{\mathbf{x}} h(\mathbf{x})]^\top F(\mathbf{x}) \mathbf{I}_{m+1}) + \alpha_h(h(\mathbf{x})) \\
&= \underline{\mathbf{u}}^\top (\mathbf{I}_{m+1} \otimes [\nabla_{\mathbf{x}} h(\mathbf{x})]^\top) \text{vec}(F(\mathbf{x})) + \alpha_h(h(\mathbf{x})) \\
&= \underline{\mathbf{u}}^\top \underbrace{(\mathbf{I}_{m+1} \otimes [\nabla_{\mathbf{x}} h(\mathbf{x})]^\top) \text{vec}(\tilde{F}(\mathbf{x}))}_{\mathbf{q}_h(\mathbf{x})} + \alpha_h(h(\mathbf{x})) \\
&\quad \underline{\mathbf{u}}^\top \underbrace{(\mathbf{I}_{m+1} \otimes [\nabla_{\mathbf{x}} h(\mathbf{x})]^\top) \mathbf{W}(\mathbf{x})}_{\mathbf{R}_h(\mathbf{x})} \boldsymbol{\xi} \\
&= \underline{\mathbf{u}}^\top \mathbf{q}_h(\mathbf{x}) + \underline{\mathbf{u}}^\top \mathbf{R}_h(\mathbf{x}) \boldsymbol{\xi} + \alpha_h(h(\mathbf{x})). \tag{18}
\end{aligned}$$

We can also write $CLC(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi}) = \underline{\mathbf{u}}^\top \mathbf{q}_V(\mathbf{x}) + \underline{\mathbf{u}}^\top \mathbf{R}_V(\mathbf{x}) \boldsymbol{\xi} + \alpha_V(V(\mathbf{x}))$ with similar definitions. Since $\tilde{F}(\mathbf{x})$, $h(\mathbf{x})$, $V(\mathbf{x})$, and $\mathbf{W}(\mathbf{x})$ are known and deterministic, both $CBC(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi})$ and $CLC(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi})$ are affine in $\boldsymbol{\xi}$.

We consider a general optimization program:

$$\begin{aligned}
&\min_{\underline{\mathbf{u}} \in \underline{\mathcal{U}}} \|\underline{\mathbf{u}} - \underline{\mathbf{k}}(\mathbf{x})\|^2, \tag{19} \\
&\text{s.t. } \sup_{\mathbb{P} \in \mathcal{M}_N} \inf_{t \in \mathbb{R}} [\epsilon^{-1} \mathbb{E}_{\mathbb{P}}[(G_l(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi}) + t_l)_+] - t_l] \leq 0, \forall l \in [M],
\end{aligned}$$

where $G_l : \mathcal{X} \times \underline{\mathcal{U}} \times \Xi \mapsto \mathbb{R}$ may represent a safety or stability constraint that is affine in $\boldsymbol{\xi}$:

$$G_l(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi}) = \underline{\mathbf{u}}^\top \mathbf{q}_l(\mathbf{x}) + \underline{\mathbf{u}}^\top \mathbf{R}_l(\mathbf{x}) \boldsymbol{\xi} + \alpha_l(J_l(\mathbf{x})), \tag{20}$$

and J_l is used to represent the certificate function (e.g. CLF, CBF). We write $G_l(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi}_i)$ as $G_l(i)$ for brevity. Depending on the information available about the uncertainty space Ξ , we propose two reformulations of (19). In either case, we assume the metric d of Ξ is the Euclidean distance.

A. Reformulation with Unbounded Uncertainty Space

First, we consider the case with no prior knowledge of Ξ , meaning that $\Xi = \mathbb{R}^k$. We show that the constraints in (19) can be reformulated as second-order cone constraints.

