

; revised .

Distributionally Robust Policy and Lyapunov-Certificate Learning

Kehan Long¹ (Student Member), Jorge Cortés¹ (Fellow), Nikolay Atanasov¹ (Senior Member)

¹Contextual Robotics Institute, University of California San Diego, La Jolla, CA 92093, USA

CORRESPONDING AUTHOR: Kehan Long (e-mail: k3long@ucsd.edu)

This work was supported by ONR Award N00014-23-1-2353 and NSF CCF-2112665 (TILOS).

ABSTRACT This article presents novel methods for synthesizing distributionally robust stabilizing neural controllers and certificates for control systems under model uncertainty. A key challenge in designing controllers with stability guarantees for uncertain systems is the accurate determination of and adaptation to shifts in model parametric uncertainty during online deployment. We tackle this with a novel distributionally robust formulation of the Lyapunov derivative chance constraint ensuring a monotonic decrease of the Lyapunov certificate. To avoid the computational complexity involved in dealing with the space of probability measures, we identify a sufficient condition in the form of deterministic convex constraints that ensures the Lyapunov derivative constraint is satisfied. We integrate this condition into a loss function for training a neural network-based controller and show that, for the resulting closed-loop system, the global asymptotic stability of its equilibrium can be certified with high confidence, even with Out-of-Distribution (OoD) model uncertainties. To demonstrate the efficacy and efficiency of the proposed methodology, we compare it with an uncertainty-agnostic baseline approach and several reinforcement learning approaches in two control problems in simulation. Open-source implementations of the examples are available at https://github.com/KehanLong/DR_Stabilizing_Policy.

INDEX TERMS Learning for control, Lyapunov methods, optimization under uncertainty, stability of nonlinear systems

I. Introduction

In control theory and robotics, the task of synthesizing stabilizing controllers with provable certificates for open-loop systems under model uncertainty remains an important challenge. A Lyapunov function (LF) [2] is a fundamental tool for asserting the asymptotic stability of ordinary differential equations (ODEs). Existing approaches in deriving stabilizing controllers for uncertain nonlinear systems usually rely on known bounds or probabilistic models of the uncertainty, which may be hard to obtain in practice. In this work, we address this challenge by exploring the potential of neural networks to synthesize distributionally robust stabilizing controllers and associated Lyapunov stability certificates, providing a new approach for ensuring robustness to model uncertainty.

The use of Lyapunov functions for ensuring the stability of nonlinear systems is well documented in the literature

A preliminary version of this paper appeared as [1] at the 2023 Learning for Dynamics and Control Conference.

[3]–[5]. The early 2000s marked a significant development in the field with the introduction of sum-of-square (SOS) polynomials to identify LFs through semi-definite programming (SDP) [6], [7]. These approaches, however, often necessitate polynomial approximations of the system dynamics and scale poorly in higher dimensional spaces. Moreover, traditional SOS methods face difficulties in the joint synthesis of stabilizers and LF certificates for open-loop systems [8]. With the remarkable advancements in using neural networks as function approximators, researchers [9]–[15] have started representing LFs and/or control policies using neural networks. This progress holds great potential for developing stabilizing controllers and certificates for complex nonlinear systems, offering a vastly expanded functional space compared to SOS polynomials. Despite these advancements, a critical gap persists in deriving controllers equipped with suitable stability guarantees for systems with model uncertainty, an aspect paramount to the practical application of control and robotic systems in real-world scenarios. To

address the challenges of ensuring stability and stabilization in systems with model uncertainty, a related class of works [16]–[21] has focused on adapting stability certificates for nominal systems by incorporating model uncertainty during deployment. These approaches commonly assume knowing the nominal certificates or controllers as well as the error bounds or distribution of the model uncertainty.

Our work aims to develop a distributionally robust Lyapunov function and controller pair for systems with model uncertainty. Instead of relying on knowledge of the distribution or error bounds of the model uncertainty, we seek to synthesize a controller that ensures the global asymptotic stability of the controlled system, even in the presence of distribution shifts during online deployment. To do this, we rely on recent advances in distributionally robust optimization (DRO) [22], [23]. Distributionally robust chance constraints handle uncertainty in constraints with a limited number of available samples. The core strategy involves constructing an ambiguity set around the empirical distribution derived from these samples, defined by a radius using a probability distance measure such as the Wasserstein distance [23], [24]. The objective is to ensure constraint satisfaction with high confidence for all distributions within this set. This key benefit of the approach is its robust performance guarantee for distributional uncertainty, i.e., uncertainty in the possible probability distribution of the samples.

The growing field of distributionally robust control (DRC) [25]–[34] offers tools for control design under uncertainty. DRC has been successfully applied to various control domains, including model predictive control [31]–[33], differential dynamic programming [29], and linear quadratic control [34]. The main objective of DRC is to design controllers that provide robust performance and constraint satisfaction under distributional shifts, making them well-suited for real-world applications where the true uncertainty distribution may differ from the estimated distribution. Building upon the principles of DRC, our work focuses on developing a distributionally robust LF and controller pair for control systems with model uncertainty.

This paper is an extended version of our previous conference publication [1], which introduced distributionally robust Lyapunov stability for closed-loop systems with model uncertainty. In this work, we extend our approach to the joint synthesis of distributionally robust controllers and Lyapunov certificates for nonlinear control systems. Our **contributions** include the following.

- Given samples of uncertain model parameters, we propose a method for joint synthesis of a distributionally robust controller and a Lyapunov certificate for a nonlinear control system subject to affine model uncertainty;
- We identify a sufficient convex constraint that ensures the distributionally robust Lyapunov condition is satisfied and employ it to learn a stabilizing neural network controller and Lyapunov function pair;

- We show that the resulting neural network controller guarantees the global asymptotic stability in probability of the closed-loop system;
- We validate in both theory and experiments that uncertain systems governed by the learned controller are stable with high confidence, even under out-of-distribution model uncertainty.

II. Related Work

This section reviews related work on synthesizing Lyapunov-stable controllers and distributionally robust optimization.

Stabilizing Control Design. Lyapunov theory plays a pivotal role in synthesizing stabilizing controllers for nonlinear systems. It offers strong asymptotic stability guarantees for the controlled system, which are particularly useful in real-world applications. For linear systems, synthesizing Lyapunov-stable controllers can be efficiently accomplished by linear quadratic regulator techniques [35]. For control-affine systems with polynomial dynamics, researchers [8] have successfully obtained controller and certificate pairs by solving SOS programs. However, the absence of a valid SOS Lyapunov function does not necessarily imply the system's instability, as certain positive-definite functions are not SOS-representable [36]. Moreover, the computational complexity of SOS techniques limits their applicability to relatively simple polynomial systems.

Recently, there has been an increasing focus on leveraging neural networks to address the limitations of SOS-based methods. Neural networks offer more general function representations compared to SOS polynomials, thereby enhancing the capacity to handle a broader class of systems.

Richards *et al.* [10] introduced a neural network-based method to learn the region of attraction for given controllers in discrete-time systems. Chang *et al.* [9] focused on synthesizing controllers and neural network Lyapunov functions for nonlinear continuous-time systems, with verification and improvement of the learned functions by satisfiability modulo theories (SMT) solvers. Boffi *et al.* [37] improved the efficiency of learning Lyapunov functions by integrating positive-definiteness and equilibrium conditions directly into the network architecture. Dai *et al.* [11] developed an approach for discrete-time systems that synthesizes neural network stabilizing controllers and Lyapunov functions, refining controller performance using mixed integer programs (MIP) verifiers. Gaby *et al.* [12] theoretically certified neural network Lyapunov functions by analyzing the approximation power of networks with specific architectures. Zhou *et al.* [38] proposed a framework for learning nonlinear systems, stabilizing controllers, and Lyapunov functions with stability guarantees verified by SMT solvers. Dawson *et al.* [39] extended these ideas to learning safety certificates as control Lyapunov-barrier functions for control-affine systems with bounded convex-hull uncertainty. A more extensive survey by Dawson *et al.* [40] details recent advancements in this line of research. Another related line of research is the use

of CLFs to improve the sample efficiency and robustness of reinforcement learning (RL) algorithms [41], [42].

Despite these developments, there has been limited focus on synthesizing Lyapunov-stable controllers with provable certificates for systems under model uncertainty, particularly in settings where bounds or the distribution of the model uncertainty are challenging to determine.

Distributionally Robust Optimization. Distributionally robust optimization has emerged as a powerful approach to tackle uncertainty in optimization problems, particularly advantageous when dealing with a limited number of samples. To account for the discrepancy between the empirical and true distributions of the uncertainty samples, researchers have introduced several uncertainty descriptors, such as moment ambiguity sets [43], Kullback–Leibler ambiguity sets [44], and Wasserstein ambiguity sets [23], [24], [45]. The DRO framework has been increasingly recognized for its robust performance guarantees against distributional uncertainty, leading to its widespread application in machine learning [46], [47], uncertainty quantification [48], control theory [25], [30], [49]–[51], and robotics [26], [52].

Long *et al.* [53] proposed a distributionally robust formulation for safe stabilizing control under model uncertainty, assuming nominal safe and stable certificates are provided. Ren and Majumdar [52] introduced a DRO framework that enhances policy robustness by iteratively training with adversarial environments generated through a learned generative model. Liviu *et al.* [31] introduced Wasserstein tube MPC for stochastic systems, enhancing robustness and efficiency by effectively analyzing state trajectory uncertainties with limited noise samples. In [27], a distributionally robust model predictive control formulation with the Wasserstein metric is proposed, with its parameters efficiently optimized through reinforcement learning. Hakobyan and Yang [29] developed a tractable distributionally robust differential dynamical programming method for stochastic systems by utilizing the Kantorovich duality principle. Distributionally robust RRT for motion planning that addresses model and environmental uncertainties was developed using moment-based ambiguity sets by Summers [54] and using Wasserstein-metric-based ambiguity sets by Lathrop *et al.* [28].

III. Background

This section introduces our notation and offers a brief review of Lyapunov functions.

A. Notation

We denote the sets of real, non-negative real, and natural numbers by \mathbb{R} , $\mathbb{R}_{\geq 0}$, and \mathbb{N} , respectively. For $N \in \mathbb{N}$, we let $[N] := \{1, 2, \dots, N\}$. For a scalar x , we define $(x)_+ := \max(x, 0)$. We denote by $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ the identity matrix and $\mathbf{0}_n \in \mathbb{R}^n$ the zero vector. For a vector \mathbf{x} and a matrix \mathbf{X} , we use $\|\mathbf{x}\|$ and $\|\mathbf{X}\|$ to denote the L_2 norm and the spectral norm, respectively. We use $B(\mathbf{x}; \delta)$ to denote the open ball centered at \mathbf{x} with radius δ . The closure of $B(\mathbf{x}; \delta)$ is denoted by $\bar{B}(\mathbf{x}; \delta)$. The gradient of a differentiable function V is

denoted by ∇V . We consider a complete separable metric space Ξ with metric d and Borel σ -algebra \mathcal{F} . The set of Borel probability measures on Ξ is denoted as $\mathcal{P}(\Xi)$. For a random variable ξ with distribution supported on Ξ , we denote the distribution and expectation by \mathbb{P}^* and $\mathbb{E}_{\mathbb{P}^*}(\xi)$, respectively.

