

Distributionally Robust Lyapunov Function Search Under Uncertainty

Kehan Long
Yinzhuang Yi
Jorge Cortés
Nikolay Atanasov

K3LONG@UCSD.EDU
YIYI@UCSD.EDU
CORTES@UCSD.EDU
NATANASOV@UCSD.EDU

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

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

This paper develops methods for proving Lyapunov stability of dynamical systems subject to disturbances with an unknown distribution. We assume only a finite set of disturbance samples is available and that the true online disturbance realization may be drawn from a different distribution than the given samples. We formulate an optimization problem to search for a sum-of-squares (SOS) Lyapunov function and introduce a distributionally robust version of the Lyapunov function derivative constraint. We show that this constraint may be reformulated as several SOS constraints, ensuring that the search for a Lyapunov function remains in the class of SOS polynomial optimization problems. For general systems, we provide a distributionally robust chance-constrained formulation for neural network Lyapunov function search. Simulations demonstrate the validity and efficiency of either formulation on non-linear uncertain dynamical systems.

Keywords: Lyapunov stability, distrib. robust optimization, sum of squares, neural networks.

Supplementary Material

Open-source implementation available at <https://github.com/KehanLong/DR-Lyapunov-Function>.

1. Introduction

A Lyapunov function (LF) is one of the main tools for analyzing the stability of nonlinear dynamical systems (Khalil, 1996). Similarly, control synthesis for open-loop control-affine systems is often done using a control Lyapunov function (CLF) (Artstein, 1983) since a stabilizing controller can be obtained from a CLF using a universal formula (Sontag, 1989a) or quadratic programming (Galloway et al., 2015). Various techniques exist for obtaining LF or CLF candidates but the majority assume that the system model is known. In this paper, we study the problem of synthesizing a Lyapunov function when the system model is uncertain.

Synthesizing a valid LF for a linear system can be formulated as a semi-definite program (SDP) (Boyd et al., 1994). Parrilo (2000); Papachristodoulou and Prajna (2002) generalized the formulation for non-linear polynomial systems by using sum-of-squares (SOS) polynomials to represent an LF candidate. For polynomial systems with uncertainty, Ahmadi and Majumdar (2016); Lasserre (2015) extended SOS techniques to find robust LFs based on known error bounds, see Laurent (2009) for a general exposition. The lack of a valid SOS LF does not imply that the system instability since there exist positive-definite functions that are not representable as SOS (Hilbert, 1888).

Using a neural network as a more general LF representation than an SOS polynomial has been gaining increasing popularity. Richards et al. (2018) proposed a neural network approach

for discrete-time non-linear systems that learns the region of attraction of a given controller. [Boffi et al. \(2020\)](#); [Gaby et al. \(2022\)](#) improved the efficiency of learning LFs by incorporating positive-definiteness and equilibrium conditions directly into the network architecture. [Chang et al. \(2019\)](#); [Dai et al. \(2021\)](#) considered learning neural network CLFs and controllers by minimizing violations of the conditions for a valid LF. [Dawson et al. \(2022b\)](#) extended the idea to learn safety certificates as control Lyapunov-barrier functions and also considered control-affine systems with convex-hull uncertainty. The survey by [Dawson et al. \(2022a\)](#) provides a recent account of this line of research. Despite their expressiveness, neural network methods do not offer theoretical guarantees for the validity of the learned Lyapunov function over the entire state space.

The stability guarantees provided by either SOS or neural network LFs are sensitive to uncertainty or disturbances in the system model. A related body of work ([Choi et al., 2020](#); [Taylor et al., 2019](#); [Castañeda et al., 2021](#); [Dhiman* et al., 2021](#); [Long et al., 2022a](#)) assumes that an LF certificate is given for a nominal system and develops approaches to adapt it by taking the model uncertainty into account during deployment. In input-to-state stability (ISS) ([Sontag, 1989b](#)), one deals directly with an uncertain dynamical system to provide robustness guarantees on graceful degradation of stability as a function of the disturbance input magnitude. This can be ensured via an ISS-Lyapunov function ([Sontag and Wang, 1995](#)), e.g., constructed using SOS techniques ([Hespanha et al., 2008](#); [Voßwinkel and Röbenack, 2020](#)). Without assuming any known distribution or error bounds on the model uncertainty, in this paper, we utilize distributionally robust constraints ([Shapiro et al., 2014](#); [Esfahani and Kuhn, 2018](#)) to enforce LF conditions for an uncertain system model with only finitely many uncertainty samples obtained offline.

Distributionally robust chance-constrained programming (DRCCP) deals with uncertain variables in the constraints using finitely many available samples. The main idea is to construct an ambiguity ball centered at the empirical distribution of the observed samples using a distribution distance function, such as Wasserstein distance ([Esfahani and Kuhn, 2018](#); [Hota et al., 2019](#)). Then, the constraints are required to be satisfied with high probability for all distributions in the ambiguity ball. Given its powerful guarantee to handle uncertainty with an unknown or shifting distribution, DRCCP has been applied in several areas in systems and control ([Coulson et al., 2019](#); [Boskos et al., 2021](#); [Long et al., 2022b](#); [Coppens et al., 2020](#)). However, its application to LF search here is novel.

