

Distributionally Robust Lyapunov-Barrier Networks for Safe and Stable Control Under Uncertainty

Ali Baheri

Abstract

This paper presents a novel approach to ensuring both stability and safety in nonlinear control systems subject to parametric uncertainties. We introduce distributionally robust Lyapunov-Barrier networks (DR-LBNs), which combine Lyapunov stability certificates and barrier functions within a neural network framework, while accounting for distributional uncertainty in system parameters. Our method provides probabilistic guarantees on both asymptotic stability and forward invariance of a safe set, even when the true distribution of uncertainties is unknown. We develop a comprehensive theoretical foundation for DR-LBNs, presenting and proving several key theorems. These include results on probabilistic stability and safety guarantees, robustness to distribution shifts, input-to-state stability, universal approximation capabilities, and sample complexity bounds.

1 Introduction

The control of nonlinear dynamical systems under uncertainty remains a fundamental challenge in robotics, aerospace, and other safety-critical domains. As autonomous systems become increasingly prevalent in real-world environments, the need for controllers that can ensure both stability and safety in the face of parametric uncertainties has never been more pressing. Traditional approaches to robust control often rely on worst-case analyses or assume known bounds on uncertainties, leading to overly conservative controllers that may significantly sacrifice performance for the sake of robustness [11].

Recent years have seen a surge of interest in learning-based control methods, particularly those leveraging neural networks for their ability to approximate complex functions [4]. These approaches have shown remarkable success in handling nonlinear dynamics and adapting to uncertain environments [21]. However, they often lack formal guarantees on stability and safety, a critical requirement for deployment in safety-critical applications. Moreover, the performance of these learned controllers can degrade significantly when faced with distribution shifts between training and deployment environments. Lyapunov theory and barrier functions have long been powerful tools in control theory for certifying stability and safety, respectively. However, their application to complex, high-dimensional systems with significant uncertainties remains challenging. The manual design of Lyapunov and barrier functions for such systems is often intractable, motivating the need for data-driven approaches to certificate synthesis. Distributionally robust optimization has emerged as a promising framework for decision-making under uncertainty, particularly when the true distribution of uncertainties is unknown. By optimizing over an ambiguity set of possible distributions, this approach provides robustness to distribution shifts while avoiding the excessive conservatism of worst-case analyses [15]. Despite its potential, the application of

distributionally robust optimization to nonlinear control synthesis, especially in conjunction with Lyapunov and barrier function approaches, remains largely unexplored.

In this paper, we propose distributionally robust Lyapunov-Barrier networks (DR-LBNs), a novel framework that addresses these challenges by unifying concepts from neural networks, Lyapunov theory, barrier functions, and distributionally robust optimization. Our approach aims to synthesize controllers for nonlinear systems that *simultaneously* ensure stability and safety under parametric uncertainties, with robustness to distribution shifts. By leveraging the expressive power of neural networks and the foundations of control theory, DR-LBNs offer a principled way to design controllers with probabilistic guarantees on performance, even when the true distribution of uncertainties is unknown. The DR-LBN framework represents a significant step towards bridging the gap between learning-based control and robust control theory. It offers the potential to combine the adaptability of neural network approaches with the formal guarantees associated with analytical control methods. This integration is crucial for the development of next-generation autonomous systems that can operate safely in uncertain, dynamic environments.

Our Contributions. In this paper, we present a series of key contributions that address challenges at the intersection of stability, safety, and robust control in learning-based systems. Our primary contributions are as follows:

- *Unified Neural Architecture for Stability and Safety:* We introduce a novel neural network architecture that simultaneously learns Lyapunov functions, barrier functions, and adaptive control gains. This unified approach enables the synthesis of controllers that jointly optimize for stability and safety, a significant advancement over traditional methods that often treat these objectives separately.
- *Distributionally Robust Formulation of Control Synthesis:* We develop a distributionally robust framework for control synthesis that explicitly accounts for uncertainties in the underlying system parameters. Unlike traditional robust control methods that rely on worst-case analyses or assume known uncertainty distributions, our approach constructs an ambiguity set around the empirical distribution of uncertainties. This formulation leads to controllers that are robust to distribution shifts between training and deployment.
- *Theoretical Guarantees:* We provide a comprehensive theoretical analysis of the DR-LBN framework, including proofs of input-to-state stability, universal approximation capabilities, and sample complexity bounds.

Paper Organization. The remainder of this paper is organized as follows. Section 2 provides a comprehensive review of related work, situating our approach within the broader context of control Lyapunov functions, control barrier functions, and distributionally robust optimization. Section 3 presents the problem formulation and control objectives. Section 4 introduces the core methodology of Distributionally Robust Lyapunov-Barrier Networks (DR-LBNs), including the neural network architecture, distributionally robust stability and safety conditions, and the optimization framework. Section 5 offers theoretical analysis, presenting key results on probabilistic stability guarantees, robustness to distribution shifts, input-to-state stability, universal approximation capabilities, and sample complexity bounds. Finally, Section 6 concludes the paper and discusses future research directions.

2 Related Work

Control Lyapunov Functions (CLFs). Control Lyapunov Functions (CLFs) have long been a key tool in nonlinear control theory, offering a structured approach to stabilizing nonlinear systems [3, 23]. CLFs provide a method for designing controllers that ensure asymptotic stability by ensuring the derivative of a positive definite function decreases along system trajectories. Over the years, the concept has evolved significantly, including the integration of neural networks with CLFs. One approach leverages deep neural networks and sum-of-squares programming to learn CLFs, enabling the design of stabilizing controllers for systems with complex, high-dimensional dynamics [22]. This method harnesses the power of neural networks to approximate CLFs for cases where analytical solutions are difficult or impossible to obtain. Additionally, data-driven approaches have gained prominence, with recent work proposing frameworks that learn CLFs directly from system trajectories, bypassing the need for explicit knowledge of system dynamics [20]. This development expands the applicability of CLFs to systems with partially known or uncertain models.