Proposition IV.1 (DRCCP formulation with unbounded support set). *Consider the optimization problem in (19) with G_l in (20), p -Wasserstein distance with $p = 1$, and $\Xi = \mathbb{R}^k$. Then, the following SOCP is equivalent to (19):*

$$\begin{aligned}
&\min_{\underline{\mathbf{u}} \in \underline{\mathcal{U}}, y \in \mathbb{R}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}} y \tag{21} \\
&\text{s.t. } r \|\underline{\mathbf{u}}^\top \mathbf{R}_l(\mathbf{x})\| + \frac{1}{N} \sum_{i=1}^N s_l(i) - t_l \epsilon \leq 0, \\
&\quad s_l(i) \geq G_l(i) + t_l, \quad s_l(i) \geq 0, \quad \forall i \in [N], \quad \forall l \in [M], \\
&\quad y + 1 \geq \sqrt{\|2(\underline{\mathbf{u}} - \underline{\mathbf{k}}(\mathbf{x}))\|^2 + (y - 1)^2}.
\end{aligned}$$

Proof. We start by considering the following program:

$$\begin{aligned}
&\min_{\underline{\mathbf{u}} \in \underline{\mathcal{U}}} \|\underline{\mathbf{u}} - \underline{\mathbf{k}}(\mathbf{x})\|^2 \tag{22} \\
&\text{s.t. } r \|\underline{\mathbf{u}}^\top \mathbf{R}_l(\mathbf{x})\| + \inf_{t \in \mathbb{R}} \left[\frac{1}{N} \sum_{i=1}^N (G_l(i) + t_l)_+ - t_l \epsilon \right] \leq 0.
\end{aligned}$$

Based on [23, Lemma V.8] and assuming $\Xi = \mathbb{R}^k$, the supremum over the Wasserstein ambiguity set (i.e. the constraint in (19)) can be written equivalently as the sample average $\inf_{t \in \mathbb{R}} \left[\frac{1}{N} \sum_{i=1}^N (G_l(i) + t_l)_+ - t_l \epsilon \right]$ and a regularization term $rL(\underline{\mathbf{u}}, \mathbf{x})$, where $L(\underline{\mathbf{u}}, \mathbf{x})$ denotes the Lipschitz constant of G_l in $\boldsymbol{\xi}$.

As defined in (20), for each \mathbf{x} , we can define the convex function $L : \underline{\mathcal{U}} \times \mathcal{X} \mapsto \mathbb{R}_{>0}$ by

$$L(\underline{\mathbf{u}}, \mathbf{x}) = \|\underline{\mathbf{u}}^\top \mathbf{R}_l(\mathbf{x})\|. \tag{23}$$

Then, the function $\boldsymbol{\xi} \mapsto G_l(\mathbf{x}, \underline{\mathbf{u}}, \boldsymbol{\xi})$ is Lipschitz in $\boldsymbol{\xi}$ with constant $L(\underline{\mathbf{u}}, \mathbf{x})$ for fixing \mathbf{x} (assuming $L(\underline{\mathbf{u}}, \mathbf{x}) < \infty$). This is because the Lipschitz constant of a differentiable affine function equals the dual-norm of its gradient [35], and the dual norm of the L_2 norm is itself. This implies that (22) is equivalent to (19).

Next, we show that the bi-level optimization in (22) is equivalent to:

$$\begin{aligned}
&\min_{\underline{\mathbf{u}} \in \underline{\mathcal{U}}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}} \|\underline{\mathbf{u}} - \underline{\mathbf{k}}(\mathbf{x})\|^2 \tag{24} \\
&\text{s.t. } r \|\underline{\mathbf{u}}^\top \mathbf{R}_l(\mathbf{x})\| + \frac{1}{N} \sum_{i=1}^N s_l(i) - t_l \epsilon \leq 0, \\
&\quad s_l(i) \geq G_l(i) + t_l, \quad s_l(i) \geq 0, \quad \forall i \in [N], \quad \forall l \in [M].
\end{aligned}$$

For $i \in [N], l \in [M]$, let $(\underline{\mathbf{u}}_1, t_l^*, s_l(i)^*)$ denote an optimal solution to (24) and $\underline{\mathbf{u}}_2$ an optimal solution to (22), with \hat{t}_l the optimizer for the inf terms in the constraint of (22).