B. System Stabilization Certified by Lyapunov Functions

Consider the controlled continuous-time system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \quad (1)$$

where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$ denotes the system state and $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ is the control input. We assume $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^n$ is locally Lipschitz, and the origin $\mathbf{x} = \mathbf{0}_n$ is the desired equilibrium of the unforced system, i.e., $\mathbf{f}(\mathbf{0}_n, \mathbf{0}_m) = \mathbf{0}_n$. Plugging a Lipschitz controller $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x})$ in (1) results in the closed-loop system

$$\dot{\mathbf{x}} = \mathbf{f}_{\text{cl}}(\mathbf{x}) := \mathbf{f}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x})). \quad (2)$$

In this context, a Lyapunov function (LF) $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is used to analyze the asymptotic stability of the equilibrium. An LF for (2) satisfies the following conditions:

$$V(\mathbf{0}_n) = 0, V(\mathbf{x}) > 0, \forall \mathbf{x} \neq \mathbf{0}_n, \quad (3a)$$

$$\dot{V}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x})) = \nabla V(\mathbf{x})^\top \mathbf{f}_{\text{cl}}(\mathbf{x}) < 0, \forall \mathbf{x} \neq \mathbf{0}_n. \quad (3b)$$

The existence of a radially unbounded ($V(\mathbf{x}) \rightarrow \infty$ as $\|\mathbf{x}\| \rightarrow \infty$) LF implies global asymptotic stability [55] of the dynamical system (2).

When the system model (1) is perfectly known, classical work [4], [56] has shown how to synthesize stabilizing controllers when a control Lyapunov function is available. This has provided the theoretical basis for work aimed at learning Lyapunov-stable controllers [8], [9], [37]. The availability of an LF-controller pair provides a complete solution, as the controller ensures system stabilization and the LF certifies its validity. However, real-world scenarios often present uncertainty in the dynamics model (1) that vary based on different physical parameters or operational conditions. Our goal is to find an LF-controller pair when uncertainty is present in the model. In particular, we do not assume any known error bounds or known distribution for the system uncertainty and instead investigate how to handle it based on finitely many realizations.

IV. Problem Formulation

We aim to synthesize a Lyapunov-stable controller for the continuous-time system with model uncertainty:

$$\dot{\mathbf{x}} = \bar{\mathbf{f}}(\mathbf{x}, \mathbf{u}, \boldsymbol{\xi}) := \mathbf{f}(\mathbf{x}, \mathbf{u}) + \mathbf{W}(\mathbf{x}, \mathbf{u})\boldsymbol{\xi}, \quad (4)$$

where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$ denotes the state vector, $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ is the control input, and the functions $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $\mathbf{W} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^{n \times k}$ are locally Lipschitz. The set $\mathcal{X} \subset \mathbb{R}^n$ represents the domain of interest, which includes the origin, and $\mathcal{U} \subset \mathbb{R}^m$ denotes the set of admissible control inputs. We assume \mathcal{X} and \mathcal{U} are compact.

Each column of \mathbf{W} represents a specific direction of uncertainty (e.g., mass, friction). The vector $\boldsymbol{\xi} =$

$[\xi_1, \xi_2, \dots, \xi_k]^\top$ characterizes the uncertainty in the system model, and the distribution \mathbb{P}^* of ξ is supported on a compact set $\Xi \subset \mathbb{R}^k$. Additionally, we assume that $\mathbf{W}(\mathbf{0}_n, \mathbf{0}_m) = \mathbf{0}_{n \times k}$, which ensures the origin is the desired equilibrium.

With the aim of stabilizing a system with model uncertainty in (4), we review chance-constraint formulations and their convex approximations. The Conditional Value-at-Risk (CVaR) is particularly useful in this context. For a random variable Q with distribution \mathbb{P}_q , the Value-at-Risk (VaR) at confidence level $1 - \epsilon$ is defined as $\text{VaR}_{1-\epsilon}^{\mathbb{P}_q}(Q) := \inf_{t \in \mathbb{R}} \{t \mid \mathbb{P}_q(Q \leq t) \geq 1 - \epsilon\}$. Based on VaR, CVaR is defined [57] as $\text{CVaR}_{1-\epsilon}^{\mathbb{P}_q}(Q) = \mathbb{E}_{\mathbb{P}_q}[Q \mid Q \geq \text{VaR}_{1-\epsilon}^{\mathbb{P}_q}(Q)]$. Fig. 1 shows an illustration of VaR and CVaR.




FIGURE 1: Illustration of VaR and CVaR within the distribution of a random variable. $\text{VaR}_{1-\epsilon}$ is the lower $1 - \epsilon$ percentile of the random variable while $\text{CVaR}_{1-\epsilon}$ computes the expected value of the realizations above VaR.

Our goal is to find a function $V : \mathbb{R}^n \mapsto \mathbb{R}$ and a controller $\pi : \mathbb{R}^n \mapsto \mathbb{R}^m$ that satisfy the Lyapunov conditions in (3). However, the uncertainty in the dynamical system (4), which manifests itself in the term $\dot{V}(\mathbf{x}, \pi(\mathbf{x}))$, poses a challenge for ensuring that the condition on the derivative of the Lyapunov function (3b) are satisfied. Therefore, we consider the chance-constrained condition,

$$\mathbb{P}^*(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}(\mathbf{x}, \pi(\mathbf{x})) + \gamma \|\mathbf{x}\|) \leq 0) \geq 1 - \epsilon, \quad (5)$$

where $\gamma \in \mathbb{R}_{>0}$ and $\epsilon \in (0, 1)$ is a user-specified risk tolerance. The chance constraint (5) requires that, with probability at least $1 - \epsilon$, the function $\dot{V}(\mathbf{x}, \pi(\mathbf{x})) + \gamma \|\mathbf{x}\|$ is less than or equal to zero for all $\mathbf{x} \in \mathcal{X}$. This condition ensures that the Lyapunov derivative is negative definite with high probability.

The feasible set of the probabilistic constraint (5) is generally non-convex. To address this, Nemirovski and Shapiro [58] proposed a convex CVaR approximation of the chance constraint in (5) as follows,

$$\text{CVaR}_{1-\epsilon}^{\mathbb{P}^*}(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}(\mathbf{x}, \pi(\mathbf{x})) + \gamma \|\mathbf{x}\|)) \leq 0. \quad (6)$$

This inequality implies the satisfaction of the original chance constraint in (5). As shown in [57], the inequality (6) can be

written equivalently as:

$$\inf_{t \in \mathbb{R}} [\epsilon^{-1} \mathbb{E}_{\mathbb{P}^*}[(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}(\mathbf{x}, \pi(\mathbf{x})) + \gamma \|\mathbf{x}\|) + t)_+] - t] \leq 0. \quad (7)$$

Using the formulations in (5) and (6) requires knowledge of the distribution \mathbb{P}^* . However, in robotics and control applications typically \mathbb{P}^* is unknown but a limited number of samples $\{\xi_i\}_{i \in [N]}$ of the uncertain system parameters is available instead. For example, these samples can be obtained by collecting state-action sequences from a real system or a high-fidelity simulation and comparing the observed states with the predicted states based on a nominal model. The differences between the true and estimated states can be used as samples of the uncertain model parameters. Moreover, the distribution of ξ may shift at system deployment time. This motivates us to consider the following problem.

Problem 1 (Distributionally robust Lyapunov function and controller learning). Consider the system (4) with nominal dynamics \mathbf{f} and perturbation \mathbf{W} . Let $\{\xi_i\}_{i \in [N]}$ be finitely many samples of the uncertainty ξ . Taking potential distribution shift into account, design a control policy $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ along with a Lyapunov certificate $V : \mathbb{R}^n \rightarrow \mathbb{R}$ that ensures that the equilibrium of the closed-loop system is globally asymptotically stable with high probability.

V. Convex Distributionally Robust Stability Formulation

In this section, we develop a distributionally robust formulation for control synthesis and analyze the stability of the resulting closed-loop systems. Our approach builds upon the concepts of distributionally robust optimization [23], [45].

A. Distributionally Robust Lyapunov-Stable Constraint

Let $\mathbb{P}_N := \frac{1}{N} \sum_{i=1}^N \delta_{\xi_i}$ denote the discrete empirical distribution obtained from the samples $\{\xi_i\}_{i \in [N]}$. We consider an ambiguity set of possible distributions for ξ , which are similar to the empirical distribution \mathbb{P}_N . 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 between two probability measures μ, ν in $\mathcal{P}_p(\Xi)$ is defined as [45]:

$$W_p(\mu, \nu) := \left(\inf_{\gamma \in \mathcal{Q}(\mu, \nu)} \left[\int_{\Xi \times \Xi} \|x - y\|^p d\gamma(x, y) \right] \right)^{\frac{1}{p}}, \quad (8)$$

where $\mathcal{Q}(\mu, \nu)$ is the set of measures on $\Xi \times \Xi$ with marginals μ and ν on the first and second factors. We define an ambiguity set $\mathcal{M}_N^r := \{\mu \in \mathcal{P}_p(\Xi) \mid W_p(\mu, \mathbb{P}_N) \leq r\}$ as a Wasserstein ball of distributions centered at \mathbb{P}_N with radius r .

Remark V.1 (Choice of Wasserstein ball radius). There is a relationship between the sample size N and the Wasserstein radius r for constructing the ambiguity set \mathcal{M}_N^r . A distribution \mathbb{P} is light-tailed if there exists an exponent ρ such that $C := \mathbb{E}_{\mathbb{P}}[\exp \|\xi\|^\rho] = \int_{\Xi} \exp \|\xi\|^\rho d\mathbb{P} < \infty$. If the true distribution \mathbb{P}^* of ξ is light-tailed, the choice of

radius $r = r_N(\bar{\epsilon})$ given in [23, Theorem 3.5],

$$r_N(\bar{\epsilon}) = \begin{cases} \left(\frac{\log(c_1 \bar{\epsilon}^{-1})}{c_2 N}\right)^{\frac{1}{\max\{k, 2\}}} & \text{if } N \geq \frac{\log(c_1 \bar{\epsilon}^{-1})}{c_2}, \\ \left(\frac{\log(c_1 \bar{\epsilon}^{-1})}{c_2 N}\right)^{\frac{1}{\rho}} & \text{else,} \end{cases} \quad (9)$$

where c_1 and c_2 are positive constants that depend on ρ, C and k (cf. [23, Theorem 3.4]), ensures that the ambiguity set $\mathcal{M}_N^{r_N(\bar{\epsilon})}$ contains \mathbb{P}^* with probability at least $1 - \bar{\epsilon}$. •

To account for the potential discrepancy between the empirical distribution \mathbb{P}_N and the true distribution \mathbb{P}^* of the system uncertainty ξ at run time, we aim to find a pair (V^*, π^*) that satisfies the following distributionally robust Lyapunov derivative constraint,

$$\inf_{\mathbb{P} \in \mathcal{M}_N^{r_N(\bar{\epsilon})}} \mathbb{P}(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) + \gamma \|\mathbf{x}\|) \leq 0) \geq 1 - \epsilon. \quad (10)$$

Compared with the chance-constrained formulation in (5), the distributionally robust chance-constrained formulation (10) requires only a finite set of samples instead of the true distribution \mathbb{P}^* . Moreover, it offers robust constraint satisfaction guarantees against potential shifts of the uncertainty distribution within the constructed ambiguity set.

B. Distributionally Robust Stability Characterization

Next, we characterize the stability properties of the closed-loop system governed by a controller satisfying (10). If a pair (V^*, π^*) satisfies (10), the following result ensures that the closed-loop system satisfies a chance constraint under the true distribution.