Robustness to uncertainty has been a key focus of recent CLF research. Long et al. developed a distributionally robust framework for CLF search, providing probabilistic stability guarantees when the true system parameters are unknown [15]. This approach bridges the gap between traditional robust control methods and modern machine learning techniques. Extensions of CLF theory to hybrid systems have also been explored. The work in [10] developed conditions for asymptotic stability in hybrid loops using CLFs, broadening the applicability of these methods to systems with both continuous and discrete dynamics. Adaptive CLFs have been introduced in [25] to handle parametric uncertainties in nonlinear systems. By combining CLFs with adaptive control techniques, their approach allows for real-time adaptation to changing system parameters. The stochastic domain has also seen advancements in CLF theory. Taylor et al. proposed a reinforcement learning approach to synthesize CLFs, combining model-free learning with formal stability guarantees [24]. This method offers a promising direction for systems where accurate models are difficult to obtain.

Control Barrier Functions (CBFs). CBFs is a powerful tool for enforcing safety constraints in nonlinear control systems. CBFs provide a means to verify and synthesize controllers that ensure forward invariance of a safe set. The modern formulation of CBFs, draws inspiration from CLFs and allows for the unification of safety and stability in a single framework. CBFs have found applications in various domains, including automotive systems [2], multi-robot coordination [26], and bipedal robots [12]. Recent advancements in CBF theory have addressed several practical challenges in safe control synthesis [1]. The work in [9] presented novel CBFs that handle safety constraints with time and input constraints under disturbances, as well as high-relative degree constraints under disturbances and input constraints. Their work also addresses the effects of adversarial inputs and sampled-data implementation. Adaptive CBFs have been developed to handle parametric uncertainties in nonlinear systems. Lopez et al. introduced a framework that guarantees safety through online parameter adaptation and data-driven model estimation [16]. This approach allows for real-time adaptation to changing system parameters, enhancing robustness in uncertain environments. The integration of CBFs with learning-based approaches has been a focus of recent research. For instance, CBFs have been extended to handle multi-agent systems. Zhang et al. introduced a neural graph CBF framework for distributed safe multi-agent control, enabling decentralized stabilization of interconnected nonlinear systems while ensuring safety

constraints [27].

Distributionally Robust Optimization (DRO). DRO has gained significant attention in recent years as a framework for decision-making under uncertainty, particularly when the true distribution of uncertainties is unknown [8, 19]. DRO aims to optimize performance while providing robustness against a set of possible distributions, typically defined as an ambiguity set around an empirical distribution. The fundamental idea behind DRO is to find solutions that perform well under the worst-case probability distribution within a specified ambiguity set. This ambiguity set is often constructed using statistical distance measures such as φ -divergences or the Wasserstein metric. By considering a range of possible distributions, DRO provides a safeguard against distributional ambiguity and potential model misspecification. DRO has found applications in various fields, including machine learning [28], operations research, and finance [13]. In machine learning, DRO has been used to develop robust classification and regression models that maintain good performance across different data distributions. In operations research, DRO has been applied to supply chain management, logistics, and resource allocation problems to handle uncertainties in demand and supply [5, 17].

3 Problem Formulation and Methodological Approach

We consider a nonlinear control-affine system with parametric uncertainties:

$$\frac{dx}{dt} = f(x, \xi) + g(x, \xi)u \quad (1)$$

where:

- $x \in X \subseteq \mathbb{R}^n$ is the state vector
- $u \in U \subseteq \mathbb{R}^m$ is the control input
- $\xi \in \Xi \subseteq \mathbb{R}^k$ represents uncertain parameters
- $f : X \times \Xi \rightarrow \mathbb{R}^n$ is the drift dynamics
- $g : X \times \Xi \rightarrow \mathbb{R}^{n \times m}$ is the control input matrix

Assumption 1: The functions f and g are locally Lipschitz continuous, ensuring the existence and uniqueness of solutions.

Assumption 2: The sets X , U , and Ξ are compact, reflecting physical limitations of the system and bounded uncertainties.

The safe set $S \subset X$ is defined as:

$$S = \{x \in X : h(x) \geq 0\} \quad (2)$$

where $h : X \rightarrow \mathbb{R}$ is a continuously differentiable function representing safety constraints. The true distribution P^* of ξ is unknown. We have access to N independent and identically distributed samples $\{\xi_1, \dots, \xi_N\}$. Let P_N denote the empirical distribution based on these samples:

$$P_N = \frac{1}{N} \sum_{i=1}^N \delta_{\xi_i} \quad (3)$$

where δ_{ξ_i} is the Dirac measure at ξ_i . To account for the discrepancy between P_N and P^* , we define an ambiguity set \mathcal{P} as a Wasserstein ball around P_N :

$$\mathcal{P} = \{P \in \mathcal{M}(\Xi) : W_p(P, P_N) \leq \rho\} \quad (4)$$

where:

- $\mathcal{M}(\Xi)$ is the set of all probability measures on Ξ
- W_p is the p -Wasserstein distance
- $\rho > 0$ is the radius of the Wasserstein ball

The p -Wasserstein distance between two probability measures P_1 and P_2 is defined as:

$$W_p(P_1, P_2) = \left(\inf_{\gamma \in \Gamma(P_1, P_2)} \int_{\Xi \times \Xi} \|\xi_1 - \xi_2\|^p d\gamma(\xi_1, \xi_2) \right)^{1/p} \quad (5)$$

where $\Gamma(P_1, P_2)$ is the set of all joint distributions with marginals P_1 and P_2 . Our objective is to design a state-feedback controller $u = \pi(x)$ that ensures both stability and safety with high probability under distributional uncertainty. Specifically, we aim to:

1. Stabilize the system to the origin: $\lim_{t \rightarrow \infty} \|x(t)\| = 0$
2. Ensure the state remains in the safe set: $x(t) \in S$ for all $t \geq 0$

Both objectives should be achieved with high probability with respect to the true (unknown) distribution P^* .