Given $(\underline{\mathbf{u}}_1, t_l^*, s_l(i)^*)$, we have $s_l(i)^* \geq (G_l(i) + t_l^*)_+$ and

$$r \|\underline{\mathbf{u}}_1^\top \mathbf{R}_l(\mathbf{x})\| + \frac{1}{N} \sum_{i=1}^N s_l(i)^* - t_l^* \epsilon \leq 0. \tag{25}$$

Thus, if $s_l(i)^*$ is replaced by $(G_l(i) + t_l^*)_+$ in (25), we conclude that the constraint in (22) is satisfied with $\underline{\mathbf{u}}_1$ and t_l^* . This implies that $\underline{\mathbf{u}}_1$ is also a solution to (22), and the cost satisfies $\|\underline{\mathbf{u}}_1 - \underline{\mathbf{k}}(\mathbf{x})\|^2 \geq \|\underline{\mathbf{u}}_2 - \underline{\mathbf{k}}(\mathbf{x})\|^2$.

Given $\underline{\mathbf{u}}_2$ and \hat{t}_l , for every $i \in [N]$, we choose $\hat{s}_l(i) = (G_l(i) + \hat{t}_l)_+$. This implies $\hat{s}_l(i) \geq G_l(i) + \hat{t}_l$, $\hat{s}_l(i) \geq 0$, and the first constraints in (24) is satisfied since

$$r \|\underline{\mathbf{u}}_2^\top \mathbf{R}_l(\mathbf{x})\| + \frac{1}{N} \sum_{i=1}^N (G_l(i) + \hat{t}_l)_+ - \hat{t}_l \epsilon \leq 0.$$

Thus, $(\underline{\mathbf{u}}_2, \hat{t}_l, \hat{s}_l(i))$ is also a solution to (24). Furthermore, the cost satisfies $\|\underline{\mathbf{u}}_1 - \underline{\mathbf{k}}(\mathbf{x})\|^2 \leq \|\underline{\mathbf{u}}_2 - \underline{\mathbf{k}}(\mathbf{x})\|^2$ since $(\underline{\mathbf{u}}_1, t_l^*, s_l(i)^*)$ is an optimal solution to (24). Therefore, both costs are equal, and (24) and (22) are equivalent.

Finally, by reformulating the objective function of (24) as a linear objective with an SOC constraint [15, Proposition IV.3], we conclude the SOCP (21) is equivalent to (19). \square

Proposition IV.1 allows control synthesis with distributionally robust safety and stability constraints without prior knowledge about the uncertainty support set Ξ . The SOCP in (21) can be solved efficiently online using an off-the-shelf solver (e.g. [36]).

B. Reformulation with Bounded Uncertainty Space

Assuming no prior knowledge about the uncertainty set Ξ may result in an overly conservative controller. This motivates us to also consider the case that the uncertainty support set Ξ is a compact polyhedron.

Proposition IV.2 (DRCCP formulation with bounded polyhedron support set). Consider the optimization problem in (19) with G_l in (20), p -Wasserstein distance with $p = 1$, and compact $\Xi = \{\boldsymbol{\xi} \in \mathbb{R}^k \mid \mathbf{C}\boldsymbol{\xi} \leq \mathbf{d}\}$, where $\mathbf{C} \in \mathbb{R}^{q \times k}$ and $\mathbf{d} \in \mathbb{R}^q$ for some $q > 0$. Then, the following SOCP is equivalent to (19),

$$\begin{aligned} & \min_{\mathbf{u} \in \mathcal{U}, y \in \mathbb{R}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}, \boldsymbol{\eta}_i \in \mathbb{R}^q} y & (26) \\ \text{s.t. } & \beta r + \frac{1}{N} \sum_{i=1}^N s_l(i) - t_l \epsilon \leq 0, \\ & s_l(i) \geq 0, \\ & s_l(i) \geq \mathbf{u}^\top \mathbf{q}_l(\mathbf{x}) + t_l + (\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C}) \boldsymbol{\xi}_i + \boldsymbol{\eta}_i^\top \mathbf{d} + \alpha_l(J_l(\mathbf{x})), \\ & \|\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C}\| \leq \beta, \quad \boldsymbol{\eta}_i \geq \mathbf{0}_q, \quad \forall i \in [N], \quad \forall l \in [M], \\ & y + 1 \geq \sqrt{\|2(\mathbf{u} - \mathbf{k}(\mathbf{x}))\|^2 + (y - 1)^2}. \end{aligned}$$