Lemma V.2 (Chance-constraint satisfaction under the true distribution). *Assume the distribution \mathbb{P}^* of ξ in (4) is light-tailed and the Wasserstein radius $r_N(\bar{\epsilon})$ is set according to (9). If the controller $\pi^*(\mathbf{x})$ and Lyapunov function $V^*(\mathbf{x})$ pair satisfies (10) with $r = r_N(\bar{\epsilon})$, then,*

$$\mathbb{P}^*(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) + \gamma \|\mathbf{x}\|) \leq 0) \geq (1 - \epsilon)(1 - \bar{\epsilon}). \quad (11)$$

Proof:

Let $A := \{\mathbb{P}^* \in \mathcal{M}_N^{r_N(\bar{\epsilon})}\}$ be the event that the true distribution \mathbb{P}^* lies within the ambiguity set $\mathcal{M}_N^{r_N(\bar{\epsilon})}$, and let $B := \{\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) + \gamma \|\mathbf{x}\|) \leq 0\}$ be the event that the Lyapunov derivative condition holds. From [23, Theorem 3.4], we have $\mathbb{P}^*(A) \geq 1 - \bar{\epsilon}$. From (10), we have that

$$\inf_{\mathbb{P} \in \mathcal{M}_N^{r_N(\bar{\epsilon})}} \mathbb{P}(B) \geq 1 - \epsilon.$$

Now, consider the probability of the event B under the true distribution \mathbb{P}^* :

$$\begin{aligned} \mathbb{P}^*(B) &\geq \mathbb{P}^*(B \cap A) = \mathbb{P}^*(B|A)\mathbb{P}^*(A) \\ &\geq \left(\inf_{\mathbb{P} \in \mathcal{M}_N^{r_N(\bar{\epsilon})}} \mathbb{P}(B) \right) \mathbb{P}^*(A) \geq (1 - \epsilon)(1 - \bar{\epsilon}) \end{aligned}$$

According to Lemma V.2, the closed-loop system satisfies the Lyapunov stability conditions pointwise in the state

space with high probability. However, to analyze the global stability of the closed-loop system, it is necessary to extend this guarantee to the trajectories of the system over time. To this end, the following result links pointwise Lyapunov stability to asymptotic stability in probability.

Lemma V.3 (Global asymptotic stability in probability). *Let $V^* : \mathbb{R}^n \rightarrow \mathbb{R}$ be a positive definite function with $V^*(\mathbf{0}_n) = 0$, and let V^* and the controller $\pi^* : \mathbb{R}^n \mapsto \mathbb{R}^m$ be a pair satisfying (11). Suppose $\xi \sim \mathbb{P}^*$, then, the origin $\mathbf{0}_n$ of the closed-loop system $\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}, \pi^*(\mathbf{x}), \xi)$ is globally asymptotically stable with probability at least $(1 - \epsilon)(1 - \bar{\epsilon})$.*

Proof:

Define the sets

$$\mathcal{A} := \{\xi \mid \sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) + \gamma \|\mathbf{x}\|) \leq 0\},$$

$$\mathcal{B} := \{\xi \mid \text{the system is asymptotically stable at } \mathbf{x} = \mathbf{0}_n\}.$$

For $\xi \in \mathcal{A}$, we have $\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) \leq -\gamma \|\mathbf{x}\|$, implying V^* decreases along the system trajectories, and leading to convergence to the equilibrium state $\mathbf{x} = \mathbf{0}$. Therefore, $\mathcal{A} \subseteq \mathcal{B}$, and we have $\mathbb{P}^*(\mathcal{B}) \geq \mathbb{P}^*(\mathcal{A}) \geq (1 - \epsilon)(1 - \bar{\epsilon})$, which concludes the result. ■

Lemmas V.2 and V.3 establish global asymptotic stability in probability of the closed-loop system. Lemma V.2 demonstrates that if a pair (V^*, π^*) satisfies the distributionally robust derivative constraint (10), then the system fulfills the Lyapunov stability condition with high probability under the true distribution \mathbb{P}^* . Subsequently, Lemma V.3 formalizes the global asymptotic stability of the system. Therefore, the satisfaction of the constraint in (10) ensures global asymptotic stability in probability.

Remark V.4 (Exponential stability under additional conditions). Under additional conditions, Lemma V.3 can ensure global exponential stability in probability. Consider the closed-loop system $\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}, \pi^*(\mathbf{x}), \xi)$ with equilibrium at $\mathbf{0}_n$. Suppose there exists a positive constant α such that the pair (V^*, π^*) satisfies

$$\begin{aligned} \inf_{\mathbb{P} \in \mathcal{M}_N^{r_N(\bar{\epsilon})}} \mathbb{P}(\sup_{\mathbf{x} \in \mathcal{X}} (\dot{V}^*(\mathbf{x}, \pi^*(\mathbf{x}), \xi) + \gamma \|\mathbf{x}\| + \alpha V^*(\mathbf{x})) < 0) \\ \geq 1 - \epsilon, \end{aligned} \quad (12)$$

and V^* satisfies

$$\alpha_1 \|\mathbf{x}\|^p \leq V^*(\mathbf{x}) \leq \alpha_2 \|\mathbf{x}\|^p, \quad (13)$$

for some constants $\alpha_1, \alpha_2, p > 0$. Similar to the reasoning presented in Lemma V.2 and V.3, we conclude that the equilibrium of the closed-loop system is globally exponentially stable [55] with probability at least $(1 - \epsilon)(1 - \bar{\epsilon})$. •

Remark V.5 (Connections to other probabilistic stability notions). In the literature, there exist other notions of stability in probability, especially for stochastic discrete-time systems, see e.g., [59]–[61]. For the stochastic discrete-time system:

$$\mathbf{x}_{i+1} = \hat{\mathbf{f}}(\mathbf{x}_i, \mathbf{u}_i, \xi_i), \quad (14)$$

a popular stability notion refers to the probability that the state remains in a bounded region [61, Definition 6]. Formally, the system (14) is bounded in probability for some $K \in \mathbb{N}$ if there exists $M > 0$ and $\epsilon \in (0, 1)$ such that

$$\mathbb{P}^* \left\{ \max_{k \leq K} \|\mathbf{x}_k\| \leq M \right\} \geq 1 - \epsilon. \quad (15)$$

This notion is widely used for ensuring finite-time stability and safety [62], [63] in stochastic systems. In typical stochastic systems, the parameter $\boldsymbol{\xi}$ varies as time progresses, which is a more general and challenging setting. In contrast, as outlined in (4), our work focuses on scenarios where $\boldsymbol{\xi}$ remains fixed over time, allowing us to derive stronger convergence results. Lemma V.3 not only certifies stability in probability over finite-time horizons but also extends to global asymptotic stability guarantees in probability. •

C. Reformulation of Distributionally Robust Stability Constraint

Synthesizing a pair (V^*, π^*) that satisfies (10) is challenging because the constraint encompasses an infimum over a set of probability measures and a supremum over the state space \mathcal{X} . We leverage recent advances in distributionally robust optimization [23], [24] to identify a sufficient condition.

To simplify the notation, we define $h : \mathbb{R}^k \rightarrow \mathbb{R}$:

$$h(\boldsymbol{\xi}) := \sup_{\mathbf{x} \in \mathcal{X}} \left(\dot{V}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}), \boldsymbol{\xi}) + \gamma \|\mathbf{x}\| \right). \quad (16)$$

Note that by (4), we have

$$\dot{V}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}), \boldsymbol{\xi}) = \nabla V(\mathbf{x})^\top (\mathbf{f}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x})) + \mathbf{W}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\boldsymbol{\xi}). \quad (17)$$

As it is the supremum of affine functions of $\boldsymbol{\xi}$, $h(\boldsymbol{\xi})$ is convex in $\boldsymbol{\xi}$. Furthermore, as \mathcal{X} is compact, V is continuously differentiable, and \mathbf{f} and \mathbf{W} are locally Lipschitz, the supremum in (16) exists and is finite for each $\boldsymbol{\xi}$.

We present a sufficient condition for the distributionally robust Lyapunov derivative constraint in (10) that relies on samples of $\boldsymbol{\xi}$. This condition provides a tractable reformulation of the constraint, making it more amenable to finding and learning the pair (V^*, π^*) .

Proposition V.6 (Distributionally Robust Stability Condition). *Let $\{\boldsymbol{\xi}_i\}_{i \in [N]}$ be a set of samples of the uncertainty parameter $\boldsymbol{\xi}$, and let h be the function defined in (16). Assume that the samples are indexed such that $h(\boldsymbol{\xi}_i) \geq h(\boldsymbol{\xi}_k)$, for all $1 \leq i < k \leq N$. For $\epsilon \in (0, 1)$, let $j \in [N]$ be the unique index that satisfies $\frac{j-1}{N} - \epsilon < 0$ and $\frac{j}{N} - \epsilon \geq 0$. Then,*

$$\begin{aligned} & \frac{r}{\epsilon} \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{W}^\top(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\nabla V(\mathbf{x})\| + \\ & \frac{1}{N\epsilon} \sum_{i=1}^{j-1} (h(\boldsymbol{\xi}_i) - h(\boldsymbol{\xi}_j)) + h(\boldsymbol{\xi}_j) \leq 0 \end{aligned} \quad (18)$$

is a sufficient condition for the distributionally robust stability constraint (10) to hold for the ambiguity set corresponding to the 1-Wasserstein distance. Furthermore, if $\epsilon \leq \frac{1}{N}$, the condition (18) simplifies to:

$$\frac{r}{\epsilon} \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{W}^\top(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\nabla V(\mathbf{x})\| + \max_i h(\boldsymbol{\xi}_i) \leq 0. \quad (19)$$

Proof:

Similar to the CVaR approximation (6) of the chance constraint (5), the following is a sufficient condition for (10):

$$\sup_{\mathbb{P} \in \mathcal{M}_N^{\mathbb{P}}} \text{CVaR}_{1-\epsilon}^{\mathbb{P}}(h(\boldsymbol{\xi})) \leq 0, \quad (20)$$

Using (7), we write (20) equivalently as,

$$\sup_{\mathbb{P} \in \mathcal{M}_N^{\mathbb{P}}} \inf_{t \in \mathbb{R}} [\mathbb{E}_{\mathbb{P}}[(h(\boldsymbol{\xi}) + t)_+] - t\epsilon] \leq 0. \quad (21)$$

Based on [24, Lemma V.8] and [23, Theorem 6.3], with the 1-Wasserstein distance, the supremum over the ambiguity set can be written conservatively as the sample average over the empirical distribution \mathbb{P}_N and a regularization term:

$$rL_h + \inf_{t \in \mathbb{R}} [\mathbb{E}_{\mathbb{P}_N}[(h(\boldsymbol{\xi}) + t)_+] - t\epsilon] \leq 0, \quad (22)$$

where L_h is the Lipschitz constant of $h(\boldsymbol{\xi})$ in $\boldsymbol{\xi}$. Based on [64, Proposition 1.32], the Lipschitz constant of the supremum of a family of Lipschitz functions is bounded by the supremum of their Lipschitz constants. As shown in (17), $\dot{V}(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}), \boldsymbol{\xi})$ is affine in $\boldsymbol{\xi}$ for each \mathbf{x} and its Lipschitz constant is given by $\|\mathbf{W}^\top(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\nabla V(\mathbf{x})\|$. Therefore,

$$L_h = \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{W}^\top(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\nabla V(\mathbf{x})\|. \quad (23)$$

Next, we write the second term in (22) as:

$$\inf_{t \in \mathbb{R}} [\mathbb{E}_{\mathbb{P}_N} [h(\boldsymbol{\xi}) + t)_+] - t\epsilon = \inf_{t \in \mathbb{R}} \left[\frac{1}{N} \sum_{i=1}^N (h(\boldsymbol{\xi}_i) + t)_+ - t\epsilon \right]. \quad (24)$$

By assumption, $h(\boldsymbol{\xi}_i) \geq h(\boldsymbol{\xi}_k)$, for all $1 \leq i < k \leq N$. Observe that the function $\frac{1}{N} \sum_{i=1}^N (h(\boldsymbol{\xi}_i) + t)_+ - t\epsilon$ is piecewise-linear in t with $N+1$ intervals and N breakpoints, given by $\{-h(\boldsymbol{\xi}_i)\}_{i \in [N]}$ and the slope for the i -th interval is $\frac{i-1}{N} - \epsilon$. Thus, the optimal solution is $t^* = -h(\boldsymbol{\xi}_j)$, where $j \in [N]$ satisfies $\frac{j-1}{N} - \epsilon < 0$ and $\frac{j}{N} - \epsilon \geq 0$. Substituting the optimal solution t^* into (24), we get:

$$\frac{1}{N} \sum_{i=1}^{j-1} (h(\boldsymbol{\xi}_i) - h(\boldsymbol{\xi}_j)) + \epsilon h(\boldsymbol{\xi}_j). \quad (25)$$

Therefore, we know that:

$$\begin{aligned} & r \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{W}^\top(\mathbf{x}, \boldsymbol{\pi}(\mathbf{x}))\nabla V(\mathbf{x})\| + \\ & \frac{1}{N} \sum_{i=1}^{j-1} (h(\boldsymbol{\xi}_i) - h(\boldsymbol{\xi}_j)) + \epsilon h(\boldsymbol{\xi}_j) \leq 0 \end{aligned} \quad (26)$$

is equivalent to (22), and thus sufficient for (10).