Contributions: 1) We formulate a distributionally robust version of the Lyapunov function derivative constraint for uncertain dynamical systems using finitely many offline samples. 2) For polynomial systems, we show that the distributionally robust constraint can be reformulated as multiple SOS constraints, ensuring that LF synthesis with uncertainty remains an SOS polynomial optimization. 3) For general nonlinear systems, we propose a distributionally robust neural network approach for learning Lyapunov functions.

2. Background

Here¹, we give an overview of sum-of-squares (SOS) techniques for Lyapunov function synthesis and distributionally robust chance constraints.

1. The sets of non-negative real and natural numbers are denoted $\mathbb{R}_{\geq 0}$ and \mathbb{N} . For $N \in \mathbb{N}$, $[N] := \{1, 2, \dots, N\}$. We denote the distribution and expectation of a random variable Y by \mathbb{P} and $\mathbb{E}_{\mathbb{P}}(Y)$, resp. For a scalar x , $(x)_+ := \max(x, 0)$. We use $\mathbf{0}$ to denote the n -dimensional vector with all entries equal to 0. The gradient of a differentiable function V is denoted by ∇V , and its Lie derivative along a vector field f by $\mathcal{L}_f V = \nabla V \cdot f$. We denote a uniform distribution on $[a, b]$ as $\mathcal{U}(a, b)$ and a Gaussian distribution with mean μ and variance σ^2 as $\mathcal{N}(\mu, \sigma^2)$.

2.1. Lyapunov Theory and Sum-of-Squares Optimization

Consider a dynamical system, $\dot{\mathbf{x}} = f(\mathbf{x})$, with state $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$. Assume $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ is locally Lipschitz and the origin $\mathbf{x} = \mathbf{0}$ is the desired equilibrium, i.e., $f(\mathbf{0}) = \mathbf{0}$. A valid Lyapunov function, ensuring the stability of the origin, satisfies:

$$V(\mathbf{0}) = 0, V(\mathbf{x}) > 0 \text{ and } \dot{V}(\mathbf{x}) < 0, \forall \mathbf{x} \neq \mathbf{0}, \quad (1)$$

where $\dot{V}(\mathbf{x}) = \mathcal{L}_f V(\mathbf{x})$. If the LF is also radially unbounded ($V(\mathbf{x}) \rightarrow \infty$ as $\|\mathbf{x}\| \rightarrow \infty$), then its existence implies global asymptotic stability. The second and third conditions in (1) are implied by

$$V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0 \text{ and } -\dot{V}(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0, \forall \mathbf{x} \neq \mathbf{0}, \quad (2)$$

for some $\epsilon \in \mathbb{R}_{>0}$. A natural way of imposing non-negativity is by using SOS polynomials. A polynomial $\eta(\mathbf{x})$ of degree $2d$ is called an SOS polynomial if and only if there exist polynomials $s_1(\mathbf{x}), \dots, s_p(\mathbf{x})$ of degree at most d such that $\eta(\mathbf{x}) = \sum_{i=1}^p s_i(\mathbf{x})^2$. Based on the positive-definiteness property of SOS polynomials, Parrilo (2000); Papachristodoulou and Prajna (2002) proposed the following SOS conditions, which are sufficient to imply (1),

$$V(\mathbf{x}) = \sum_{k=0}^{2d} c_k \mathbf{x}^k, c_0 = 0; \quad V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}); \quad -\dot{V}(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}), \quad (3)$$

where $\text{SOS}(\mathbf{x})$ denotes the set of SOS polynomials in variable \mathbf{x} . By fixing a polynomial degree d , one can search for an SOS LF using a semidefinite program (Laurent, 2009) enforcing (3).

2.2. Conditional Value-at-Risk and Distributionally Robust Chance Constraint

We review chance-constraint formulations that will be useful to handle model uncertainty. Consider a complete separable metric space Ξ with metric d , and associate to it a Borel σ -algebra \mathcal{F} and the set $\mathcal{P}(\Xi)$ of Borel probability measures on Ξ . A chance constraint can be written as,

$$\mathbb{P}^*(G(\mathbf{z}, \boldsymbol{\xi}) \leq 0) \geq 1 - \beta, \quad (4)$$

where the constraint function $G(\mathbf{z}, \boldsymbol{\xi}) \in \mathbb{R}^n \times \Xi \mapsto \mathbb{R}$ depends both on a decision vector \mathbf{z} and a random variable $\boldsymbol{\xi}$ with distribution $\mathbb{P}^* \in \mathcal{P}(\Xi)$, and $\beta \in (0, 1)$ is a user-specified risk tolerance. The feasible set for \mathbf{z} defined by (4) is not convex. (4), Nemirovski and Shapiro (2006) proposed a Conditional Value-at-Risk (CVaR) approximation of the chance constraint, which results in a convex feasible set and is sufficient for (4) to hold:

$$\text{CVaR}_{1-\beta}^{\mathbb{P}^*}(G(\mathbf{z}, \boldsymbol{\xi})) \leq 0. \quad (5)$$