4 Distributionally Robust Lyapunov-Barrier Networks

To address this challenge, we propose the Distributionally Robust Lyapunov-Barrier Networks (DR-LBNs) approach. We introduce a neural network $\Phi_\theta : X \rightarrow \mathbb{R}^4$ parameterized by θ , which simultaneously learns:

$$\Phi_\theta(x) = [V(x), h(x), k_V(x), k_h(x)] \quad (6)$$

where:

- $V(x)$ is a candidate Lyapunov function
- $h(x)$ is a candidate barrier function
- $k_V(x)$ and $k_h(x)$ are adaptive gain functions

The architecture consists of shared layers followed by separate heads:

$$z = \varphi(x; \theta_{\text{shared}}) \quad (7)$$

$$V(x) = \|\psi_V(z; \theta_V)\|^2 + \varepsilon \|x\|^2 \quad (8)$$

$$h(x) = \text{softplus}(\psi_h(z; \theta_h)) \quad (9)$$

$$k_V(x) = \text{softplus}(\psi_{k_V}(z; \theta_{k_V})) \quad (10)$$

$$k_h(x) = \text{softplus}(\psi_{k_h}(z; \theta_{k_h})) \quad (11)$$

where $\varepsilon > 0$ is a small constant and $\text{softplus}(a) = \log(1 + e^a)$.

Distributionally Robust Stability and Safety Conditions. We formulate the distributionally robust Lyapunov condition as:

$$\inf_{P \in \mathcal{P}} P(L_V(x, \xi, \pi) < 0) \geq 1 - \varepsilon_V, \quad \forall x \in X \setminus \{0\} \quad (12)$$

where $L_V(x, \xi, \pi) = \nabla V(x)^\top (f(x, \xi) + g(x, \xi)\pi(x)) + \alpha_1(\|x\|)$, and α_1 is a class \mathcal{K} function. Similarly, the distributionally robust barrier condition is:

$$\inf_{P \in \mathcal{P}} P(L_h(x, \xi, \pi) > 0) \geq 1 - \varepsilon_h, \quad \forall x \in S \quad (13)$$

where $L_h(x, \xi, \pi) = \nabla h(x)^\top (f(x, \xi) + g(x, \xi)\pi(x)) + \alpha_2(h(x))$, and α_2 is a class \mathcal{K} function.

Control Law Synthesis. We propose a control law that leverages both the Lyapunov and barrier functions:

$$\pi(x) = \pi_V(x) + \pi_h(x) \quad (14)$$

where:

$$\pi_V(x) = -k_V(x)g(x, \hat{\xi})^\top \nabla V(x) \quad (15)$$

$$\pi_h(x) = k_h(x)g(x, \hat{\xi})^\top \nabla h(x) \quad (16)$$

Here, $\hat{\xi}$ is an estimate of ξ (e.g., the mean of the empirical distribution). We formulate the DR-LBN training as an optimization problem:

$$\min_{\theta} \mathcal{L}(\theta) = \lambda_V \mathcal{L}_V(\theta) + \lambda_h \mathcal{L}_h(\theta) + \lambda_u \mathcal{L}_u(\theta) + \lambda_c \mathcal{L}_c(\theta) + \lambda_r \mathcal{L}_r(\theta) \quad (17)$$

where:

- $\mathcal{L}_V(\theta)$ and $\mathcal{L}_h(\theta)$ are losses for Lyapunov and barrier conditions
- $\mathcal{L}_u(\theta)$ penalizes excessive control effort
- $\mathcal{L}_c(\theta)$ penalizes constraint violations
- $\mathcal{L}_r(\theta)$ is a regularization term

Using the Wasserstein duality theorem [18], we reformulate the chance constraints:

$$\mathcal{L}_V(\theta) = \mathbb{E}_x \left[\sup_{\lambda_V \geq 0, \eta_V} \Phi_V(x, \theta, \lambda_V, \eta_V) \right] \quad (18)$$

$$\mathcal{L}_h(\theta) = \mathbb{E}_x \left[\sup_{\lambda_h \geq 0, \eta_h} \Phi_h(x, \theta, \lambda_h, \eta_h) \right] \quad (19)$$

where Φ_V and Φ_h are dual functions defined as:

$$\Phi_V(x, \theta, \lambda_V, \eta_V) = -\lambda_V \varepsilon_V - \eta_V - \frac{1}{N} \sum_{i=1}^N (\lambda_V (L_V(x, \xi_i, \pi_\theta) - \eta_V)^+ - \rho \lambda_V^{1/p}) \quad (20)$$

$$\Phi_h(x, \theta, \lambda_h, \eta_h) = -\lambda_h \varepsilon_h - \eta_h - \frac{1}{N} \sum_{i=1}^N (\lambda_h (-L_h(x, \xi_i, \pi_\theta) - \eta_h)^+ - \rho \lambda_h^{1/p}) \quad (21)$$

This formulation allows to solve the distributionally robust optimization problem using standard gradient-based optimization techniques, providing a novel approach to designing controllers with probabilistic stability and safety guarantees under uncertainty.

5 Theoretical Analysis

Having established the framework for DR-LBNs, we now turn our attention to the theoretical foundations underpinning this framework. The following section presents a series of key theoretical results that provide guarantees on the performance, robustness, and generalization capabilities of DR-LBNs.

Theorem 1. Consider the system (1) with uncertainty ξ following an unknown true distribution P^* . Let Φ_{θ^*} be the optimal DR-LBN obtained from solving the optimization problem, and let π^* be the corresponding controller. Assume that:

(i) The Wasserstein radius ρ is chosen such that $P^* \in \mathcal{P}$ with high probability. (ii) The distributional robustness conditions are satisfied.

Then, for any initial state $x_0 \in S$, the closed-loop system satisfies:

$$P^*(\|x(t)\| \rightarrow 0 \text{ as } t \rightarrow \infty \text{ and } x(t) \in S \text{ for all } t \geq 0) \geq 1 - (\varepsilon_V + \varepsilon_h) \quad (22)$$

where ε_V and ε_h are the risk tolerances for stability and safety, respectively.