Furthermore, if $\epsilon \leq \frac{1}{N}$, only the first interval has negative slope and (25) can be written equivalently as $\epsilon h(\boldsymbol{\xi}_1) = \max_i \epsilon h(\boldsymbol{\xi}_i)$. Thus, given that $\epsilon \leq \frac{1}{N}$, we conclude that (19) is sufficient for (10). ■

Based on Proposition V.6, our goal is to find a controller $\pi^*(\mathbf{x})$ and Lyapunov function $V^*(\mathbf{x})$ pair that fulfills condition (18). Once such a pair is determined, leveraging Lemmas V.2 and V.3, we can ascertain with high confidence that the system (4), governed by $\pi^*(\mathbf{x})$, exhibits (exponential) asymptotic stability of the equilibrium.

VI. Distributionally Robust Lyapunov-Stable Controller Learning

In this section, we present a learning-based approach for finding a distributionally robust controller $\pi^*(\mathbf{x})$ and Lyapunov function $V^*(\mathbf{x})$ pair satisfying (18). We make the following assumptions.

Assumption 1 (Lipschitz Continuity and Boundedness). Assume the nominal system dynamics $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and the perturbation function $\mathbf{W} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^{n \times k}$ are Lipschitz on $\mathcal{X} \times \mathcal{U}$ with constants L_f and L_W , respectively, i.e., for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and $\mathbf{u}, \mathbf{u}' \in \mathcal{U}$,

$$\begin{aligned} \|\mathbf{f}(\mathbf{x}, \mathbf{u}) - \mathbf{f}(\mathbf{x}', \mathbf{u}')\| &\leq L_f(\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{u} - \mathbf{u}'\|), \\ \|\mathbf{W}(\mathbf{x}, \mathbf{u}) - \mathbf{W}(\mathbf{x}', \mathbf{u}')\| &\leq L_W(\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{u} - \mathbf{u}'\|). \end{aligned}$$

Lipschitzness together with the compactness of the domain implies that both \mathbf{f} and \mathbf{W} are uniformly bounded on $\mathcal{X} \times \mathcal{U}$. Specifically, there exist constants $B_f, B_W \in \mathbb{R}_{>0}$ such that, for any $\mathbf{x} \in \mathcal{X}$ and $\mathbf{u} \in \mathcal{U}$,

$$\|\mathbf{f}(\mathbf{x}, \mathbf{u})\| \leq B_f \quad \text{and} \quad \|\mathbf{W}(\mathbf{x}, \mathbf{u})\| \leq B_W.$$

The Lipschitz continuity assumption is particularly useful in our setting, as the neural networks are trained with discrete data, and we aim to certify stability for the entire state space. Lipschitzness ensures that the system dynamics do not change drastically for nearby states, which is crucial for the stability proofs presented later in this work.

A. Model Architecture and Loss Function

To parameterize the desired pair $(V^*(\mathbf{x}), \pi^*(\mathbf{x}))$, we consider the following neural network representations. For the Lyapunov function, termed NN-LF, we set

$$V_{\theta_1}(\mathbf{x}) := \|\phi_{\theta_1}(\mathbf{x}) - \phi_{\theta_1}(\mathbf{0}_n)\|^2 + \hat{\alpha}\|\mathbf{x}\|^2, \quad (27)$$

where $\phi_{\theta_1} : \mathbb{R}^n \mapsto \mathbb{R}$ is a fully-connected neural network with parameters θ_1 and tanh activations, and $\hat{\alpha}$ is a user-chosen parameter [11], [12]. By construction, this function is positive definite and $V_{\theta_1}(\mathbf{0}_n) = 0$. Furthermore, the use of tanh activations ensures continuous differentiability and bounded derivatives. Therefore, the gradient $\nabla V_{\theta_1}(\mathbf{x})$ is Lipschitz with some constant $L_{\nabla V}$ and its norm $\|\nabla V_{\theta_1}(\mathbf{x})\|$ is bounded within the compact set \mathcal{X} by some constant B_V . As shown in [65], these constants can be estimated explicitly, but we will only rely on their existence for stability analysis.

For the controller, we set

$$\pi_{\theta_2}(\mathbf{x}) := \varphi_{\theta_2}(\mathbf{x}) - \varphi_{\theta_2}(\mathbf{0}_n), \quad (28)$$

where φ_{θ_2} is a neural network with parameters θ_2 and tanh activations. By construction, this formulation guarantees $\pi_{\theta_2}(\mathbf{0}_n) = \mathbf{0}_m$ and ensures that $\pi_{\theta_2}(\mathbf{x})$ is locally Lipschitz in \mathcal{X} , with constant L_π .

When employing a neural-network Lyapunov function, the Lyapunov derivative condition in (3b) may not hold within a small neighborhood of the equilibrium due to numerical inaccuracies, see e.g., [9], [12]. These inaccuracies stem from a neural network's inability to precisely approximate the derivative's behavior near the equilibrium, where its values

are close to zero. This observation leads us to focus on establishing the validity of learned controller and Lyapunov function pairs outside of a neighborhood close to the origin. Therefore, our analysis concentrates on the region $\mathcal{X}_\delta := \mathcal{X} \setminus \overline{B}(\mathbf{0}_n; \delta)$, which excludes a δ -radius ball around the equilibrium.

Definition VI.1 (δ -accurate Lyapunov function over \mathcal{X} [9]). A Lyapunov function $V(\mathbf{x})$ for a dynamical system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \pi(\mathbf{x}))$ is δ -accurate if it is positive definite, $V(\mathbf{0}_n) = 0$, and satisfies the Lyapunov derivative conditions (3b) everywhere in \mathcal{X}_δ .

The existence of a δ -accurate Lyapunov function implies the ultimate boundedness of the closed-loop system trajectories within a closed ball $\overline{B}(\mathbf{0}_n; \delta)$. For a neural network representation, $\delta > 0$ can be chosen arbitrarily small, potentially requiring a larger network size for V_{θ_1} and π_{θ_2} and greater sampling density for the training set. For our distributionally robust formulation, we introduce a distributionally robust δ -accurate Lyapunov function as follows.

Definition VI.2 (Distributionally robust δ -accurate Lyapunov function with ϵ margin of error over \mathcal{X}). A Lyapunov function $V(\mathbf{x})$ for a dynamical system $\dot{\mathbf{x}} = \bar{\mathbf{f}}(\mathbf{x}, \pi(\mathbf{x}), \boldsymbol{\xi})$ in (4) is distributionally robust δ -accurate if it is positive definite, $V(\mathbf{0}_n) = 0$, and satisfies:

$$\inf_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(\sup_{\mathbf{x} \in \mathcal{X}_\delta} (\dot{V}(\mathbf{x}, \pi(\mathbf{x}), \boldsymbol{\xi}) + \gamma\|\mathbf{x}\|) \leq 0) \geq 1 - \epsilon.$$

Arguments analogous to the ones we employed in the proofs of Lemmas V.2 and V.3 show that the existence of a distributionally robust δ -accurate Lyapunov function implies the stability of the closed-loop system to the closed ball $\overline{B}(\mathbf{0}_n; \delta)$ with high probability.

To facilitate the training process of the pair $(V_{\theta_1}, \pi_{\theta_2})$, a training set $\mathcal{D}_{\text{LF}} := \{\mathbf{x}_i\}_{i \in [M]}$ is constructed by sampling $\mathcal{D}_{\mathcal{X}} := \{\mathbf{x}_i\}_{i \in [M]}$ uniformly from \mathcal{X}_δ . We also assume the availability of the uncertainty training set $\mathcal{D}_\xi := \{\boldsymbol{\xi}_i\}_{i \in [N]}$, collected by offline measurements of the system.

For a system without any uncertainty, we minimize the following empirical loss function that encourages the satisfaction of the Lyapunov derivative constraint (3b):

$$\ell_{\text{Nominal}}(\boldsymbol{\theta}) = \frac{1}{M} \sum_{i=1}^M (\dot{V}_{\theta_1}(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i)) + \gamma\|\mathbf{x}_i\|)_+, \quad (29)$$

where $\boldsymbol{\theta} = [\theta_1, \theta_2]$, and $\gamma \in \mathbb{R}_{>0}$ is a user-chosen parameter.

To address the model uncertainty in (4), we aim to find a distributionally robust (DR)-NN LF V_{θ_1} and a controller π_{θ_2} . We focus on the case where $\epsilon < \frac{1}{N}$ for two reasons. First, the DRO framework is particularly useful when the available data is limited, as it does not require a large number of samples to provide performance guarantees. In fact, the Wasserstein DRO approach is motivated by the fact that the true distribution is often only indirectly observable through a finite training dataset [23]. When the number of samples N is relatively small, the condition $\epsilon < \frac{1}{N}$ is also easy to

satisfy, making it a practical choice for scenarios with limited data. Second, setting $\epsilon < \frac{1}{N}$ ensures a higher probability of stability guarantees, as it requires the Lyapunov derivative condition to hold for all uncertainty samples in the training set \mathcal{D}_ξ . Based on the stability condition in (19), we consider the following loss function for learning a DR-NN LF and controller:

$$\ell_{\text{DR}}(\theta) = \left(\frac{r}{\epsilon} \max_{\mathbf{x}_i \in \mathcal{D}_{\text{LF}}} \|\mathbf{W}^\top(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i)) \nabla V_{\theta_1}(\mathbf{x}_i)\| + \max_j \left(\sum_{i=1}^M (\dot{V}_{\theta_1}(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i), \xi_j)) + \gamma \|\mathbf{x}_i\| \right) \right)_+ \quad (30)$$

In the following, for a sufficiently large M , we aim to show that the minimizer of the loss function (30) yields a pair $(V_{\theta_1}^*, \pi_{\theta_2}^*)$ satisfying the condition in (10).

B. Nominal Neural Lyapunov-Stable Control

First, we show based on results from [12] that for a sufficiently large M , the nominal stabilizing controller learned from (29) stabilizes the nominal system. We formalize the necessary sampling density required to ensure stability using the learned LF.