Proof. Based on [23, Proposition V.1] and [20, Corollary 5.1], we know the following program is equivalent to (19),

$$\begin{aligned} & \min_{\mathbf{u} \in \mathcal{U}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}, \boldsymbol{\eta}_i \in \mathbb{R}^q} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|^2 & (27) \\ \text{s.t. } & \beta r + \frac{1}{N} \sum_{i=1}^N s_l(i) - t_l \epsilon \leq 0, \\ & (\mathbf{u}^\top \mathbf{q}_l(\mathbf{x}) + t_l + (\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C}) \boldsymbol{\xi}_i + \boldsymbol{\eta}_i^\top \mathbf{d} + \alpha_l(J_l(\mathbf{x})))_+ \leq s_l(i), \\ & \|\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C}\| \leq \beta, \quad \boldsymbol{\eta}_i \geq \mathbf{0}_q, \quad \forall i \in [N] \end{aligned}$$

Next, we aim to rewrite (27) as a SOCP. The ReLU-type inequality $(w_i)_+ = (\mathbf{u}^\top \mathbf{q}_l(\mathbf{x}) + t_l + (\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C}) \boldsymbol{\xi}_i + \boldsymbol{\eta}_i^\top \mathbf{d} + \alpha_l(J_l(\mathbf{x})))_+ \leq s_l(i)$ can be written equivalently as two constraints: $s_l(i) \geq w_i$ and $s_l(i) \geq 0$. Following the same technique as in Proposition IV.1, we conclude that (26) is equivalent to (19). \square

Remark IV.3 (Comparison between the two formulations). If $\mathbf{C} = \mathbf{0}_{q \times k}$ and $\mathbf{d} = \mathbf{0}_q$ in Proposition IV.2, then $\Xi = \mathbb{R}^k$ and the SOCP in (26) reduces to (21). \bullet

Remark IV.4 (Different choice of metric d). If instead of L_2 norm, we take the metric d of Ξ to be the L_1 norm, then the optimization problems in Propositions IV.1 and IV.2 become QPs. Details are provided in Appendix I. \bullet

V. EVALUATION

We evaluate the proposed distributionally robust approach for safe and stable control synthesis in an adaptive cruise control problem introduced in [3].

A. Cruise Control Model

Consider a simplified adaptive cruise control model that consists of two vehicles, one leading vehicle traveling at a constant speed and one following vehicle using our control synthesis methodology. The objective is have the following

TABLE I: Parameters used in the simulation results

Variable	Description	Value
g	Gravitational acceleration	9.81
m	Mass of vehicle	1650
f_0	Coefficient in $F_r(v)$	0.1
f_1	Coefficient in $F_r(v)$	5
f_2	Coefficient in $F_r(v)$	0.25
v_d	Desired speed	35
v_0	Speed of leading vehicle	20
c_a	Max accelerate constant	0.3
c_d	Max decelerate constant	-0.3

vehicle achieve a desired speed while keeping a safe distance from the leading vehicle. The system model is:

$$\underbrace{\begin{bmatrix} \dot{p} \\ \dot{z} \end{bmatrix}}_{\dot{\mathbf{x}}} = \underbrace{\begin{bmatrix} v \\ -\frac{1}{m} F_r(v) \\ v_0 - v \end{bmatrix}}_{f(\mathbf{x})} + \underbrace{\begin{bmatrix} 0 \\ \frac{1}{m} \\ 0 \end{bmatrix}}_{g(\mathbf{x})} u, \quad (28)$$

where v and v_0 are the velocities of the following and leading vehicles, respectively, $F_r(v) = f_0 + f_1 v + f_2 v^2$ is the air drag, p is the following vehicle position, and z is the distance to the leading vehicle. The input is constrained by $-c_d g \leq \frac{u}{m} \leq c_a g$, where c_d and c_a denote the factor of g for deceleration and acceleration, respectively. We define a CLF, $V(\mathbf{x}) = (v - v_d)^2$, where v_d is the desired speed of the following vehicle. The safety requirements is specified by the CBF $h(\mathbf{x}) = z - \frac{1}{2} \frac{(v_0 - v)^2}{c_a g} - 1.8v$. We assume that the system (28) is uncertain with the following parametric uncertainty,