Proof. We proceed in several steps:

Step 1: Probabilistic Lyapunov Decrease Condition

From the distributionally robust Lyapunov condition (49), we have:

$$\inf_{P \in \mathcal{P}} P(L_V(x, \xi, \pi^*) < 0) \geq 1 - \varepsilon_V, \quad \forall x \in X \setminus \{0\} \quad (23)$$

By assumption (i) and the definition of the infimum, this implies:

$$P^*(L_V(x, \xi, \pi^*) < 0) \geq 1 - \varepsilon_V, \quad \forall x \in X \setminus \{0\} \quad (24)$$

Step 2: Probabilistic Barrier Increase Condition Similarly, from the distributionally robust barrier condition (50):

$$\inf_{P \in \mathcal{P}} P(L_h(x, \xi, \pi^*) > 0) \geq 1 - \varepsilon_h, \quad \forall x \in S \quad (25)$$

This implies:

$$P^*(L_h(x, \xi, \pi^*) > 0) \geq 1 - \varepsilon_h, \quad \forall x \in S \quad (26)$$

Step 3: Joint Probability

Let E_V be the event that $L_V(x, \xi, \pi^*) < 0$ for all $x \in X \setminus \{0\}$, and E_h be the event that $L_h(x, \xi, \pi^*) > 0$ for all $x \in S$. By the union bound:

$$P^*(E_V \cap E_h) \geq 1 - P^*(E_V^c) - P^*(E_h^c) \quad (27)$$

$$\geq 1 - \varepsilon_V - \varepsilon_h \quad (28)$$

where E_V^c and E_h^c denote the complements of E_V and E_h , respectively.

Step 4: Stability Analysis Conditioned on the event E_V , we have:

$$\frac{dV(x)}{dt} = \nabla V(x)^\top (f(x, \xi) + g(x, \xi)\pi^*(x)) < -\alpha_1(\|x\|), \quad \forall x \in X \setminus \{0\} \quad (29)$$

where α_1 is a class \mathcal{K} function. Let $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a class \mathcal{K}_∞ function such that $\gamma(\|x\|) \leq V(x)$ for all $x \in X$. Such a function exists because $V(x)$ is positive definite. Then, from (29),

$$\frac{d}{dt}\gamma^{-1}(V(x)) \leq -\frac{\alpha_1(\gamma^{-1}(V(x)))}{\gamma'(\gamma^{-1}(V(x)))} \quad (30)$$

By the comparison principle, this implies that $\|x(t)\| \rightarrow 0$ as $t \rightarrow \infty$.

Step 5: Safety Analysis Conditioned on the event E_h , we have:

$$\frac{dh(x)}{dt} = \nabla h(x)^\top (f(x, \xi) + g(x, \xi)\pi^*(x)) > -\alpha_2(h(x)), \quad \forall x \in S \quad (31)$$

where α_2 is a class \mathcal{K} function. By the comparison principle, this implies that if $h(x_0) \geq 0$, then $h(x(t)) \geq 0$ for all $t \geq 0$, ensuring that $x(t)$ remains in S for all $t \geq 0$. From steps 4 and 5, we can conclude that when both E_V and E_h occur, the system state converges to the origin while remaining in the safe set S for all time. From step 3, we know that $P^*(E_V \cap E_h) \geq 1 - (\varepsilon_V + \varepsilon_h)$. Therefore:

$$P^*(\|x(t)\| \rightarrow 0 \text{ as } t \rightarrow \infty \text{ and } x(t) \in S \text{ for all } t \geq 0) \geq 1 - (\varepsilon_V + \varepsilon_h) \quad (32)$$

which completes the proof. \square

Implication. Theorem 1 establishes probabilistic guarantees on both stability and safety for systems controlled by DR-LBNs. This result is particularly significant in the context of uncertain nonlinear systems, where deterministic guarantees are often unattainable. The probabilistic nature of these guarantees allows for the treatment of uncertainty. The joint consideration of stability and safety within a single framework represents an advancement over existing methods that often treat these objectives separately. This approach enables the synthesis of controllers that can navigate the complex trade-offs between stability and safety in uncertain environments. However, it is important to note that the probabilistic nature of these guarantees may limit the applicability of DR-LBNs in settings where hard, deterministic constraints are required.

Theorem 2 (Robustness to Distribution Shift). Let P^* be the true distribution of the uncertainty ξ , and let \tilde{P} be a shifted distribution such that $W_p(\tilde{P}, P^*) \leq \Delta$, where W_p is the p -Wasserstein distance. Let π^* be the DR-LBN controller trained using the empirical distribution P_N . Assume that:

(i) The Wasserstein radius ρ in the ambiguity set \mathcal{P} is chosen such that $P^* \in \mathcal{P}$ with probability at least $1 - \delta$. (ii) The distributional robustness conditions are satisfied for P^* with risk tolerances ε_V and ε_h for stability and safety, respectively. Then, under the controller π^* , for any initial state $x_0 \in S$, we have:

$$\tilde{P}(\|x(t)\| \rightarrow 0 \text{ as } t \rightarrow \infty \text{ and } x(t) \in S \text{ for all } t \geq 0) \geq 1 - (\varepsilon_V + \varepsilon_h + 2C\Delta^q + \delta) \quad (33)$$

where $C > 0$ is a constant that depends on the Lipschitz constants of the system dynamics and the DR-LBN, and $q = p/(p+1)$.