For a random variable $\xi \in \mathbb{R}$ with distribution $\hat{\mathbb{P}}$, the Value-at-risk (VaR) at confidence level $1 - \beta$ is $\text{VaR}_{1-\beta}^{\hat{\mathbb{P}}}(\xi) := \inf_{t \in \mathbb{R}} \{t \mid \hat{\mathbb{P}}(\xi \leq t) \geq 1 - \beta\}$. The CVaR of ξ is $\text{CVaR}_{1-\beta}^{\hat{\mathbb{P}}}(\xi) := \mathbb{E}_{\hat{\mathbb{P}}}[\xi \mid \xi \geq \text{VaR}_{1-\beta}^{\hat{\mathbb{P}}}(\xi)]$ and can be formulated as a convex program (Rockafellar and Uryasev, 2000):

$$\text{CVaR}_{1-\beta}^{\hat{\mathbb{P}}}(\xi) = \inf_{t \in \mathbb{R}} [\beta^{-1} \mathbb{E}_{\hat{\mathbb{P}}}[(\xi + t)_+] - t]. \quad (6)$$

The chance constraint in (4) or (5) cannot be specified if the distribution \mathbb{P}^* of $\boldsymbol{\xi}$ is unknown. In robotics and control applications, it is common that only finitely many samples $\{\boldsymbol{\xi}_i\}_{i=1}^N$ from

\mathbb{P}^* are available. This motivates a distributionally robust formulation of the chance constraint (Es-fahani and Kuhn, 2018; Xie, 2021). 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 (see, for example Chen et al., 2018; Xie, 2021): $W_p(\mu, \nu) := \left(\inf_{\gamma \in \mathcal{Q}(\mu, \nu)} \left[\int_{\Xi \times \Xi} d(x, y)^p d\gamma(x, y) \right]^{\frac{1}{p}}$, where $\mathcal{Q}(\mu, \nu)$ denotes the collection of all measures on $\Xi \times \Xi$ with marginals μ and ν on the first and second factors, and d denotes the metric in Ξ .

We denote by $\hat{\mathbb{P}}_N := \frac{1}{N} \sum_{i=1}^N \delta_{\xi_i}$ the discrete empirical distribution of samples $\{\xi_i\}_{i=1}^N$. Using the Wasserstein distance, we define an ambiguity set $\mathcal{M}_N^r := \{\mu \in \mathcal{P}_p(\Xi) \mid W_p(\mu, \hat{\mathbb{P}}_N) \leq r\}$ as a ball of distributions with radius r centered at $\hat{\mathbb{P}}_N$. We write a distributionally robust version of the chance constraint in (4) as $\inf_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(G(\mathbf{z}, \boldsymbol{\xi}) \leq 0) \geq 1 - \beta$ or equivalently $\sup_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(G(\mathbf{z}, \boldsymbol{\xi}) \geq 0) \leq \beta$. Thus, similar to the CVaR approximation in (5)-(6), one considers the sufficient constraint $\sup_{\mathbb{P} \in \mathcal{M}_N^r} \inf_{t \in \mathbb{R}} [\beta^{-1} \mathbb{E}_{\mathbb{P}}[(G(\mathbf{z}, \boldsymbol{\xi}) + t)_+] - t] \leq 0$, which is convex in \mathbf{z} .

3. Problem Formulation

We aim to analyze Lyapunov stability for a dynamical system subject to model uncertainty:

$$\dot{\mathbf{x}} = f(\mathbf{x}) + \sum_{i=1}^m d_i(\mathbf{x}) \xi_i = f(\mathbf{x}) + d(\mathbf{x}) \boldsymbol{\xi}, \quad (7)$$

where $d_i : \mathbb{R}^n \mapsto \mathbb{R}^n$ is locally Lipschitz. We assume that $d(\mathbf{x}) = [d_1(\mathbf{x}), \dots, d_m(\mathbf{x})] \in \mathbb{R}^{n \times m}$ is known or estimated from state-control trajectories (Harrison et al., 2020; Duong and Atanasov, 2022). We do not assume any known error bounds or distribution for the parameter $\boldsymbol{\xi} \in \Xi \subseteq \mathbb{R}^m$. Instead, we consider a finite data set of samples $\{\xi_i\}_{i=1}^N$ that may be used for LF synthesis. The uncertainty model in (7) captures the commonly considered additive disturbance, which in our formulation corresponds to $m = n$ and $d(\mathbf{x}) = \mathbf{I}_n$. The matrix $d(\mathbf{x})$ allows specifying particular system modes affected by the disturbance $\boldsymbol{\xi}$ depending on the state \mathbf{x} .

Problem 1 (Lyapunov Function Search For Uncertain Systems) *Given a finite set of uncertainty samples $\{\xi_i\}_{i=1}^N$ from the uncertain system in (7), obtain a Lyapunov function $V : \mathbb{R}^n \mapsto \mathbb{R}$ that can be used to verify the stability of the origin while taking the uncertainty into account.*

4. Lyapunov Function Search For Systems with Model Uncertainty

We present an SOS approach (Sec. 4.1) and a neural network approach (Sec. 4.2) to address Problem 1. Our methodology is based on finding a function $V : \mathbb{R}^n \mapsto \mathbb{R}$ that satisfies the Lyapunov conditions in (1). The uncertainty in the dynamical system (7) appears in the term $\dot{V}(\mathbf{x})$, which presents a challenge for ensuring that the condition $\dot{V}(\mathbf{x}) < 0, \forall \mathbf{x} \neq \mathbf{0}$ is satisfied.