$$\dot{\mathbf{x}} = (\tilde{F}(\mathbf{x}) + \sum_{i=1}^3 W_i(\mathbf{x}) \boldsymbol{\xi}_i) \mathbf{u} \quad (29)$$

where $\tilde{F}(\mathbf{x}) = [f(\mathbf{x}) \ g(\mathbf{x})]$, $\mathbf{u} = [1 \ u]^\top$, and:

$$W_1(\mathbf{x}) = \begin{bmatrix} 0 & 0 \\ \frac{v}{20} & 0 \\ 0 & 0 \end{bmatrix}, \quad W_2(\mathbf{x}) = \begin{bmatrix} 0 & 0 \\ 0 & \frac{0.05}{m} \\ 0 & 0 \end{bmatrix}, \quad W_3(\mathbf{x}) = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{2z}{25} & 0 \end{bmatrix},$$

where W_1 , W_2 , and W_3 represent the model perturbations in the drag, input force, and leading vehicle distance, respectively. Table I reports the parameter values used in the simulation.

B. Results

We evaluate our distributionally robust control synthesis approach and illustrate its versatility in handling model uncertainty. We report simulation results from the unbounded uncertainty formulation (Proposition IV.1) and the bounded uncertainty formulation (Proposition IV.2). For comparison, we include results from the CLF-CBF-QP (which takes no model uncertainty into account) formulation in [3] with baseline controller $\mathbf{k}(\mathbf{x}) = [1 \ F_r(v)]^\top$, the robust (which requires prior knowledge on the error bound) and the chance-constrained (which assumes the uncertainty distribution to be Gaussian) formulations in [15]. In the simulation, the error bounds are provided by the support set information and the Gaussian parameters are estimated via offline uncertainty

samples. In all cases, we use the value of the CBF as a measure of the safety ensured by the corresponding approach. We consider different choices of Wasserstein radius $r_1 = r_2 = r$, confidence level ϵ , support set Ξ , offline uncertainty samples $\{\xi_i\}_{i=1}^N$, and online true uncertainty realization ξ^* . To demonstrate that our formulation ensures safety for out-of-sample uncertainty, we use different distributions for sampling offline observations ξ_i and a true online uncertainty realization ξ^* .

We consider 8 cases with different parameter choices, where \mathcal{N} and \mathcal{B} denote normal and beta distributions, respectively. For each case, we conduct 50 simulations with the same $\{\xi_i\}_{i=1}^{N=10}$ and different ξ^* .

Case 1 (Gaussian Distribution): $r = 0.3$, $\epsilon = 0.1$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim \mathcal{N}(-1, 1/3)$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 2 (Confident in Sample): $r = 0.001$, $\epsilon = 0.1$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim \mathcal{N}(-1, 1/3)$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 3 (Out of Sample): $r = 0.3$, $\epsilon = 0.1$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim \mathcal{N}(0, 0.2)$, $\xi^* \sim \mathcal{B}(0.1, 2) - 1$.

Case 4 (Baseline Radius and Confidence): $r = 0.3$, $\epsilon = 0.1$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim 4\mathcal{B}(3, 0.1) - 2$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 5 (Larger Radius): $r = 0.5$, $\epsilon = 0.1$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim 4\mathcal{B}(3, 0.1) - 2$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 6 (Higher Confidence): $r = 0.3$, $\epsilon = 0.05$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim 4\mathcal{B}(3, 0.1) - 2$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 7 (Larger Radius and Higher Confidence): $r = 0.5$, $\epsilon = 0.05$, $\Xi = [-2_3, 2_3]$, $\xi_i \sim 4\mathcal{B}(3, 0.1) - 2$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

Case 8 (Out of Support): $r = 0.3$, $\epsilon = 0.1$, $\Xi = [-0.5_3, 0.5_3]$, $\xi_i \sim \mathcal{B}(2, 0.1) - 0.5$, $\xi^* \sim \mathcal{N}(-1, 1/3)$.