Lemma VI.3. ([12, Lemma 4]). *For any $\delta > 0$ and a chosen parameter $c > 0$, there exists a minimum number of samples $M(\delta, c) \in \mathbb{N}$ such that, for all $M \geq M(\delta, c)$, a uniformly sampled training dataset $\mathcal{D}_{\text{LF}} = \{\mathbf{x}_i\}_{i \in [M]} \subset \mathcal{X}_\delta$ ensures that the domain \mathcal{X}_δ can be covered by the union of balls $B(\mathbf{x}_i; c\|\mathbf{x}_i\|)$. Specifically, this guarantees that for any $\mathbf{x} \in \mathcal{X}_\delta$, there exists $\mathbf{x}_i \in \mathcal{D}_{\text{LF}}$ satisfying $\|\mathbf{x} - \mathbf{x}_i\| \leq c\|\mathbf{x}_i\|$.*

This result underscores the relation between δ , c , and the sampling density M , crucial for approximating the controller and Lyapunov function effectively.

Lemma VI.4 (Neural Lyapunov-stable control). *Let $\mathcal{D}_{\text{LF}} = \{\mathbf{x}_i\}_{i \in [M]} \subset \mathcal{X}_\delta$ be a uniformly sampled training set, with $M \geq M(\delta, c)$ as defined in Lemma VI.3. Let $\theta^* = (\theta_1^*, \theta_2^*)$ be the trained parameters such that $\ell_{\text{Nominal}}(\theta^*) = 0$. Denote by $L_{\nabla V}$ and L_π the Lipschitz constants of $\nabla V_{\theta_1}^*$ and $\pi_{\theta_2}^*$, respectively, and let B_V be the bound on $\|\nabla V_{\theta_1}^*(\mathbf{x})\|$ for $\mathbf{x} \in \mathcal{X}$. Let L_f and B_f be as defined in Assumption 1. If $c > 0$ is sufficiently small to ensure $\gamma - (L_f(L_\pi + 1)B_V + L_{\nabla V}B_f)c > 0$, then the controller $\pi_{\theta_2}^*$ stabilizes the nominal system \mathbf{f} to the closed ball $\bar{B}(\mathbf{0}_n; \delta)$ as certified by $V_{\theta_1}^*$, which is a δ -accurate Lyapunov function for the controlled system over \mathcal{X} .*

Proof:

Given that $\ell_{\text{Nominal}}(\theta^*) = 0$, we have

$$\dot{V}_{\theta_1}^*(\mathbf{x}_i, \pi_{\theta_2}^*(\mathbf{x}_i)) + \gamma \|\mathbf{x}_i\| \leq 0, \quad (31)$$

for all $\mathbf{x}_i \in \mathcal{D}_{\text{LF}}$. Since $\nabla V_{\theta_1}^*(\mathbf{x})$ and $\pi_{\theta_2}^*$ are Lipschitz, based on the Assumption 1, we have for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_\delta$,

$$\begin{aligned} \|\nabla V_{\theta_1}^*(\mathbf{x}) - \nabla V_{\theta_1}^*(\mathbf{y})\| &\leq L_{\nabla V} \|\mathbf{x} - \mathbf{y}\|, \quad (32) \\ \|\mathbf{f}(\mathbf{x}, \pi_{\theta_2}^*(\mathbf{x})) - \mathbf{f}(\mathbf{y}, \pi_{\theta_2}^*(\mathbf{y}))\| &\leq L_f(L_\pi + 1) \|\mathbf{x} - \mathbf{y}\|. \end{aligned}$$

We denote $L_q := L_f(L_\pi + 1)$ and write $\mathbf{f}(\mathbf{x}, \pi_{\theta_2}^*(\mathbf{x}))$ as $\mathbf{f}(\mathbf{x})$ for brevity. Based on (32), we have for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_\delta$,

$$\begin{aligned} &|\dot{V}_{\theta_1}^*(\mathbf{x}, \pi_{\theta_2}^*(\mathbf{x})) - \dot{V}_{\theta_1}^*(\mathbf{y}, \pi_{\theta_2}^*(\mathbf{y}))| = \\ &|[\nabla V_{\theta_1}^*(\mathbf{x})]^\top \mathbf{f}(\mathbf{x}) - [\nabla V_{\theta_1}^*(\mathbf{y})]^\top \mathbf{f}(\mathbf{y})| \leq \\ &\|\nabla V_{\theta_1}^*(\mathbf{x})\| \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})\| + \|\nabla V_{\theta_1}^*(\mathbf{x}) - \nabla V_{\theta_1}^*(\mathbf{y})\| \|\mathbf{f}(\mathbf{y})\| \\ &\leq (L_q \|\nabla V_{\theta_1}^*(\mathbf{x})\| + L_{\nabla V} \|\mathbf{f}(\mathbf{y})\|) \|\mathbf{x} - \mathbf{y}\| \\ &\leq (L_q B_V + L_{\nabla V} B_f) \|\mathbf{x} - \mathbf{y}\|. \end{aligned} \quad (33)$$

This establishes that $L_{\dot{V}}$, the Lipschitz constant for the Lyapunov derivative, can be bounded by $L_q B_V + L_{\nabla V} B_f$. By Lemma VI.3, for any $\mathbf{x} \in \mathcal{X}_\delta$, there exists a nearby sampled point \mathbf{x}_i such that $\|\mathbf{x} - \mathbf{x}_i\| \leq c\|\mathbf{x}_i\|$. Using (31) and (33),

$$\begin{aligned} \dot{V}_{\theta_1}^*(\mathbf{x}, \pi_{\theta_2}^*(\mathbf{x})) &\leq \dot{V}_{\theta_1}^*(\mathbf{x}_i, \pi_{\theta_2}^*(\mathbf{x}_i)) + L_{\dot{V}} c \|\mathbf{x}_i\| \quad (34) \\ &\leq -\gamma \|\mathbf{x}_i\| + L_{\dot{V}} c \|\mathbf{x}_i\| \\ &\leq -(\gamma - L_{\dot{V}} c) \delta. \end{aligned}$$

Since $\gamma - L_{\dot{V}} c > 0$, this ensures that $\dot{V}_{\theta_1}^*(\mathbf{x}, \pi_{\theta_2}^*(\mathbf{x}))$ satisfies the Lyapunov derivative condition for all $\mathbf{x} \in \mathcal{X}_\delta$, and thus it is a δ -accurate Lyapunov function. Therefore, the learned controller provides δ -accurate stabilization of the nominal system \mathbf{f} as certified by the Lyapunov function. ■

The stability assurances provided by Lemma VI.4 may not hold in the presence of the uncertainties in (4), which motivates our study of the system under the distributionally robust controller learned from (30).

C. Neural Distributionally Robust Lyapunov-Stable Control

For a sufficiently large sampling density, we establish probabilistic δ -accurate stability guarantees for the uncertain system (4) governed by the controller learned from (30).

Proposition VI.5 (Distributionally robust neural Lyapunov-stable control). *Let $\mathcal{D}_{\text{LF}} = \{\mathbf{x}_i\}_{i \in [M]} \subset \mathcal{X}_\delta$ be a uniformly sampled training set, with $M \geq M(\delta, c)$ as defined in Lemma VI.3. Let $\theta^* = (\theta_1^*, \theta_2^*)$ be the trained parameters such that $\ell_{\text{DR}}(\theta^*) = 0$. Let $L_{\nabla V}$ and L_π denote the Lipschitz constants of $\nabla V_{\theta_1}^*$ and $\pi_{\theta_2}^*$, respectively, and let B_V be the bound on $\|\nabla V_{\theta_1}^*(\mathbf{x})\|$ for $\mathbf{x} \in \mathcal{X}$ and $\epsilon \leq \frac{1}{N}$. Let $c > 0$ be sufficiently small such that $\gamma - ((L_W(L_\pi + 1))B_V + L_{\nabla V}B_W)B_\xi \frac{r}{\epsilon} c - (L_f(L_\pi + 1)B_V + L_{\nabla V}B_f)c > 0$, where L_W , L_f , B_W , and B_f are as defined in Assumption 1, and $B_\xi \in \mathbb{R}_{>0}$ is such that $\|\xi\| \leq B_\xi$ for all $\xi \in \Xi$. Then, the controller $\pi_{\theta_2}^*$ stabilizes the uncertain system (4) to the closed ball $\bar{B}(\mathbf{0}_n; \delta)$ with high probability, as certified by $V_{\theta_1}^*$, which is a distributionally robust δ -accurate Lyapunov function for the controlled system over \mathcal{X} .*

Proof:

Given that $\ell_{\text{DR}}(\theta^*) = 0$, we have for all $\mathbf{x}_i \in \mathcal{D}_{\text{LF}}$,

$$\begin{aligned} &\frac{r}{\epsilon} \max_{\mathbf{x}_i \in \mathcal{D}_{\text{LF}}} \|\mathbf{W}^\top(\mathbf{x}_i, \pi_{\theta_2}^*(\mathbf{x}_i)) \nabla V_{\theta_1}^*(\mathbf{x}_i)\| + \\ &\max_j (\dot{V}_{\theta_1}^*(\mathbf{x}_i, \pi_{\theta_2}^*(\mathbf{x}_i), \xi_j)) + \gamma \|\mathbf{x}_i\| \leq 0. \end{aligned}$$

Based on Assumption 1 and the Lipschitz constant L_π of the controller $\pi_{\theta_2^*}$, we have for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_\delta$:

$$\|\mathbf{W}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x})) - \mathbf{W}(\mathbf{y}, \pi_{\theta_2^*}(\mathbf{x}))\| \leq L_W(L_\pi + 1)\|\mathbf{x} - \mathbf{y}\|. \quad (35)$$

For brevity, we denote $L_p = L_W(L_\pi + 1)$ and write $\mathbf{W}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}))$ and $V_{\theta_1^*}(\mathbf{x})$ as $\mathbf{W}(\mathbf{x})$ and $V(\mathbf{x})$, respectively. Using (32) and (35), we have for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_\delta$:

$$\begin{aligned} & \|\mathbf{W}^\top(\mathbf{x})\nabla V(\mathbf{x}) - \mathbf{W}^\top(\mathbf{y})\nabla V(\mathbf{y})\| \leq \\ & \|(\mathbf{W}^\top(\mathbf{x}) - \mathbf{W}^\top(\mathbf{y}))\nabla V(\mathbf{x})\| + \|\mathbf{W}^\top(\mathbf{y})(\nabla V(\mathbf{x}) - \nabla V(\mathbf{y}))\| \\ & \leq (L_p\|\nabla V(\mathbf{x})\| + L_{\nabla V}\|\mathbf{W}^\top(\mathbf{y})\|)\|\mathbf{x} - \mathbf{y}\| \\ & \leq (L_p B_V + L_{\nabla V} B_W)\|\mathbf{x} - \mathbf{y}\|. \end{aligned} \quad (36)$$

Thus, the term $\mathbf{W}^\top(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}))\nabla V_{\theta_1^*}(\mathbf{x})$ is Lipschitz with constant $L_{\dot{W}}$, bounded by $L_p B_V + L_{\nabla V} B_W$.