4.1. Sum-of-Squares Approach For Lyapunov Function Search

We first introduce our SOS approach for LF synthesis under model uncertainty. The Lyapunov conditions in (2), taking the uncertainty in (7) into account, become:

$$V(\mathbf{0}) = 0; \forall \mathbf{x} \neq \mathbf{0}, V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0 \text{ and } \mathbb{P}^*(-\dot{V}(\mathbf{x}, \boldsymbol{\xi}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0) \geq 1 - \beta, \quad (8)$$

where \mathbb{P}^* denotes the true distribution of $\boldsymbol{\xi}$.

To simplify the presentation, let $G(\mathbf{x}, \boldsymbol{\xi}) = \dot{V}(\mathbf{x}, \boldsymbol{\xi}) + \epsilon \|\mathbf{x}\|_2^2 = \nabla V(\mathbf{x})^\top (f(\mathbf{x}) + d(\mathbf{x})\boldsymbol{\xi}) + \epsilon \|\mathbf{x}\|_2^2$, so that the chance-constraint in (8) becomes $\mathbb{P}^*(-G(\mathbf{x}, \boldsymbol{\xi}) \geq 0) \geq 1 - \beta, \forall \mathbf{x} \neq \mathbf{0}$. Based on the discussion in Sec. 2.2, the CVaR approximation provides a sufficient condition for enforcing the chance constraint: $\inf_{t \in \mathbb{R}} [\beta^{-1} \mathbb{E}_{\mathbb{P}^*} [(G(\mathbf{x}, \boldsymbol{\xi}) + t)_+] - t] \leq 0$, for all $\mathbf{x} \neq \mathbf{0}$. If the true distribution \mathbb{P}^* were known, this formulation could be used to deal with the uncertainty. However, we are only provided with samples $\{\boldsymbol{\xi}_i\}_{i=1}^N$ from \mathbb{P}^* . We thus rewrite the condition by multiplying by β on both sides and using the empirical expectation to approximate the true expectation,

$$\inf_{t \in \mathbb{R}} \left[\frac{1}{N} \sum_{i=1}^N (G(\mathbf{x}, \boldsymbol{\xi}_i) + t)_+ - t\beta \right] \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}. \quad (9)$$

Due to the infimum term in the constraint, one cannot directly write (9) as an SOS condition, as in (3). The following result provides an alternative SOS condition that ensures (9) holds.

Proposition 1 (CC-SOS Condition) *Assume $\beta \leq \frac{1}{N}$, the constraint in (9) is equivalent to:*

$$\max_i \beta (\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) + \epsilon \|\mathbf{x}\|_2^2) \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}, \quad (10)$$

Furthermore, if f and d_i are polynomials, the following N SOS conditions are sufficient for (10),

$$-\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}), \quad \forall i = 1, 2, \dots, N. \quad (11)$$

Proof Denote by t^* the value when the infimum is attained in (9). Without loss of generality, we assume that for a given \mathbf{x} , $G(\mathbf{x}, \boldsymbol{\xi}_i) \geq G(\mathbf{x}, \boldsymbol{\xi}_j)$, for all $1 \leq i < j \leq N$. Observe that for each $\mathbf{x} \neq \mathbf{0}$, the function $\frac{1}{N} \sum_{i=1}^N (G(\mathbf{x}, \boldsymbol{\xi}_i) + t)_+ - t\beta$ is piecewise-linear in t with $N + 1$ intervals and N breakpoints, given by $\{-G(\mathbf{x}, \boldsymbol{\xi}_i)\}_{i=1}^N$ and the slope for the i -th interval is $\frac{i-1}{N} - \beta$. Thus, the optimal solution is $t^* = -G(\mathbf{x}, \boldsymbol{\xi}_k)$, where k satisfies $\frac{k-1}{N} - \beta < 0$ and $\frac{k}{N} - \beta \geq 0$. The constraint in (9) can be rewritten as $\frac{1}{N} \sum_{i=1}^k (G(\mathbf{x}, \boldsymbol{\xi}_i) - G(\mathbf{x}, \boldsymbol{\xi}_k)) + \beta G(\mathbf{x}, \boldsymbol{\xi}_k) \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}$. Since $\beta \leq \frac{1}{N}$, only the first interval has negative slope and this constraint can be written as (10). Inspired by the SOS formulation in (3), (10) is implied by the N SOS constraints in (11). \blacksquare