In Table II, we report the failure rate and the average CBF values for the 8 cases above. In *Cases 1* and *2*, under Gaussian uncertainty in the dynamics model, all formulations ensure safety except the CLF-CBF-QP. When we set the Wasserstein radius small ($r = 0.001$), meaning that we are confident in the offline uncertainty samples, the unbounded DRCCP and bounded DRCCP formulations have the same mean CBF values. In *Case 3*, we verify that if the uncertainty distribution shifts during the online phase (e.g., the online uncertainty no longer from a Gaussian distribution), then the Gaussian CLF-CBF-SOCP formulation fails, while the other three formulations ensure safety. *Cases 4* to *7* demonstrate the effects of the Wasserstein distance and confidence level in our bounded and unbounded DRCCP formulations. On the one hand, the unbounded DRCCP formulations tend to be more conservative if we increase the Wasserstein radius r and/or the confidence level, as shown in Fig. 1. On the other hand, only increasing the confidence level makes the bounded DRCCP controller more conservative, since support information provides a tighter bound than the Wasserstein radius. In *Case 8*, we see that the unbounded DRCCP formulation works well even with out-of-support uncertainty, while the robust CLF-CBF-SOCP and bounded DRCCP both fail due to the provided incorrect support set information, as Fig. 2 shows.

Generally, the controller provided by the bounded DRCCP formulation has the best performance in ensuring safety




Fig. 1: CBF value of one of the 50 simulations corresponding to *Case 5*. Both the unbounded and bounded DRCCP formulations ensure safety while the CLF-CBF-QP and the Gaussian formulation fail. This demonstrates that either the bounded or unbounded DRCCP formulation ensures safety for out-of-sample uncertainty. The unbounded DRCCP formulation is more conservative since it does not take the uncertainty support set information into account.

while not being too conservative (smaller average CBF values). However, if one fails to provide reliable support set information, then the controller provided by the unbounded DRCCP formulation is the safe choice.

TABLE II: Failure rate and average CBF values. The results are shown in the following format: $a\% \mid b$, where $a\%$ denotes the violation rate of each formulation: (simulations with unsafe state)/(total simulations), and b denotes the average value of CBF over all simulations. The average CBF value is computed based on stabilized CBF values, e.g., for $5 \leq t \leq 15$ in Fig. 2.

Case	Unbounded DRCCP	Bounded DRCCP	Robust CLF-CBF-SOCP	Gaussian CLF-CBF-SOCP	Original CLF-CBF-QP
1	0% 2.11	0% 0.56	0% 1.01	0% 0.55	100% -0.48
2	0% 0.35	0% 0.35	0% 1.01	0% 0.56	100% -0.48
3	0% 1.23	0% 0.55	0% 0.96	98% -0.17	100% -0.50
4	0% 0.56	0% 0.51	0% 0.89	100% -0.92	100% -0.53
5	0% 1.42	0% 0.57	0% 1.01	100% -0.88	100% -0.48
6	0% 1.94	0% 0.57	0% 1.01	100% -0.65	100% -0.48
7	0% 4.49	0% 0.57	0% 1.01	100% -0.64	100% -0.49
8	0% 1.29	98% -0.26	84% -0.16	100% -0.57	100% -0.51

VI. CONCLUSIONS

We considered the problem of enforcing safety and stability of uncertain control-affine systems. Compared with previous approaches, we derive new distributionally robust chance constrained formulations of safe and stable control synthesis that do not require any prior knowledge of error bounds or uncertainty distributions. Using only offline model uncertainty samples, we show that our formulations ensure safety and stability with out-of-sample errors during online execution. Future work will consider deploying the algorithms on real autonomous systems and learning the perturbation matrices and uncertainty samples from offline state-control sequences.




Fig. 2: CBF value of one of the 50 simulations corresponding to Case 8. The offline uncertainty distribution is set to be within the uncertainty support set Ξ , while the online uncertainty distribution is outside of Ξ . The controller obtained with the bounded DRCCP formulation (26) fails to guarantee safety because the assumptions in Proposition IV.2 are violated. However, the controller obtained with the unbounded DRCCP formulation (21) still guarantees safety.