Moreover, as $\xi \in \Xi$ and Ξ is compact, we know that $\|\xi\| \leq B_\xi$ for some constant $B_\xi \in \mathbb{R}_{>0}$. Using (33) and (35), for the term $\dot{V}_{\theta_1^*}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}), \xi)$, we have that, for each ξ , and for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_\delta$,

$$\begin{aligned} & |\dot{V}_{\theta_1^*}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}), \xi) - \dot{V}_{\theta_1^*}(\mathbf{y}, \pi_{\theta_2^*}(\mathbf{y}), \xi)| \\ & = |[\nabla V_{\theta_1^*}(\mathbf{x})]^\top \mathbf{f}(\mathbf{x}) - [\nabla V_{\theta_1^*}(\mathbf{y})]^\top \mathbf{f}(\mathbf{y}) \\ & \quad + ([\nabla V_{\theta_1^*}(\mathbf{x})]^\top \mathbf{W}(\mathbf{x}) - [\nabla V_{\theta_1^*}(\mathbf{y})]^\top \mathbf{W}(\mathbf{y}))\xi| \\ & \leq L_{\dot{V}}\|\mathbf{x} - \mathbf{y}\| + L_{\dot{W}}\|\xi\|\|\mathbf{x} - \mathbf{y}\| \\ & \leq (L_{\dot{V}} + L_{\dot{W}}B_\xi)\|\mathbf{x} - \mathbf{y}\|. \end{aligned} \quad (37)$$

We denote $L_{\max} := L_{\dot{V}} + L_{\dot{W}}B_\xi$. We define $G(\mathbf{x}) = \max_j \dot{V}_{\theta_1^*}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}), \xi_j)$. From (37), we deduce $|G(\mathbf{x}) - G(\mathbf{y})| \leq L_{\max}\|\mathbf{x} - \mathbf{y}\|$, and thus G is also Lipschitz with constant L_{\max} .

Lastly, following the covering argument in Lemma VI.3, we have that for any $\mathbf{x} \in \mathcal{X}_\delta$, there exists a $\mathbf{x}_i \in \mathcal{D}_{\text{LF}}$ such that $\|\mathbf{x} - \mathbf{x}_i\| \leq c\|\mathbf{x}_i\|$. Therefore, $\forall \mathbf{x} \in \mathcal{X}_\delta$,

$$\begin{aligned} & \frac{r}{\epsilon} \sup_{\mathbf{x} \in \mathcal{X}_\delta} \|\mathbf{W}^\top(\mathbf{x})\nabla V(\mathbf{x})\| + \max_j (\dot{V}_{\theta_1^*}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}), \xi_j)) \leq \\ & \frac{r}{\epsilon} \max_{\mathbf{x}_i \in \mathcal{D}_{\text{LF}}} \|\mathbf{W}^\top(\mathbf{x}_i)\nabla V(\mathbf{x}_i)\| + \max_j (\dot{V}_{\theta_1^*}(\mathbf{x}_i, \pi_{\theta_2^*}(\mathbf{x}_i), \xi_j)) + \\ & \left(\frac{r}{\epsilon} L_{\dot{W}}c + L_{\max}c\right)\|\mathbf{x}_i\| \leq -(\gamma - \frac{r}{\epsilon} L_{\dot{W}}c - L_{\max}c)\delta. \end{aligned} \quad (38)$$

Note that $\gamma - \frac{r}{\epsilon} L_{\dot{W}}c - L_{\max}c > 0$. Thus, we have that $\frac{r}{\epsilon} \sup_{\mathbf{x} \in \mathcal{X}_\delta} \|\mathbf{W}^\top(\mathbf{x})\nabla V(\mathbf{x})\| + \max_j (\dot{V}_{\theta_1^*}(\mathbf{x}, \pi_{\theta_2^*}(\mathbf{x}), \xi_j)) < 0$ holds for all $\mathbf{x} \in \mathcal{X}_\delta$.

Therefore, reasoning as in Proposition V.6, we conclude that, for any $\xi \sim \mathbb{P}^*$, (10) is satisfied for all $\mathbf{x} \in \mathcal{X}_\delta$. This implies that $V_{\theta_1^*}$ is a distributionally robust δ -accurate Lyapunov function. Therefore, the learned controller provides distributionally robust δ -accurate stabilization of the uncertain system, as certified by the Lyapunov function. ■

The universal approximation property of neural networks [66], [67] suggests that, given sufficient network width and depth, the optimal Lyapunov function and controller can be approximated to arbitrary accuracy, making it possible to achieve zero training loss in the nominal (29) and distributionally robust (30) settings.

The learning process for the DR-NN-LF and controller pair, as defined in (30), can be made more efficient by

initializing the neural network weights with those obtained from the nominal training (29). This initialization strategy significantly reduces the required training time and computational resources for the DR formulation.

Remark VI.6 (Neural verification of stability guarantees).

While the universal approximation property of neural networks suggests that the optimal Lyapunov function and controller can be approximated to arbitrary accuracy, achieving zero training loss may not always be possible in practice. In such cases, the stability guarantees provided by Lemma VI.4 and Proposition VI.5 may not hold exactly. However, similar to [68], by employing neural network verification techniques, such as alpha-beta-CROWN [69] and Marabou [70], one can formally verify whether the learned Lyapunov function and controller pair satisfy the required derivative conditions over a compact state set. •

Remark VI.7 (Pointwise distributionally robust formulation).

The formulation in Proposition V.6 and the corresponding loss (30) may be conservative and challenging to train. The conservativeness arises from the term $\frac{r}{\epsilon} \max_{\mathbf{x}_i \in \mathcal{D}_{\text{LF}}} \|\mathbf{W}^\top(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i))\nabla V_{\theta_1}(\mathbf{x}_i)\|$, which is a positive constant at any \mathbf{x}_i in the training set \mathcal{D}_{LF} . This term requires the Lyapunov derivative condition to be the negative of that constant for all training data, which can be a stringent requirement, especially for states near the origin.

An alternative approach is to consider a pointwise distributionally robust formulation, which enforces the chance constraint at each state $\mathbf{x} \in \mathcal{X}$ individually:

$$\mathbb{P}^*(\dot{V}(\mathbf{x}, \pi(\mathbf{x}), \xi) + \gamma\|\mathbf{x}\| \leq 0) \geq 1 - \epsilon, \quad \forall \mathbf{x} \in \mathcal{X}. \quad (39)$$

Based on this pointwise formulation, we can define a pointwise loss function for training the neural network Lyapunov function and controller:

$$\begin{aligned} \ell_{\text{DR}}(\theta) = & \frac{1}{M} \sum_{i=1}^M \left(\frac{r}{\epsilon} \|\mathbf{W}^\top(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i))\nabla V_{\theta_1}(\mathbf{x}_i)\| + \right. \\ & \left. \max_j (\dot{V}_{\theta_1}(\mathbf{x}_i, \pi_{\theta_2}(\mathbf{x}_i), \xi_j)) + \gamma\|\mathbf{x}_i\| \right)_+. \end{aligned} \quad (40)$$

This pointwise loss function is more flexible and easier to train compared to the uniform loss (30) because it allows for more variability in the Lyapunov derivative values across different states. By enforcing the distributionally robust chance constraint individually for each state, the pointwise formulation allows the learned Lyapunov function and controller to adapt to the specific characteristics of each state, potentially leading to faster training convergence.

This pointwise formulation would ensure that the Lyapunov function decreases with high probability at each individual state. However, providing a global stability guarantee based on this requires notions of stochastic stability [60], [61]. We leave the analysis of the stability properties of the closed-loop system under the controller learned using the pointwise loss function (40) for future work. •

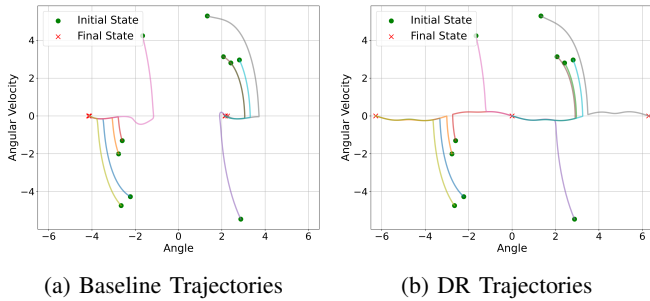


FIGURE 2: Comparison of trajectories for the inverted pendulum system with test case parameters: mass $m = 1.1$, length $l = 1.0$, and damping $b = 0.18$. The 10 random sampled initial states are marked as green dots, while the final states are marked as red crosses. The states $(\theta, \dot{\theta}) = (2k\pi, 0)$ for $k \in \mathbb{N}$ are stable equilibrium states, while the points $((2k - 1)\pi, 0)$ are unstable equilibrium points, corresponding to the upside-down position of the pendulum. In (a), the baseline controller, trained using the average mass and damping from offline observations, fails to stabilize the pendulum to upright. In (b), the distributionally robust (DR) controller successfully stabilizes the pendulum to an upright position for every initial state, demonstrating improved robustness to distributional shifts in the system parameters.

VII. Evaluation

This section evaluates our approach for synthesizing distributionally robust Lyapunov-stable controllers. We focus on two classic control problems from Gymnasium [71]: the Inverted Pendulum and the Continuous-time Mountain Car. For each system, we consider multiple instances with varying characteristics, each defined by its own physical parameters such as mass, length, and friction. These instances are used for training our distributionally robust Lyapunov-stable controller (30).

We use PyTorch [72] to train our neural network Lyapunov function and controller pair, both for the baseline case (29) and for the distributionally robust (DR) formulation (30). The models are trained using the Adam optimizer [73] with a learning rate of 0.002. The PyTorch clamp function is used to enforce input bounds for the control policy network, which is implemented with a smooth approximation and enables gradient back-propagation. The training was done on hardware with an Nvidia GeForce RTX 2070 Super GPU and an Intel i7 9700K CPU.

During testing, we assume the physical parameters of each system may be drawn from distributions different from those of the training instances. This setup aims to assess the performance of the learned controller under realistic scenarios, where distributional shifts in the system parameters are common. By doing so, we effectively evaluate the robustness and adaptability of our controller in the presence of distributional uncertainty in the system models.

A. Inverted Pendulum

The inverted pendulum is a standard nonlinear control problem for testing control methods. The system consists of two state variables, angular position θ and angular velocity $\dot{\theta}$, and one control input u . The system dynamics are:

$$\underbrace{\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \end{bmatrix}}_{\dot{\mathbf{x}}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) := \begin{bmatrix} \dot{\theta} \\ \frac{mgl \sin \theta - b\dot{\theta}}{ml^2} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{ml^2} \end{bmatrix} u \quad (41)$$

where $g = 9.81 \text{ m/s}^2$ is the gravity acceleration, $m = 1.0 \text{ kg}$ is the mass, $l = 1.0 \text{ m}$ is the length, and $b = 0.13 \text{ N}\cdot\text{m}\cdot\text{s}/\text{rad}$ is the damping coefficient. The equilibrium state is defined as $[\theta, \dot{\theta}] = [0, 0]$, which corresponds to the upright equilibrium position.

For training both the baseline and the DR controllers, we generate a training dataset by uniformly sampling $\{\mathbf{x}_i\}_{i \in [3600]}$ from the box region defined by $0 \leq \theta \leq 2\pi$ and $-8 \leq \dot{\theta} \leq 8$.

We assume uncertainties in mass and damping, represented by ξ_1 and ξ_2 , respectively. The uncertain system dynamics are:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\theta} \\ \frac{(m+\xi_1)gl \sin \theta - (b+\xi_2)\dot{\theta}}{(m+\xi_1)l^2} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{(m+\xi_1)l^2} \end{bmatrix} u. \quad (42)$$

Through first-order Taylor expansion around the nominal parameters, we derive the perturbation matrices $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2]$ for the inverted pendulum subject to model uncertainty:

$$\mathbf{w}_1(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} 0 \\ \frac{b\dot{\theta} - u}{m^2 l^2} \end{bmatrix}, \quad \mathbf{w}_2(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} 0 \\ -\frac{\dot{\theta}}{ml^2} \end{bmatrix}. \quad (43)$$

We assume that a set of offline samples $\{\xi_i\}_{i=1}^5$ of the uncertainties $\xi = [\xi_1, \xi_2]^T$ are available for training the controller and certificate pairs. We take $\xi_1 \sim \mathcal{U}(-0.04, 0.08)$, $\xi_2 \sim \mathcal{N}(0.0, 0.02)$, where \mathcal{U} and \mathcal{N} denote uniform and normal distributions, respectively. However, during test time, the test uncertainty parameters are set to $\xi_1 = 0.1$ and $\xi_2 = 0.05$. These test values represent a significant deviation from the training distributions, allowing us to evaluate the controllers' robustness to distributional shifts in parameters.