Using Proposition 1, we propose a chance-constrained (CC)-SOS formulation to search for a valid Lyapunov function for the uncertain system in (7):

$$V(\mathbf{x}) = \sum_{k=0}^{2d} c_k \mathbf{x}^k, \quad c_0 = 0; \quad V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}); \quad -\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}), \quad (12)$$

for all $i \in [N]$. Note that by using CVaR approximations in (9) and assuming $\beta \leq \frac{1}{N}$, the CC-SOS formulation becomes equivalent to the formulation that is robust against the provided samples $\{\boldsymbol{\xi}_i\}_{i=1}^N$, as shown in (11). This CC-SOS formulation overcomes the lack of knowledge of the true uncertainty distribution \mathbb{P}^* by using the available samples $\boldsymbol{\xi}_i$ to conservatively approximate the probabilistic constraint in (8) with N SOS conditions. Nonetheless, the test-time validity of a Lyapunov function satisfying (12) is not guaranteed because the CC-SOS condition does not account for the error between the empirical $\hat{\mathbb{P}}_N$ and the true \mathbb{P}^* distributions. Moreover, the distribution \mathbb{P}^* that generates the uncertainty samples may change at deployment time. This motivates the following distributionally robust chance-constrained formulation:

$$V(\mathbf{0}) = 0; \quad \forall \mathbf{x} \neq \mathbf{0}, \quad V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0 \quad \text{and} \quad \inf_{\mathbb{P} \in \mathcal{M}_N^r} \mathbb{P}(-\dot{V}(\mathbf{x}, \boldsymbol{\xi}) - \epsilon \|\mathbf{x}\|_2^2 \geq 0) \geq 1 - \beta, \quad (13)$$

where \mathcal{M}_N^r denotes the Wasserstein ambiguity set around the empirical distribution $\hat{\mathbb{P}}_N$ with user-defined radius r . Based on the discussion in Sec. 2.2, the following constraint is a sufficient condition for the distributionally robust chance constraint in (13) to hold,

$$\sup_{\mathbb{P} \in \mathcal{M}_N^r} \inf_{t \in \mathbb{R}} [\mathbb{E}_{\mathbb{P}}[G(\mathbf{x}, \boldsymbol{\xi}) + t]_+ - t\beta] \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}. \quad (14)$$

As before, (14) is not amenable to a SOS formulation. The following result presents SOS conditions which are sufficient to ensure that (14) holds.

Proposition 2 (DRCC-SOS Condition) *Assume $\beta \leq \frac{1}{N}$, consider the 1-Wasserstein distance with L_1 norm as the metric d . The following is a sufficient condition for (14) to hold,*

$$r \max_{1 \leq j \leq k} |\nabla V(\mathbf{x})^\top d_j(\mathbf{x})| + \max_i \beta(\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) + \epsilon \|\mathbf{x}\|_2^2) \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}, \quad (15)$$

where $\nabla V(\mathbf{x})^\top d_j(\mathbf{x})$ denotes the j -th element of the row vector. If $\Xi = \mathbb{R}^m$, then (15) is equivalent to (14). Also, if f and d_i are polynomials, (15) is implied by the following SOS conditions,

$$\pm r \nabla V(\mathbf{x})^\top d_j(\mathbf{x}) - \beta(\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) - \epsilon \|\mathbf{x}\|_2^2) \in \text{SOS}(\mathbf{x}), \quad \forall i = 1, 2, \dots, N, \quad \forall j = 1, 2, \dots, m. \quad (16)$$

Proof Based on (Hota et al., 2019, Lemma V.8) and (Esfahani and Kuhn, 2018, Theorem 6.3), the supremum over the Wasserstein ambiguity set, i.e., condition (14), can be written conservatively as the sample average $\inf_{t \in \mathbb{R}} \left[\frac{1}{N} \sum_{i=1}^N (G(\mathbf{x}, \boldsymbol{\xi}_i) + t)_+ - t\beta \right]$ and a regularization term $rL_G(\mathbf{x})$, where $L_G(\mathbf{x}) : \mathcal{X} \mapsto \mathbb{R}_{>0}$ is the Lipschitz constant of $G(\mathbf{x}, \boldsymbol{\xi})$ in $\boldsymbol{\xi}$. If $\Xi = \mathbb{R}^m$, then (14) is equivalent to the sample average plus $rL_G(\mathbf{x})$. Since the Lipschitz constant of a differentiable affine function equals the dual norm of its gradient, and the dual norm of the L_1 norm is the L_∞ norm, we can define the convex function $L_G : \mathcal{X} \mapsto \mathbb{R}_{>0}$ as $L_G(\mathbf{x}) = \|\nabla V(\mathbf{x})^\top d(\mathbf{x})\|_\infty = \max_{1 \leq j \leq m} |\nabla V(\mathbf{x})^\top d_j(\mathbf{x})|$, which satisfies the property that $\boldsymbol{\xi} \mapsto G(\mathbf{x}, \boldsymbol{\xi})$ is Lipschitz in $\boldsymbol{\xi}$ with Lipschitz constant $L_G(\mathbf{x})$. With the assumption that $\beta \leq \frac{1}{N}$, we use Proposition 1 and conclude that (15) is a sufficient condition for (14) and they are equivalent if $\Xi = \mathbb{R}^m$. Finally, inspired by the SOS relaxations of (1) to (3), we can relax (15) to the $2Nm$ SOS constraints in (16). ■

Based on Proposition 2, we propose a DRCC-SOS formulation to find a Lyapunov function,

$$\begin{aligned} V(\mathbf{x}) &= \sum_{k=0}^{2d} c_k \mathbf{x}^k, \quad c_0 = 0; \quad V(\mathbf{x}) - \epsilon \|\mathbf{x}\|_2^2 \in \text{SOS}(\mathbf{x}); \\ &\pm r [\nabla V(\mathbf{x})^\top d_j(\mathbf{x}) - \beta(\dot{V}(\mathbf{x}, \boldsymbol{\xi}_i) - \epsilon \|\mathbf{x}\|_2^2)] \in \text{SOS}(\mathbf{x}), \quad \forall i = 1, 2, \dots, N, \quad \forall j = 1, 2, \dots, m. \end{aligned} \quad (17)$$

The next result identifies conditions under which the resulting Lyapunov function solves Problem 1.