Proof. We approach this in several steps:

Step 1: Wasserstein Distance Properties

We first recall a key property of the Wasserstein distance:

Lemma 3 (Wasserstein Distance Property). *For any Lipschitz continuous function f with Lipschitz constant L , and for any two probability measures P and Q , we have:*

$$|\mathbb{E}_P[f] - \mathbb{E}_Q[f]| \leq LW_p(P, Q) \quad (34)$$

Proof. This is a well-known result in optimal transport theory. A proof can be found in [7]. \square

Step 2: Lipschitz Continuity of Lyapunov and Barrier Conditions

Let $L_V(x, \xi, \pi^*)$ and $L_h(x, \xi, \pi^*)$ be the Lyapunov and barrier conditions, respectively. We can show that these functions are Lipschitz continuous in ξ :

Lemma 4 (Lipschitz Continuity of Conditions). *There exist constants $L_V, L_h > 0$ such that for all $x \in X, \xi_1, \xi_2 \in \Xi$:*

$$|L_V(x, \xi_1, \pi^*) - L_V(x, \xi_2, \pi^*)| \leq L_V \|\xi_1 - \xi_2\| \quad (35)$$

$$|L_h(x, \xi_1, \pi^*) - L_h(x, \xi_2, \pi^*)| \leq L_h \|\xi_1 - \xi_2\| \quad (36)$$

Proof. This follows from the Lipschitz continuity of f, g, V , and h , and the composition of Lipschitz functions. \square

Step 3: Bounding the Probability Difference

Let $E_V = \{L_V(x, \xi, \pi^*) < 0\}$ and $E_h = \{L_h(x, \xi, \pi^*) > 0\}$. We want to bound $|\tilde{P}(E_V) - P^*(E_V)|$ and $|\tilde{P}(E_h) - P^*(E_h)|$.

Lemma 5 (Probability Difference Bound). *For any event E defined by a Lipschitz condition with constant L , we have:*

$$|\tilde{P}(E) - P^*(E)| \leq 2C(LW_p(\tilde{P}, P^*))^q \quad (37)$$

where $C > 0$ is a constant and $q = p/(p+1)$.

Proof. This follows from the dual form of the Wasserstein distance and the Kantorovich-Rubinstein duality. A detailed proof can be found in [7]. \square

Applying this lemma to our Lyapunov and barrier conditions:

$$|\tilde{P}(E_V) - P^*(E_V)| \leq 2C(L_V W_p(\tilde{P}, P^*))^q \leq 2C(L_V \Delta)^q \quad (38)$$

$$|\tilde{P}(E_h) - P^*(E_h)| \leq 2C(L_h W_p(\tilde{P}, P^*))^q \leq 2C(L_h \Delta)^q \quad (39)$$

Step 4: Combining the Bounds

From the assumptions of the theorem and Theorem 1, we have:

$$P^*(E_V \cap E_h) \geq 1 - (\varepsilon_V + \varepsilon_h + \delta) \quad (40)$$

Now, we can bound the probability under \tilde{P} :

$$\tilde{P}(E_V \cap E_h) = P^*(E_V \cap E_h) + [\tilde{P}(E_V \cap E_h) - P^*(E_V \cap E_h)] \quad (41)$$

$$\geq 1 - (\varepsilon_V + \varepsilon_h + \delta) - |\tilde{P}(E_V) - P^*(E_V)| - |\tilde{P}(E_h) - P^*(E_h)| \quad (42)$$

$$\geq 1 - (\varepsilon_V + \varepsilon_h + \delta) - 2C(L_V \Delta)^q - 2C(L_h \Delta)^q \quad (43)$$

$$\geq 1 - (\varepsilon_V + \varepsilon_h + 2C\Delta^q + \delta) \quad (44)$$

where in the last step we use $L = \max(L_V, L_h)$. The event $E_V \cap E_h$ implies that both the Lyapunov condition and the barrier condition are satisfied. Following the same reasoning as in Theorem 1, this implies that $\|x(t)\| \rightarrow 0$ as $t \rightarrow \infty$ and $x(t) \in S$ for all $t \geq 0$. Therefore,

$$\tilde{P}(\|x(t)\| \rightarrow 0 \text{ as } t \rightarrow \infty \text{ and } x(t) \in S \text{ for all } t \geq 0) \geq 1 - (\varepsilon_V + \varepsilon_h + 2C\Delta^q + \delta) \quad (45)$$

which completes the proof. \square

Implication. Theorem 2, which quantifies the degradation of performance guarantees under distribution shifts, has implications for the deployment of learning-based controllers in real-world scenarios. This result provides a theoretical justification for the empirically observed robustness of distributionally robust methods to changes in operating conditions. The sublinear dependence of the performance degradation on the Wasserstein distance between distributions suggests that DR-LBNs can maintain reasonable performance even under significant shifts. This property is particularly valuable in applications where the true distribution of uncertainties may drift over time or differ between training and deployment environments. However, the theorem also highlights a fundamental trade-off: increasing the radius of the ambiguity set enhances robustness to distribution shifts but may lead to more conservative control policies.

Theorem 3 (Input-to-State Stability). *Consider the system with additive disturbances:*

$$\frac{dx}{dt} = f(x, \xi) + g(x, \xi)\pi^*(x) + w(t) \quad (46)$$

where $\|w(t)\| \leq W$ for all $t \geq 0$. Under the DR-LBN controller π^* , the system is input-to-state stable with probability at least $1 - (\varepsilon_V + \varepsilon_h)$.

Proof. We follow a series of steps:

Step 1: Probabilistic Lyapunov Condition

From the distributionally robust Lyapunov condition of our DR-LBN, we have:

$$\inf_{P \in \mathcal{P}} P(L_V(x, \xi, \pi^*) < 0) \geq 1 - \varepsilon_V, \quad \forall x \in X \setminus \{0\} \quad (47)$$

where $L_V(x, \xi, \pi^*) = \nabla V(x)^\top (f(x, \xi) + g(x, \xi)\pi^*(x)) + \alpha_1(\|x\|)$.

By the definition of infimum and our assumption on the true distribution P^* , this implies:

$$P^*(L_V(x, \xi, \pi^*) < 0) \geq 1 - \varepsilon_V, \quad \forall x \in X \setminus \{0\} \quad (48)$$

Step 2: Lyapunov Function Derivative with Disturbance

Now, consider the time derivative of $V(x)$ along the trajectories of the disturbed system:

$$\dot{V}(x) = \nabla V(x)^\top (f(x, \xi) + g(x, \xi)\pi^*(x) + w(t)) \quad (49)$$

$$= L_V(x, \xi, \pi^*) - \alpha_1(\|x\|) + \nabla V(x)^\top w(t) \quad (50)$$

Step 3: Bounding the Disturbance Term

Using the Cauchy-Schwarz inequality and the bound on $w(t)$:

$$\nabla V(x)^\top w(t) \leq \|\nabla V(x)\| \cdot \|w(t)\| \leq L_V W \quad (51)$$

where L_V is the Lipschitz constant of $V(x)$, which exists due to our neural network construction.