To train the baseline controller, we compute the average mass \tilde{m} and the average damping \tilde{b} from the five offline samples. These average values are then used in the baseline system dynamics and the loss function (29) to learn the baseline controller and Lyapunov function pair.

In contrast, for training the DR controller, we set the Wasserstein radius to $r = 0.01$ and the risk tolerance to $\epsilon = 0.1$. The 5 uncertainty samples $\{\xi_i\}_{i=1}^5$ are then used in the loss function (30) to learn the DR controller and its corresponding Lyapunov function.

Fig. 2 illustrates the performance of the baseline and DR controllers with the test uncertainty parameters $\xi_1 = 0.1$ and $\xi_2 = 0.05$. To assess the controllers' performance, we randomly sample 10 initial states within the box region defined by $0 \leq \theta \leq 2\pi$ and $-6 \leq \dot{\theta} \leq 6$. Each controller is applied to the system and the resulting trajectories are simulated.




FIGURE 3: Comparison of Lyapunov function values, control inputs, and RL value functions for the initial state $(\pi, 0)$. (a) Lyapunov function values over time for the baseline and DR controllers. (b) Control inputs over time for the baseline and DR controllers. (c) Value function values over time for the SAC and PPO algorithms. (d) Control inputs generated by the SAC and PPO algorithms.




FIGURE 4: Comparison of Lyapunov function values, control inputs, and RL value functions for the initial state $(-\pi/2, 5.5)$. (a) Lyapunov function values over time for the baseline and DR controllers. (b) Control inputs over time for the baseline and DR controllers. (c) Value function values over time for the SAC and PPO algorithms. (d) Control inputs generated by the SAC and PPO algorithms.

In Fig. 2a, we observe that the baseline controller, trained using only the average mass and damping values, fails to stabilize the inverted pendulum system to the desired upright equilibrium. For all the 10 initial states, the baseline controller's trajectories converge to states that are not the desired equilibrium. This can be attributed to the increased damping ($\xi_2 = 0.05$) and the heavier pendulum mass ($\xi_1 = 0.1$) in the test scenario. The baseline controller, designed based on the average parameter values, lacks the necessary robustness to compensate for the increased damping and mass.

On the other hand, Fig. 2b shows that all trajectories converge to the desired equilibrium under the DR controller. Despite the presence of distributional shift in the model uncertainty, the DR controller successfully stabilizes the inverted pendulum in an upright position. This robustness can be attributed to the DR controller's training process, which explicitly takes into account the distributional information of the uncertainty and optimizes for worst-case performance within the constructed ambiguity set.

To further investigate the performance of the proposed approach, we compare it with state-of-the-art reinforcement learning (RL) algorithms, such as the Soft Actor-Critic algorithm (SAC) [74] and the Proximal Policy Optimization algorithm (PPO) [75]. The reward function used for training the RL algorithms is defined as $r = -(\theta^2 + 0.1\dot{\theta}^2 + 0.001u^2)$.

It is worth noting that our approach and the RL algorithms differ in their sample complexity. We use a dataset of 3600 state samples uniformly sampled from the region of interest for training both the baseline and the DR controllers. In

contrast, the stable baselines3 implementations [76] of SAC and PPO require a significantly larger number of environment interactions, with a total of 2.048 billion steps for SAC and 4.096 billion steps for PPO. This higher sample complexity of the RL algorithms can be attributed to their model-free nature, as they learn directly from experience without explicitly using a dynamics model. Our approach incorporates the system dynamics model through Lyapunov theory, leading to a lower sample complexity.

Figs. 3 and 4 present a qualitative comparison of the baseline controller, our proposed DR controller, and the two RL policies over time for two different initial states. In Fig. 3, the initial state is the downright position $(\theta, \dot{\theta}) = (\pi, 0)$, from which the pendulum has to swing up to the desired equilibrium. We observe that the Lyapunov function corresponding to the baseline controller fails to converge to zero, while the proposed DR Lyapunov function values successfully converge to zero, as illustrated in Fig. 3a. Although the value functions learned by the two RL algorithms eventually reach zero, they exhibit non-monotonic behavior, indicating that the stability of the system cannot be certified using these learned value functions (Fig. 3c).

In Fig. 4, we change the initial state to $(-\pi/2, 5.5)$. In this scenario, the non-monotonic behavior of the value functions learned by the RL algorithms becomes more apparent, as shown in Fig. 4c. This observation highlights the advantage of the proposed DR approach, which provides provable stability guarantees through the use of a certified Lyapunov function. The DR controller ensures that the Lyapunov

function values monotonically decrease towards zero with high probability (Fig. 4a), thereby enabling robustness to distributional shifts in the system parameters.

In conclusion, while the RL algorithms can learn control policies effectively, their lack of provable stability guarantees may limit their applicability in real-world control and robot systems. In contrast, the proposed DR approach offers a principled way to design controllers that are robust to distributional shifts and provide certifiable stability guarantees.

Remark VII.1 (Impact of Input Constraints on Controller Performance). In practical applications, control inputs are often subject to physical limitations, such as actuator saturation. These constraints can significantly impact the performance and stability of the controlled system. In our implementation, we enforce control input constraints using the $\text{clamp}(\cdot, \mathbf{u}_{\min}, \mathbf{u}_{\max})$ function, which ensures that the output remains within the specified control limits. Specifically, for the inverted pendulum system, we constrain the input u within the bounds $-15 \leq u \leq 15$.

To investigate the effect of input constraints on our Lyapunov-based approach and on the reinforcement learning algorithms, we also considered a scenario where the control bounds are tightened to $-9 \leq u \leq 9$. Under these stricter input constraints, the Lyapunov-based approach failed to learn a controller with a valid Lyapunov function certificate. This failure occurs because the limited control authority hinders the ability of the Lyapunov-based controller to drive the system towards the desired equilibrium state, which is having the pendulum in the upright position.

Interestingly, some reinforcement learning algorithms, such as PPO [75], can still learn effective control policies even under tightened input constraints. These RL algorithms have the flexibility to discover alternative control strategies that exploit the system dynamics to achieve the desired goal. For instance, the RL controller may learn to swing the pendulum in one direction, leveraging gravity to gain acceleration, and then stabilize the pendulum in the upright position from the opposite direction. This swinging behavior allows the RL controller to overcome the limitations imposed by the reduced control bounds and successfully drive the system towards the equilibrium. •

B. Mountain Car

We consider a car placed at the bottom of a sinusoidal valley. The goal is to drive the car to the top of the mountain. The system has two states, velocity v and position x of the car, and one control input u , representing the directional force applied on the car. The dynamics model is given by [77]:

$$\underbrace{\begin{bmatrix} \dot{x} \\ \dot{v} \end{bmatrix}}_{\dot{\mathbf{x}}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) := \begin{bmatrix} v \\ -0.0025 \cos(3x) \end{bmatrix} + \begin{bmatrix} 0 \\ p \end{bmatrix} u \quad (44)$$

where $p = 0.0015$ denotes the constant power of the car. The desired equilibrium is $(x, v) = (\pi/6, 0)$, corresponding to the peak of the sinusoidal valley. For training both the baseline and the DR controllers, we generate a training

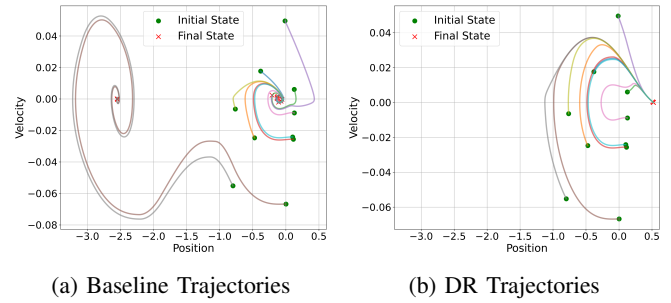


FIGURE 5: Comparison of trajectories for the mountain car system with test case power parameter $p = 0.0012$. The 10 random sampled initial states are marked as green dots while the final states are marked as red crosses. In (a), the baseline controller, trained using the average power from offline observations, fails to stabilize the car to the desired equilibrium $(\pi/6, 0)$. Instead, (b) shows the distributionally robust (DR) controller successfully stabilizing the car to the top of the mountain for every initial state, demonstrating improved robustness to distributional shifts in system parameters.

dataset by uniformly sampling $\{\mathbf{x}_i\}_{i \in [1600]}$ from the box region defined by $-2 \leq x \leq 2$ and $-0.4 \leq v \leq 0.4$. The control input is constrained within the bounds $-2 \leq u \leq 2$. We consider uncertainties in the car power p , represented by ξ . The uncertain system dynamics can be expressed as:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) + \begin{bmatrix} 0 \\ u \end{bmatrix} \xi. \quad (45)$$

We take $\{\xi_i\}_{i=1}^3$ offline uncertainty samples from the normal distribution $\mathcal{N}(0.0, 0.0002)$ are available. However, during test time, the uncertainty parameters are set to $\xi = -0.0003$.

For the baseline controller, we calculate the average power \bar{p} from the available power samples and incorporate it into the system dynamics (44) during the learning process of the baseline Lyapunov-stable controller. In contrast, for training the DR controller, we set the Wasserstein radius to $r = 0.0001$ and the risk tolerance to $\epsilon = 0.1$. The offline samples $\{\xi_i\}_{i \in [3]}$ are then utilized in the loss function (30) to learn the controller and Lyapunov function pair simultaneously.

To evaluate the performance of the baseline and DR controllers on the mountain car system, we consider a test scenario with uncertainty parameters set to $\xi = -0.0003$, representing a decrease in the car's power. We randomly initialize 10 starting states within the box region defined by $-1 \leq x \leq 1$ and $-0.1 \leq v \leq 0.1$. The controllers are then applied to the system, and the resulting trajectories are simulated to assess their effectiveness.

Fig. 5a reveals that the baseline controller, which is trained on the average power values, fails to drive the car to the desired equilibrium at the top of the mountain. The trajectories corresponding to all 10 initial states fail to converge to the desired goal. The baseline controller's design, based on average parameter values, lacks the robustness necessary to adapt to the decreased power.

In contrast, Fig. 5b showcases the superior performance of the DR controller. Despite the presence of distributional shift in the system dynamics, the DR controller successfully drives all trajectories to the desired equilibrium at the mountain's peak. By explicitly considering the potential variations in the car's power during training, the DR controller is able to adapt and maintain its effectiveness for out-of-distribution parameter uncertainties.

VIII. Conclusions

We have presented an approach for the joint synthesis of distributionally robust controllers and Lyapunov certificates suitable for nonlinear systems with model uncertainty. Our key contributions lie in the development of a distributionally robust Lyapunov stability formulation, leveraging finite samples of the model uncertainty parameters to address the challenge of stabilizing the uncertain system with guarantees. Through both theoretical analysis and experimental validation, we demonstrated that the system governed by the learned controller exhibits asymptotic stability with high probability, even in the presence of out-of-distribution model uncertainties.

Our approach contributes a new perspective to controlling systems with model uncertainty, offering a promising direction for certificate-based reinforcement learning in real systems. However, scalability challenges arise from the need to satisfy the Lyapunov condition for uniformly sampled data in the state space, leading to an exponential growth in the number of training samples as the dimension of the system state increases. Future research will focus on adapting our methods to more complex robotic systems. This includes exploring more efficient learning algorithms, considering relaxed stability notions, and establishing connections between Lyapunov theory and model-free reinforcement learning.