Proposition 3 (Stability guarantee of DRCC-SOS formulation) *Let the distribution \mathbb{P}^* of $\boldsymbol{\xi}$ in (7) be light-tailed, i.e., there exists an exponent ρ such that $C := \mathbb{E}_{\mathbb{P}^*}[\exp(\|\boldsymbol{\xi}\|^\rho)] < \infty$. Let the Wasserstein radius r^* be given by:*

$$r_N^*(\alpha) := \begin{cases} \left(\frac{\log(c_1 \alpha^{-1})}{c_2 N} \right)^{1/\max(m,2)}, & N \geq \frac{\log(c_1 \alpha^{-1})}{c_2}, \\ \left(\frac{\log(c_1 \alpha^{-1})}{c_2 N} \right)^{1/\rho}, & N < \frac{\log(c_1 \alpha^{-1})}{c_2}, \end{cases} \quad (18)$$

for $N \geq 1$, $m \neq 2$, and $\alpha \in (0, 1)$ being a user-specified risk parameter. The constants c_1, c_2 are positive and only depend on ρ, C and m . Under those conditions, the Lyapunov function obtained from the DRCC-SOS formulation (17) satisfies $\mathbb{P}^*(-\dot{V}(\mathbf{x}, \boldsymbol{\xi}) - \epsilon\|\mathbf{x}\|_2^2 \geq 0) \geq 1 - \alpha - \beta$.

Proof Consider the events $A := \{\mathbb{P}^* \in \mathcal{M}_N^{r^*}\}$ and $B := \{-\dot{V}(\mathbf{x}, \boldsymbol{\xi}) - \epsilon\|\mathbf{x}\|_2^2 \geq 0\}$. Note that $\mathbb{P}^*(A \cap B) = \mathbb{P}^*(A) + \mathbb{P}^*(B) - \mathbb{P}^*(A \cup B) \geq \mathbb{P}^*(A) + \mathbb{P}^*(B) - 1$. On the one hand, we have from (Esfahani and Kuhn, 2018, Theorem 3.4) that, under (18), $\mathbb{P}^*(\mathbb{P}^* \in \mathcal{M}_N^{r^*}) \geq 1 - \alpha$. On the other, from Proposition 2, the LF resulting from (17) with r^* satisfies $\inf_{\mathbb{P} \in \mathcal{M}_N^{r^*}} \mathbb{P}(-\dot{V}(\mathbf{x}, \boldsymbol{\xi}) - \epsilon\|\mathbf{x}\|_2^2 \geq 0) \geq 1 - \beta$. Hence, $\mathbb{P}^*(A \cap B) \geq 1 - \alpha + 1 - \beta - 1 = 1 - \alpha - \beta$. ■

The DRCC-SOS formulation (17) provides a stability guarantee (Proposition 3) if there is no uncertainty distributional shift, i.e., \mathbb{P}^* does not shift outside of $\mathcal{M}_N^{r^*}$ at deployment time. However, similar to other SOS approaches, the formulation is restricted to polynomial systems and the non-existence of an SOS LF does not imply the non-existence of other valid LFs. This motivates us to consider next a more general candidate LF candidate, represented as a neural network.

4.2. Neural Network Approach For Lyapunov Function Search

We propose a neural network approach that encourages the satisfaction of Lyapunov conditions by minimizing a loss function that quantifies their violation. Consider a neural network Lyapunov function (NN-LF) representation of the form $V_\theta(\mathbf{x}) := \|\phi_\theta(\mathbf{x}) - \phi_\theta(\mathbf{0})\|^2 + \hat{\alpha}\|\mathbf{x}\|^2$, where $\phi_\theta : \mathbb{R}^m \mapsto \mathbb{R}$ is a fully-connected neural network with parameters θ and tanh activations, and $\hat{\alpha}$ is a user-chosen parameter (Gaby et al., 2022). By construction, this function is positive definite and $V_\theta(\mathbf{0}) = 0$. We obtain a training set $\mathcal{D}_{\text{LF}} := \{\mathbf{x}_i\}_{i=1}^M$ by sampling uniformly from the domain of interest \mathcal{X}_δ and then minimize the following empirical loss function:

$$\ell_{\text{LF}}(\theta) = \frac{1}{M} \sum_{i=1}^M (\dot{V}_\theta(\mathbf{x}_i) + \gamma\|\mathbf{x}_i\|)_+, \quad (19)$$

where γ is user-defined. This loss encourages a decrease of V_θ along the system trajectories. To deal with the model uncertainty in (7), we develop chance-constrained (CC) NN-LF and distributionally robust chance-constrained (DRCC) NN-LF formulations. In both cases, we also have the offline uncertainty training set $\mathcal{D}_\xi := \{\boldsymbol{\xi}_i\}_{i=1}^N$. For the CC-NN-LF formulation, we require

$$\mathbb{P}^*(\dot{V}_\theta(\mathbf{x}, \boldsymbol{\xi}) + \gamma\|\mathbf{x}\| \leq 0) \geq 1 - \beta, \quad \forall \mathbf{x} \in \mathcal{X}_\delta. \quad (20)$$