Step 4: Probabilistic ISS Condition

Combining the results from steps 2 and 3, with probability at least $1 - \varepsilon_V$:

$$\dot{V}(x) \leq L_V(x, \xi, \pi^*) - \alpha_1(\|x\|) + L_V W \quad (52)$$

$$< -\alpha_1(\|x\|) + L_V W \quad (53)$$

Now, let $\alpha_3(\|x\|) = \frac{1}{2}\alpha_1(\|x\|)$. Then:

$$\dot{V}(x) < -\alpha_3(\|x\|) \quad \text{whenever} \quad \alpha_1(\|x\|) > 2L_V W \quad (54)$$

Let $\gamma(s) = \alpha_1^{-1}(2L_V s)$. Then:

$$\dot{V}(x) < -\alpha_3(\|x\|) \quad \text{whenever} \quad \|x\| > \gamma(W) \quad (55)$$

Step 5: Safety Consideration

From the distributionally robust barrier condition, we have:

$$P^*(x(t) \in S \text{ for all } t \geq 0) \geq 1 - \varepsilon_h \quad (56)$$

Step 6: Combining Stability and Safety

Using the union bound on the events in (48) and (56), we can conclude that with probability at least $1 - (\varepsilon_V + \varepsilon_h)$, both the ISS condition (55) and the safety condition hold. The condition in (55) is the defining property of input-to-state stability. It shows that the state x is ultimately bounded by a class \mathcal{K} function of the disturbance bound W . Moreover, this bound holds while the system remains in the safe set S . Therefore, we can conclude that under the DR-LBN controller π^* , the system is input-to-state stable with probability at least $1 - (\varepsilon_V + \varepsilon_h)$. \square

Implication. The input-to-state stability result extends the applicability of DR-LBNs to systems subject to bounded external disturbances. While this result is promising, it assumes a known bound on the disturbance magnitude. In practice, characterizing this bound may be challenging, especially for complex, high-dimensional systems.

Theorem 4 (Universal Approximation for DR-LBNs). *For any continuous Lyapunov function $V^*(x)$ and barrier function $h^*(x)$, and for any $\epsilon > 0$, there exists a DR-LBN with a finite number of neurons that can approximate $V^*(x)$ and $h^*(x)$ within an ϵ -error in the supremum norm over any compact subset of the state space.*

Proof. We proceed in several steps:

Step 1: Universal Approximation Theorem for Feedforward Networks

We begin by recalling the classical universal approximation theorem for feedforward neural networks with a single hidden layer:

Lemma 5 ([6]). *Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous sigmoid-type activation function. Then finite sums of the form*

$$G(x) = \sum_{i=1}^N \alpha_i \sigma(w_i^\top x + b_i) \quad (57)$$

are dense in $C(I^n)$ in the topology of uniform convergence, where I^n is the n -dimensional unit cube.

This lemma establishes that a single hidden layer neural network with sigmoid activation can approximate any continuous function on a compact domain to arbitrary accuracy.

Step 2: Extension to Lyapunov and Barrier Functions

Let $X \subset \mathbb{R}^n$ be a compact subset of the state space. Without loss of generality, we can assume X is contained in a hypercube $[-M, M]^n$ for some $M > 0$. Consider the functions $\tilde{V}(x) = V^*(x) - V^*(0)$ and $\tilde{h}(x) = h^*(x)$. Note that $\tilde{V}(0) = 0$, preserving the Lyapunov function property at the origin. By the Stone-Weierstrass theorem, there exist polynomials $p_V(x)$ and $p_h(x)$ such that:

$$\sup_{x \in X} |\tilde{V}(x) - p_V(x)| < \frac{\epsilon}{3} \quad (58)$$

$$\sup_{x \in X} |\tilde{h}(x) - p_h(x)| < \frac{\epsilon}{3} \quad (59)$$

Step 3: Approximation by Neural Networks

By Lemma 1, there exist single hidden layer neural networks $N_V(x)$ and $N_h(x)$ with sigmoid activation functions such that:

$$\sup_{x \in X} |p_V(x) - N_V(x)| < \frac{\epsilon}{3} \quad (60)$$

$$\sup_{x \in X} |p_h(x) - N_h(x)| < \frac{\epsilon}{3} \quad (61)$$

Step 4: Construction of DR-LBN

We now construct our DR-LBN as follows:

$$V_{NN}(x) = N_V(x) + V^*(0) + \delta \|x\|^2 \quad (62)$$

$$h_{NN}(x) = \text{softplus}(N_h(x)) \quad (63)$$

where $\delta > 0$ is a small constant chosen such that $\delta \sup_{x \in X} \|x\|^2 < \frac{\epsilon}{3}$.

Step 5: Error Bounds

For the Lyapunov function approximation:

$$|V^*(x) - V_{NN}(x)| = |\tilde{V}(x) + V^*(0) - N_V(x) - V^*(0) - \delta \|x\|^2| \quad (64)$$

$$\leq |\tilde{V}(x) - p_V(x)| + |p_V(x) - N_V(x)| + \delta \|x\|^2 \quad (65)$$

$$< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon \quad (66)$$

For the barrier function approximation:

$$|h^*(x) - h_{NN}(x)| = |\tilde{h}(x) - \text{softplus}(N_h(x))| \quad (67)$$

$$\leq |\tilde{h}(x) - p_h(x)| + |p_h(x) - N_h(x)| + |N_h(x) - \text{softplus}(N_h(x))| \quad (68)$$

$$< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon \quad (69)$$

The last inequality holds because $|\text{softplus}(z) - z| < 1$ for all z , and we can always scale $N_h(x)$ to make this difference arbitrarily small without changing its approximation properties.