REFERENCES

- [1] K. Long, Y. Yi, J. Cortes, and N. Atanasov, "Distributionally robust Lyapunov function search under uncertainty," in *Learning for Dynamics and Control Conference*, pp. 864–877, PMLR, 2023.
- [2] J.-J. E. Slotine, W. Li, et al., *Applied nonlinear control*, vol. 199. Prentice hall Englewood Cliffs, NJ, 1991.
- [3] Z. Artstein, "Stabilization with relaxed controls," *Nonlinear Analysis-theory Methods & Applications*, vol. 7, pp. 1163–1173, 1983.
- [4] E. Sontag, "A 'universal' construction of Artstein's theorem on nonlinear stabilization," *Systems & Control Letters*, vol. 13, no. 2, pp. 117–123, 1989.
- [5] W. M. Haddad and V. Chellaboina, *Nonlinear dynamical systems and control: a Lyapunov-based approach*. Princeton university press, 2008.
- [6] P. Parrilo, *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.
- [7] A. Papachristodoulou and S. Prajna, "On the construction of Lyapunov functions using the sum of squares decomposition," in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, vol. 3, pp. 3482–3487, 2002.
- [8] Z. Jarvis-Wloszek, R. Feeley, W. Tan, K. Sun, and A. Packard, "Some controls applications of sum of squares programming," in *42nd IEEE International Conference on Decision and Control*, vol. 5, pp. 4676–4681, 2003.
- [9] Y.-C. Chang, N. Roohi, and S. Gao, "Neural Lyapunov control," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [10] S. M. Richards, F. Berkenkamp, and A. Krause, "The Lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems," in *Conference on Robot Learning*, pp. 466–476, PMLR, 2018.
- [11] H. Dai, B. Landry, L. Yang, M. Pavone, and R. Tedrake, "Lyapunov-stable neural-network control," in *Proceedings of Robotics: Science and Systems*, (Virtual), July 2021.
- [12] N. Gaby, F. Zhang, and X. Ye, "Lyapunov-net: A deep neural network architecture for Lyapunov function approximation," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 2091–2096, 2022.
- [13] T. Li and N. Figueroa, "Task generalization with stability guarantees via elastic dynamical system motion policies," in *Conference on Robot Learning*, pp. 3485–3517, PMLR, 2023.
- [14] J. Wu, A. Clark, Y. Kantaros, and Y. Vorobeychik, "Neural Lyapunov control for discrete-time systems," *Advances in neural information processing systems*, vol. 36, pp. 2939–2955, 2023.
- [15] S. Zhang and C. Fan, "Learning to stabilize high-dimensional unknown systems using Lyapunov-guided exploration," in *6th Annual Learning for Dynamics & Control Conference*, pp. 52–67, PMLR, 2024.
- [16] A. Taylor, V. Dorobantu, H. Le, Y. Yue, and A. Ames, "Episodic learning with control Lyapunov functions for uncertain robotic systems," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6878–6884, 2019.
- [17] F. Castañeda, J. Choi, B. Zhang, C. Tomlin, and K. Sreenath, "Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects and dynamics," in *2021 American Control Conference (ACC)*, pp. 3683–3690, 2021.
- [18] 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.
- [19] P. Mestres and J. Cortés, "Feasibility and regularity analysis of safe stabilizing controllers under uncertainty," *Automatica*, vol. 167, p. 111800, 2024.
- [20] P. Mestres, K. Long, N. Atanasov, and J. Cortés, "Feasibility analysis and regularity characterization of distributionally robust safe stabilizing controllers," *IEEE Control Systems Letters*, vol. 8, pp. 91–96, 2024.
- [21] K. Long, C. Qian, J. Cortés, and N. Atanasov, "Learning barrier functions with memory for robust safe navigation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4931–4938, 2021.
- [22] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming*. Society for Industrial and Applied Mathematics, 2009.
- [23] 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.
- [24] A. R. Hota, A. K. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under Wasserstein ambiguity sets," in *2019 American Control Conference (ACC)*, pp. 1501–1506, 2019.
- [25] I. Yang, "Wasserstein distributionally robust stochastic control: A data-driven approach," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3863–3870, 2020.
- [26] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3289–3304, 2021.
- [27] A. B. Kordabad, R. Wisniewski, and S. Gros, "Safe reinforcement learning using Wasserstein distributionally robust MPC and chance constraint," *IEEE Access*, vol. 10, pp. 130058–130067, 2022.
- [28] P. Lathrop, B. Boardman, and S. Martínez, "Distributionally safe path planning: Wasserstein safe RRT," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 430–437, 2021.
- [29] A. Hakobyan and I. Yang, "Distributionally robust differential dynamic programming with Wasserstein distance," *IEEE Control Systems Letters*, vol. 7, pp. 2329–2334, 2023.
- [30] B. Li, Y. Tan, A.-G. Wu, and G.-R. Duan, "A distributionally robust optimization based method for stochastic model predictive control," *IEEE Transactions on Automatic Control*, vol. 67, no. 11, pp. 5762–5776, 2021.
- [31] L. Aolaritei, M. Fochesato, J. Lygeros, and F. Dörfler, "Wasserstein tube MPC with exact uncertainty propagation," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 2036–2041, IEEE, 2023.
- [32] B. Li, T. Guan, L. Dai, and G.-R. Duan, "Distributionally robust model predictive control with output feedback," *IEEE Transactions on Automatic Control*, vol. 69, no. 5, pp. 3270–3277, 2024.

- [33] F. Micheli, T. Summers, and J. Lygeros, "Data-driven distributionally robust mpc for systems with uncertain dynamics," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 4788–4793, IEEE, 2022.
- [34] B. Taskesen, D. Iancu, Ç. Koçyiğit, and D. Kuhn, "Distributionally robust linear quadratic control," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [35] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.
- [36] D. Hilbert, "Ueber die darstellung definiter formen als summe von formenquadraten," *Mathematische Annalen*, vol. 32, pp. 342–350, September 1888.
- [37] N. Boffi, S. Tu, N. Matni, J.-J. Slotine, and V. Sindhvani, "Learning stability certificates from data," in *Conference on Robot Learning*, pp. 1341–1350, PMLR, 2021.
- [38] R. Zhou, T. Quartz, H. De Sterck, and J. Liu, "Neural Lyapunov control of unknown nonlinear systems with stability guarantees," in *Advances in Neural Information Processing Systems*, vol. 35, 2022.
- [39] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural Lyapunov-barrier functions," in *Conference on Robot Learning*, vol. 164, pp. 1724–1735, PMLR, 2022.
- [40] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural Lyapunov, barrier, and contraction methods for robotics and control," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1749–1767, 2023.
- [41] T. Westenbroek, F. Castaneda, A. Agrawal, S. Sastry, and K. Sreenath, "Lyapunov design for robust and efficient robotic reinforcement learning," *arXiv preprint arXiv:2208.06721*, 2022.
- [42] A. Lopez and D. Fridovich-Keil, "Decomposing control lyapunov functions for efficient reinforcement learning," *arXiv preprint arXiv:2403.12210*, 2024.
- [43] B. P. G. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2016.
- [44] R. Jiang and Y. Guan, "Data-driven chance constrained stochastic program," *Mathematical Programming*, vol. 158, pp. 291–327, 2016.
- [45] W. Xie, "On distributionally robust chance constrained programs with Wasserstein distance," *Math. Program.*, vol. 186, pp. 115–155, 2021.
- [46] S. Sagawa*, P. W. Koh*, T. B. Hashimoto, and P. Liang, "Distributionally robust neural networks," in *International Conference on Learning Representations*, 2020.
- [47] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.
- [48] F. Boso, D. Boskos, J. Cortés, S. Martínez, and D. M. Tartakovsky, "Dynamics of data-driven ambiguity sets for hyperbolic conservation laws with uncertain inputs," *SIAM Journal on Scientific Computing*, vol. 43, no. 3, pp. A2102–A2129, 2021.
- [49] 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.
- [50] A. Cherukuri and A. R. Hota, "Consistency of distributionally robust risk- and chance-constrained optimization under Wasserstein ambiguity sets," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1729–1734, 2021.
- [51] D. Boskos, J. Cortés, and S. Martínez, "High-confidence data-driven ambiguity sets for time-varying linear systems," *IEEE Transactions on Automatic Control*, vol. 69, no. 2, pp. 797–812, 2024.
- [52] A. Z. Ren and A. Majumdar, "Distributionally robust policy learning via adversarial environment generation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1379–1386, 2022.
- [53] K. Long, Y. Yi, J. Cortés, and N. Atanasov, "Safe and stable control synthesis for uncertain system models via distributionally robust optimization," in *2023 American Control Conference (ACC)*, pp. 4651–4658, 2023.
- [54] T. Summers, "Distributionally robust sampling-based motion planning under uncertainty," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6518–6523, 2018.
- [55] S. Sastry, *Nonlinear systems: analysis, stability, and control*, vol. 10. Springer Science & Business Media, 2013.
- [56] R. Freeman and P. Kototovic, *Robust Nonlinear Control Design*. Cambridge, MA, USA: Birkhäuser Boston Inc., 1996.
- [57] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–41, 2000.
- [58] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM J. Optim.*, vol. 17, pp. 969–996, 2006.
- [59] H. J. Kushner and Kushner, *Stochastic stability and control*, vol. 33. Academic press New York, 1967.
- [60] A. R. Teel, J. P. Hespanha, and A. Subbaraman, "A converse Lyapunov theorem and robustness for asymptotic stability in probability," *IEEE Transactions on Automatic Control*, vol. 59, no. 9, pp. 2426–2441, 2014.
- [61] P. Culbertson, R. K. Cosner, M. Tucker, and A. D. Ames, "Input-to-state stability in probability," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 5796–5803, IEEE, 2023.
- [62] J. Steinhardt and R. Tedrake, "Finite-time regional verification of stochastic non-linear systems," *The International Journal of Robotics Research*, vol. 31, no. 7, pp. 901–923, 2012.
- [63] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, p. 109439, 2021.
- [64] N. Weaver, *Lipschitz algebras*. World Scientific, 2018.
- [65] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [66] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [67] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken, "Multilayer feedforward networks with a nonpolynomial activation function can approximate any function," *Neural networks*, vol. 6, no. 6, pp. 861–867, 1993.
- [68] L. Yang, H. Dai, Z. Shi, C.-J. Hsieh, R. Tedrake, and H. Zhang, "Lyapunov-stable neural control for state and output feedback: A novel formulation," in *Forty-first International Conference on Machine Learning*, 2024.
- [69] S. Wang, H. Zhang, K. Xu, X. Lin, S. Jana, C.-J. Hsieh, and J. Z. Kolter, "Beta-CROWN: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network verification," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [70] H. Wu, O. Isac, A. Zeljić, T. Tagomori, M. Daggitt, W. Kokke, I. Re-faeli, G. Amir, K. Julian, S. Bassan, *et al.*, "Marabou 2.0: A versatile formal analyzer of neural networks," *arXiv preprint arXiv:2401.14461*, 2024.
- [71] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. d. Cola, T. Deleu, M. Goulão, A. Kallinteris, A. KG, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, "Gymnasium," Mar. 2023.
- [72] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in neural information processing systems*, vol. 32, 2019.
- [73] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2015.
- [74] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*, pp. 1861–1870, PMLR, 2018.
- [75] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [76] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dornmann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021.
- [77] A. W. Moore, "Efficient memory-based learning for robot control," tech. rep., University of Cambridge, 1990.