However, we are only given samples \mathcal{D}_ξ from \mathbb{P}^* . Assuming $\beta \leq \frac{1}{N}$, similarly to Proposition 1, we approximate (20) conservatively as, $\forall \mathbf{x}_i \in \mathcal{D}_{\text{LF}}, \max_j (\dot{V}_\theta(\mathbf{x}_i, \boldsymbol{\xi}_j) + \gamma\|\mathbf{x}_i\|) \leq 0$. Thus, to aim for the satisfaction of (20) for the training set \mathcal{D}_{LF} , we construct the loss function,

$$\ell_{\text{CC-LF}}(\theta) = \frac{1}{M} \sum_{i=1}^M (\max_j (\dot{V}_\theta(\mathbf{x}_i, \boldsymbol{\xi}_j) + \gamma\|\mathbf{x}_i\|))_+. \quad (21)$$

For the DRCC-NN-LF formulation, to account for errors between the empirical distribution $\hat{\mathbb{P}}_N$ and the true distribution \mathbb{P}^* as well as possible distribution shift during deployment, we require:

$$\inf_{\mathbb{P} \in \mathcal{M}_N^{r^*}} \mathbb{P}(\dot{V}_\theta(\mathbf{x}, \boldsymbol{\xi}) + \gamma\|\mathbf{x}\| \leq 0) \geq 1 - \beta, \quad \forall \mathbf{x} \in \mathcal{X}_\delta. \quad (22)$$

Table 1: Comparison of Cases 1 and 2 under different online true distributions. Here, “vio. rate” denotes violation rate: (validations with $\dot{V} > 0$)/(total validations), and “vio. area” denotes average violation area over all simulations: (data points with $\dot{V} > 0$)/(total data points). 5000 realizations of the online true uncertainty $\boldsymbol{\xi}^*$ are sampled from uniform and Gaussian distributions: $\boldsymbol{\xi}^* \sim [\mathcal{U}(1, 4), \mathcal{U}(1, 2)]^\top$ and $\boldsymbol{\xi}^* \sim [\mathcal{N}(4, 1.5), \mathcal{N}(1, 1.5)]^\top$ for Case 1, $\boldsymbol{\xi}^* \sim [\mathcal{U}(5, 7), \mathcal{U}(-1, 1)]^\top$ and $\boldsymbol{\xi}^* \sim [\mathcal{N}(7, 1), \mathcal{N}(1, 1)]^\top$ for Case 2.

Formulations	Case 1 Uniform		Case 1 Gaussian		Case 2 Uniform		Case 2 Gaussian	
	vio. rate	vio. area	vio. rate	vio. area	vio. rate	vio. area	vio. rate	vio. area
SOS	14.28%	0.94%	12.14%	1.53%	100%	15.52%	100%	18.55%
CC-SOS	11.78%	0.89%	8.30%	1.24%	0.00%	0.00%	5.10%	0.04%
DRCC-SOS	0.02%	0.00%	5.24%	0.80%	0.00%	0.00%	1.64%	0.01%
NN	31.80%	1.95%	16.66%	1.65%	100%	17.10%	100%	19.53%
CC-NN	1.82%	0.01%	6.24%	0.72%	0.00%	0.00%	1.26%	0.01%
DRCC-NN	0.00%	0.00%	3.22%	0.38%	0.00%	0.00%	0.72%	0.00%

Note that (22) can be tightened in terms of the CVaR approximation as:

$$\sup_{\mathbb{P} \in \mathcal{M}_N^r} \inf_{t \in \mathbb{R}} [\mathbb{E}_{\mathbb{P}}(\dot{V}_{\theta}(\mathbf{x}, \boldsymbol{\xi}) + \gamma \|\mathbf{x}\| + t)_+ - t\beta] \leq 0, \quad \forall \mathbf{x} \in \mathcal{X}_{\delta}. \quad (23)$$

Next, using the uncertainty set \mathcal{D}_{ξ} and the training dataset \mathcal{D}_{LF} and assuming $\beta \leq \frac{1}{N}$, similarly to Proposition 2, we rewrite the inequality conservatively as (equivalently if $\Xi = \mathbb{R}^m$), $\forall \mathbf{x}_i \in \mathcal{D}_{\text{LF}}$, $r \|\nabla V(\mathbf{x}_i)^\top d(\mathbf{x}_i)\|_{\infty} + \beta \max_j (\dot{V}_{\theta}(\mathbf{x}_i, \boldsymbol{\xi}_j)) + \gamma \|\mathbf{x}_i\| \leq 0$. Thus, we design the following empirical loss function for the DRCC-NN-LF formulation,

$$\ell_{\text{DRCC-LF}}(\boldsymbol{\theta}) = \frac{1}{M} \sum_{i=1}^M (r \|\nabla V(\mathbf{x}_i)^\top d(\mathbf{x}_i)\|_{\infty} + \beta \max_j (\dot{V}_{\theta}(\mathbf{x}_i, \boldsymbol{\xi}_j)) + \gamma \|\mathbf{x}_i\|)_+. \quad (24)$$

The neural network approach, with the novel loss function designs in (21) and (24), overcomes the issues noted above for the SOS approach. In particular, we do not require the dynamics to be described by polynomials and avoid scalability problems.