Step 6: Preservation of Lyapunov and Barrier Properties

1. For the Lyapunov function:

- $V_{NN}(0) = V^*(0)$
- $V_{NN}(x) > 0$ for $x \neq 0$ due to the $\delta\|x\|^2$ term

2. For the barrier function:

- $h_{NN}(x) \geq 0$ for all x due to the softplus function

We have constructed a DR-LBN with a finite number of neurons that approximates both $V^*(x)$ and $h^*(x)$ within ϵ -error in the supremum norm over the compact set X , while preserving the essential properties of Lyapunov and barrier functions. \square

Implication. Theorem 4, which establishes the universal approximation capabilities of DR-LBNs, provides theoretical justification for their use in approximating complex Lyapunov and barrier functions. This result suggests that, with sufficient capacity, DR-LBNs can represent a wide class of stability certificates and safety constraints. However, the theorem is existential in nature and does not provide guarantees on the learnability of these functions or the sample complexity required to achieve a given approximation error. Moreover, the potential need for large network architectures to achieve good approximations may pose computational challenges in high-dimensional state spaces.

Theorem 5 (Sample Complexity). *For any $\delta > 0$, if the number of samples N satisfies*

$$N \geq \mathcal{O}\left(\frac{d}{\epsilon^2} \log \frac{1}{\delta}\right), \quad (70)$$

where d is the VC-dimension of the hypothesis class, then with probability at least $1 - \delta$, the true distribution P^* lies within the Wasserstein ball of radius ρ around the empirical distribution P_N .

Proof. We proceed in several steps:

Step 1: Wasserstein Distance and Empirical Processes

Let $W_p(P^*, P_N)$ denote the p -Wasserstein distance between the true distribution P^* and the empirical distribution P_N . We start by relating this to an empirical process:

$$W_p^p(P^*, P_N) = \sup_{f \in \mathcal{F}} \left| \int f dP^* - \int f dP_N \right| \quad (71)$$

where \mathcal{F} is the set of 1-Lipschitz functions.

Step 2: Symmetrization

Introduce an independent copy P'_N of the empirical measure. By Jensen's inequality:

$$\mathbb{E}[W_p^p(P^*, P_N)] \leq 2\mathbb{E}\left[\sup_{f \in \mathcal{F}} \left| \int f dP_N - \int f dP'_N \right| \right] \quad (72)$$

Step 3: Rademacher Complexity

Let $\{\sigma_i\}_{i=1}^N$ be i.i.d. Rademacher random variables. By the symmetrization principle:

$$\mathbb{E}\left[\sup_{f \in \mathcal{F}} \left| \int f dP_N - \int f dP'_N \right| \right] \leq 2\mathbb{E}\left[\sup_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \sigma_i f(X_i)\right] \quad (73)$$

The right-hand side is known as the Rademacher complexity $\mathcal{R}_N(\mathcal{F})$.

Step 4: Dudley's Entropy Integral

We can bound the Rademacher complexity using Dudley's entropy integral [14]:

$$\mathcal{R}_N(\mathcal{F}) \leq \frac{12}{\sqrt{N}} \int_0^\infty \sqrt{\log \mathcal{N}(\mathcal{F}, \varepsilon)} d\varepsilon \quad (74)$$

where $\mathcal{N}(\mathcal{F}, \varepsilon)$ is the ε -covering number of \mathcal{F} .

Step 5: VC-dimension and Covering Numbers

For a function class \mathcal{F} with VC-dimension d , we have the following bound on the covering number:

$$\log \mathcal{N}(\mathcal{F}, \varepsilon) \leq cd \log\left(\frac{1}{\varepsilon}\right) \quad (75)$$

for some constant c .

Step 6: Bounding the Integral

Substituting this into Dudley's integral:

$$\mathcal{R}_N(\mathcal{F}) \leq \frac{12}{\sqrt{N}} \int_0^\infty \sqrt{cd \log\left(\frac{1}{\varepsilon}\right)} d\varepsilon \quad (76)$$

$$\leq C \sqrt{\frac{d}{N}} \quad (77)$$

for some constant C .

Step 7: Concentration Inequality

By McDiarmid's inequality, for any $t > 0$:

$$P(W_p^p(P^*, P_N) - \mathbb{E}[W_p^p(P^*, P_N)] \geq t) \leq \exp\left(-\frac{2Nt^2}{L^2}\right) \quad (78)$$

where L is the Lipschitz constant of the loss function.

Step 8: Putting It All Together

Combining the bounds from steps 6 and 7, we have:

$$P(W_p^p(P^*, P_N) \geq C \sqrt{\frac{d}{N}} + t) \leq \exp\left(-\frac{2Nt^2}{L^2}\right) \quad (79)$$

Setting the right-hand side to δ and solving for N :

$$N \geq \mathcal{O} \left(\frac{d}{\varepsilon^2} \log \frac{1}{\delta} \right) \quad (80)$$

where $\varepsilon = C \sqrt{\frac{d}{N}} + t$ is the desired Wasserstein ball radius. Therefore, if the number of samples N satisfies the given inequality, then with probability at least $1 - \delta$, the true distribution P^* lies within the Wasserstein ball of radius $\rho = \varepsilon$ around the empirical distribution P_N . \square

Implication. The sample complexity bound provided in Theorem 5 offers insights into the data requirements for learning DR-LBNs with probabilistic guarantees. This result is important for understanding the practical feasibility of the approach, especially in data-scarce environments. The logarithmic dependence on the confidence parameter is favorable, allowing for high-probability guarantees without excessive data requirements. However, the quadratic dependence on the inverse of the desired accuracy suggests that achieving very high precision may require a substantial number of samples. It is worth noting that this bound is based on worst-case analysis and may be pessimistic for many practical scenarios.