REFERENCES

- [1] E. D. Sontag, "A 'universal' construction of Artstein's theorem on nonlinear stabilization," *Systems and Control Letters*, vol. 13, no. 2, pp. 117–123, 1989.
- [2] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," in *IFAC Proceedings Volumes*, pp. 462–467, 2007.
- [3] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *IEEE Conference on Decision and Control*, pp. 6271–6278, 2014.
- [4] A. Ames, X. Xu, J. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [5] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.
- [6] G. Wu and K. Sreenath, "Safety-critical control of a planar quadrotor," in *American Control Conference*, pp. 2252–2258, 2016.
- [7] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3d dynamic walking on stepping stones with control barrier functions," in *IEEE Conference on Decision and Control*, pp. 827–834, 2016.
- [8] R. A. Freeman and P. V. Kokotovic, *Robust Nonlinear Control Design*. Cambridge, MA, USA: Birkhäuser Boston Inc., 1996.
- [9] I. R. Petersen and R. Tempo, "Robust control of uncertain systems: Classical results and recent developments," *Automatica*, vol. 50, pp. 1315–1335, 2014.
- [10] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram, "Chance-constrained dynamic programming with application to risk-aware robotic space exploration," *Autonomous Robots*, vol. 39, no. 4, p. 555–571, 2015.
- [11] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, pp. 776–783, 2019.
- [12] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier–value functions for safety-critical control," in *IEEE Conference on Decision and Control*, pp. 6814–6821, 2021.
- [13] Q. Nguyen and K. Sreenath, "Robust safety-critical control for dynamic robotics," *IEEE Transactions on Automatic Control*, vol. 67, pp. 1073–1088, 2022.
- [14] V. Dhiman*, M. J. Khojasteh*, M. Franceschetti, and N. Atanasov, "Control barriers in bayesian learning of system dynamics," *IEEE Transactions on Automatic Control*, 2021.
- [15] K. Long, V. Dhiman, M. Leok, J. Cortés, and N. Atanasov, "Safe control synthesis with uncertain dynamics and constraints," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7295–7302, 2022.

- [16] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming*, vol. 16. Philadelphia, PA: SIAM, 2014.
- [17] A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski, *Robust Optimization*. Applied Mathematics Series, Princeton, NJ: Princeton University Press, 2009.
- [18] A. Shapiro, "Distributionally robust stochastic programming," *SIAM Journal on Optimization*, vol. 27, no. 4, p. 2258–2275, 2017.
- [19] R. Jiang and Y. Guan, "Data-driven chance constrained stochastic program," *Mathematical Programming*, vol. 158, pp. 291–327, 2016.
- [20] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, pp. 115–166, 2018.
- [21] Z. Chen, D. Kuhn, and W. Wiesemann, "Data-driven chance constrained programs over Wasserstein balls," *arXiv: Optimization and Control*, 2018.
- [22] W. Xie, "On distributionally robust chance constrained programs with Wasserstein distance," *Mathematical Programming*, vol. 186, pp. 115–155, 2021.
- [23] A. R. Hota, A. K. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under Wasserstein ambiguity sets," *American Control Conference*, pp. 1501–1506, 2019.
- [24] D. Boskos, J. Cortés, and S. Martínez, "Data-driven ambiguity sets with probabilistic guarantees for dynamic processes," *IEEE Transactions on Automatic Control*, vol. 66, no. 7, pp. 2991–3006, 2021.
- [25] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone, "Risk-constrained reinforcement learning with percentile risk criteria," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6070–6120, 2017.
- [26] E. Smirnova, E. Dohmatob, and J. Mary, "Distributionally robust reinforcement learning," *arXiv preprint arXiv:1902.08708*, 2019.
- [27] H. Xu and S. Mannor, "Distributionally Robust Markov Decision Processes," *Mathematics of Operations Research*, vol. 37, no. 2, pp. 288–300, 2012.
- [28] Z. Chen, P. Yu, and W. B. Haskell, "Distributionally robust optimization for sequential decision-making," *Optimization*, vol. 68, no. 12, pp. 2397–2426, 2019.
- [29] M. Petrik and R. H. Russel, "Beyond confidence regions: Tight Bayesian ambiguity sets for robust MDPs," *Advances in neural information processing systems*, vol. 32, 2019.
- [30] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, pp. 969–996, 2006.
- [31] H. Mausser and D. Rosen, "Beyond var: from measuring risk to managing risk," in *Proceedings of the IEEE/IAFE 1999 Conference on Computational Intelligence for Financial Engineering*, pp. 163–178, 1999.
- [32] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–41, 2000.
- [33] T. Duong and N. Atanasov, "Adaptive control of se(3) hamiltonian dynamics with learned disturbance features," *IEEE Control Systems Letters*, vol. 6, pp. 2773–2778, 2022.
- [34] S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, and C. Schütte, "Data-driven approximation of the koopman generator: Model reduction, system identification, and control," *Physica D: Nonlinear Phenomena*, vol. 406, p. 132416, 2020.
- [35] S. Shalev-Shwartz, "Online learning and online convex optimization," *Found. Trends Mach. Learn.*, vol. 4, no. 2, p. 107–194, 2012.
- [36] M. ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019.