6 Conclusion and Future Work

This paper has proposed distributionally robust Lyapunov-Barrier networks (DR-LBNs), a framework for synthesizing controllers that ensure both stability and safety for nonlinear systems under parametric uncertainty. By using recent advances in distributionally robust optimization and neural network architectures, our approach provides probabilistic guarantees on system performance even when the true distribution of uncertainties is unknown. The core contributions of this work lie in the development of a unified neural network architecture that simultaneously learns Lyapunov functions, barrier functions, and adaptive control gains. This integration enables the synthesis of controllers that balance stability and safety objectives in a principled manner.

Future Directions. While this paper has focused on establishing the theoretical foundations of DR-LBNs, several important directions for future research emerge. First, the development of comprehensive numerical examples and case studies is important to validate the practical efficacy of DR-LBNs. Such examples should span a range of nonlinear systems, from low-dimensional benchmarks like the inverted pendulum to high-dimensional, safety-critical applications such as autonomous vehicles or multi-agent robotics. Beyond numerical validation, future work should explore the online adaptation of DR-LBNs to time-varying uncertainties, potentially incorporating techniques from adaptive control and online learning. The extension of the framework to handle partial state observations and measurement noise is another critical area for investigation. Additionally, the integration of DR-LBNs with model predictive control frameworks could yield powerful tools for long-horizon planning under uncertainty. From a theoretical perspective, characterizing the finite-sample behavior of DR-LBNs and developing tighter generalization bounds remain open challenges.

References

- [1] Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European control conference (ECC)*, pages 3420–3431. IEEE, 2019.
- [2] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2016.
- [3] Zvi Artstein. Stabilization with relaxed controls. *Nonlinear Analysis: Theory, Methods & Applications*, 7(11):1163–1173, 1983.
- [4] Lukas Brunke, Melissa Greeff, Adam W Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.
- [5] Chun Cheng, Qinxiao Yu, Yossiri Adulyasak, and Louis-Martin Rousseau. Distributionally robust facility location with uncertain facility capacity and customer demand. *Omega*, 122:102959, 2024.
- [6] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- [7] Nicolas Fournier and Arnaud Guillin. On the rate of convergence in wasserstein distance of the empirical measure. *Probability theory and related fields*, 162(3):707–738, 2015.
- [8] Virginie Gabrel, Cécile Murat, and Aurélie Thiele. Recent advances in robust optimization: An overview. *European journal of operational research*, 235(3):471–483, 2014.
- [9] Kunal Garg, James Usevitch, Joseph Breeden, Mitchell Black, Devansh Agrawal, Hardik Parwana, and Dimitra Panagou. Advances in the theory of control barrier functions: Addressing practical challenges in safe control synthesis for autonomous and robotic systems. *Annual Reviews in Control*, 57:100945, 2024.
- [10] Rafal Goedel, Ricardo G Sanfelice, and Andrew R Teel. Hybrid dynamical systems: modeling stability, and robustness. *Princeton, NJ, USA*, 2012.
- [11] Dimitar Ho, Hoang Le, John Doyle, and Yisong Yue. Online robust control of nonlinear systems with large uncertainty. In *International Conference on Artificial Intelligence and Statistics*, pages 3475–3483. PMLR, 2021.
- [12] Shao-Chen Hsu, Xiangru Xu, and Aaron D Ames. Control barrier function based quadratic programs with application to bipedal robotic walking. In *2015 American Control Conference (ACC)*, pages 4542–4548. IEEE, 2015.
- [13] Ripeng Huang, Shaojian Qu, Xiaoguang Yang, and Zhimin Liu. Multi-stage distributionally robust optimization with risk aversion. *Journal of Industrial & Management Optimization*, 17(1), 2021.

- [14] Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. *Advances in neural information processing systems*, 21, 2008.
- [15] Kehan Long, Yinzhuang Yi, Jorge Cortés, and Nikolay Atanasov. Distributionally robust Lyapunov function search under uncertainty. In *Learning for Dynamics and Control Conference*, pages 864–877. PMLR, 2023.
- [16] Brett T Lopez, Jean-Jacques E Slotine, and Jonathan P How. Robust adaptive control barrier functions: An adaptive and data-driven approach to safety. *IEEE Control Systems Letters*, 5(3):1031–1036, 2020.
- [17] Ruozhen Qiu, Yue Sun, and Minghe Sun. A distributionally robust optimization approach for multi-product inventory decisions with budget constraint and demand and yield uncertainties. *Computers & Operations Research*, 126:105081, 2021.
- [18] Svetlozar Todorov Rachev et al. Duality theorems for Kantorovich-Rubinstein and Wasserstein functionals. 1990.
- [19] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.
- [20] Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning control Lyapunov functions from counterexamples and demonstrations. *Autonomous Robots*, 43:275–307, 2019.
- [21] Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):253–279, 2019.
- [22] Spencer M Richards, Felix Berkenkamp, and Andreas Krause. The Lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems. In *Conference on Robot Learning*, pages 466–476. PMLR, 2018.
- [23] Eduardo D Sontag. A ‘universal’construction of artstein’s theorem on nonlinear stabilization. *Systems & control letters*, 13(2):117–123, 1989.
- [24] Andrew J Taylor, Victor D Dorobantu, Hoang M Le, Yisong Yue, and Aaron D Ames. Episodic learning with control Lyapunov functions for uncertain robotic systems. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6878–6884. IEEE, 2019.
- [25] Wei Xiao and Calin Belta. Control barrier functions for systems with high relative degree. In *2019 IEEE 58th conference on decision and control (CDC)*, pages 474–479. IEEE, 2019.
- [26] Xiangru Xu, Paulo Tabuada, Jessy W Grizzle, and Aaron D Ames. Robustness of control barrier functions for safety critical control. *IFAC-PapersOnLine*, 48(27):54–61, 2015.
- [27] Songyuan Zhang, Oswin So, Kunal Garg, and Chuchu Fan. Gcbf+: A neural graph control barrier function framework for distributed safe multi-agent control. *arXiv preprint arXiv:2401.14554*, 2024.
- [28] Leyang Zhao, Guoyin Li, and Spiridon Penev. Regularized distributionally robust optimization with application to the index tracking problem. *Annals of Operations Research*, pages 1–28, 2024.