APPENDIX I

DIFFERENT CHOICE OF METRIC d

We show that when the metric d of the uncertainty support set Ξ is the L_1 norm (instead of the Euclidean norm as in Propositions IV.1 and IV.2), then (19) becomes a QP for both the cases of unbounded and bounded uncertainty sets.

Proposition I.1 (DRCCP formulation with unbounded support set under L_1 norm). *Consider the optimization problem in (19) with G_l in (20), p -Wasserstein distance with*

$p = 1$, and $\Xi = \mathbb{R}^k$ with metric $d(\xi, \xi') = \|\xi - \xi'\|_1$. Then, the following QP is equivalent to (19):

$$\begin{aligned} & \min_{\mathbf{u} \in \mathcal{U}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|^2 & (30) \\ \text{s.t. } & r|\mathbf{R}_l^\top(\mathbf{x})\mathbf{u}| \leq (t_l\epsilon - \frac{1}{N} \sum_{i=1}^N s_l(i))\mathbf{1}_k, \\ & s_l(i) \geq G_l(i) + t_l, \quad s_l(i) \geq 0, \quad \forall i \in [N], l \in [M]. \end{aligned}$$

Proposition I.2 (DRCCP formulation with bounded polyhedron set under L_1 norm). Consider the optimization problem in (19) with G_l in (20), p -Wasserstein distance with $p = 1$, and compact $\Xi = \{\xi \in \mathbb{R}^k \mid \mathbf{C}\xi \leq \mathbf{d}\}$ with $\mathbf{C} \in \mathbb{R}^{q \times k}$ and $\mathbf{d} \in \mathbb{R}^q$ for some $q > 0$ and metric $d(\xi, \xi') = \|\xi - \xi'\|_1$. Then, the following QP is equivalent to (19):

$$\begin{aligned} & \min_{\mathbf{u} \in \mathcal{U}, t_l \in \mathbb{R}, s_l(i) \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}, \boldsymbol{\eta}_i \in \mathbb{R}^q} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|^2 & (31) \\ \text{s.t. } & \beta r + \frac{1}{N} \sum_{i=1}^N s_l(i) - t_l\epsilon \leq 0, \\ & s_l(i) \geq 0, \\ & s_l(i) \geq \mathbf{u}^\top \mathbf{q}_l(\mathbf{x}) + t_l + (\mathbf{u}^\top \mathbf{R}_l(\mathbf{x}) - \boldsymbol{\eta}_i^\top \mathbf{C})\xi_i + \boldsymbol{\eta}_i^\top \mathbf{d} + \alpha_l(J_l(\mathbf{x})), \\ & |\mathbf{R}_l^\top(\mathbf{x})\mathbf{u} - \mathbf{C}^\top \boldsymbol{\eta}_i| \leq \beta \mathbf{1}_k, \quad \boldsymbol{\eta}_i \geq \mathbf{0}_q, \quad \forall i \in [N], l \in [M]. \end{aligned}$$

We provide a proof sketch for these results. When d is the L_1 norm, the Lipschitz constant in (23) is defined by the L_∞ norm, since the dual norm of L_1 is L_∞ . Similarly, the L_2 norm in the fourth constraint in Proposition IV.2 is replaced by the L_∞ norm. This means that we no longer need to reformulate the objective function, since all constraints are linear in the decision variables and both optimization problems are QPs.