5. Evaluation

We apply the SOS approach (Sec. 4.1) and the neural network approach (Sec. 4.2) to synthesize LFs for a polynomial system and a pendulum system under model uncertainty.

Third-degree Polynomial System: Consider a two-dimensional polynomial system (Jasour, 2019):

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2}x_1^3 - \frac{3}{2}x_1^2 - x_2 \\ 6x_1 - x_2 \end{bmatrix} + \sum_{i=1}^2 d_i(\mathbf{x})\xi_i, \quad (25)$$

with two cases for the model uncertainty:

- Case 1: $r = 0.25$, $d_1(\mathbf{x}) = -[x_1, x_2]^\top$, $d_2(\mathbf{x}) = -[x_2, 0]^\top$, $\boldsymbol{\xi} \sim [\mathcal{N}(5, 1), \mathcal{N}(3, 1)]^\top$.
- Case 2: $r = 0.15$, $d_1(\mathbf{x}) = -[(x_1^3 + x_2), x_2]^\top$, $d_2(\mathbf{x}) = -[x_2, x_1]^\top$, $\boldsymbol{\xi} \sim [\mathcal{N}(6, 1), \mathcal{N}(0, 1)]^\top$.

Suppose that 9 samples $\{\boldsymbol{\xi}_i\}_{i=1}^9$ are available offline and set the confidence level to $\beta = 0.1$.

We compare the SOS search results with polynomial degree of 4 for the original SOS formulation in (3), the CC-SOS formulation in (12), and the DRCC-SOS formulation in (17). We also

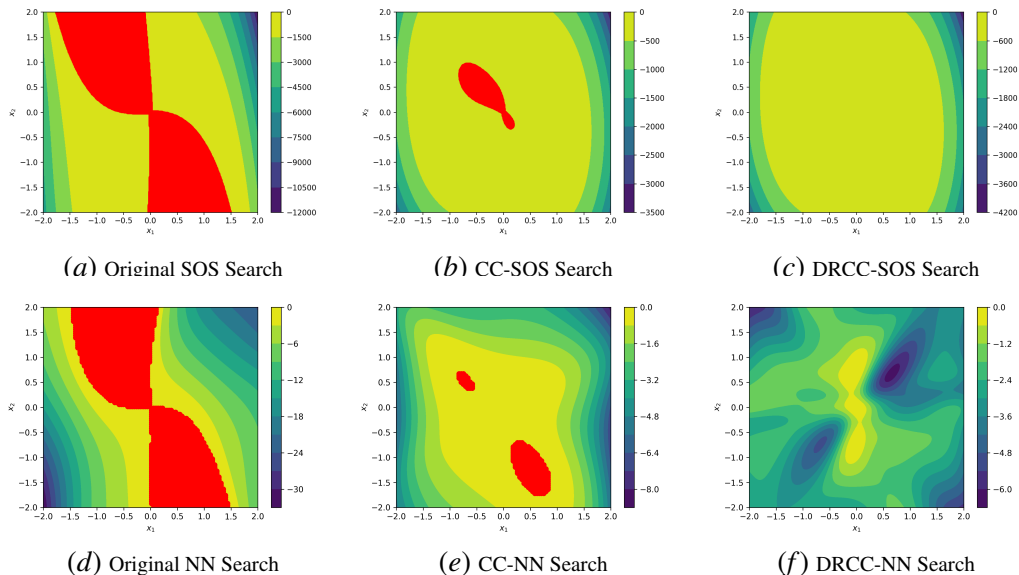


Figure 1: Results from SOS and NN formulations to design LF certificates for the polynomial system with Case 2 perturbations and online uncertainty $\xi^* = [1.9, 3.0]^\top$. The plots display the value of \dot{V} over the domain, where the red areas indicate positive values (violation of the LF derivative requirements).

include results from the NN formulation in (19), the CC-NN formulation in (21), and the DRCC-NN formulation in (24). For the neural network approach discussed in Sec. 4.2, we parametrize $V_\theta(\mathbf{x}) = |\phi_\theta(\mathbf{x}) - \phi_\theta(\mathbf{0})| + \hat{\alpha}\|\mathbf{x}\|$, where $\phi_\theta(\mathbf{x})$ is a fully connected three-layer neural network with 2-D input, two 16-D hidden layers, and 1-D output, with tanh activations. We train the network with the ADAM optimizer (Kingma and Ba, 2015) with learning rate 0.005 and Xavier initializer, and set the parameter $\hat{\alpha} = 0.05$.

We report qualitative results in Fig. 1 for Case 2 with online uncertainty $\xi^* = [1.9, 3.0]^\top$. We uniformly sample $\{\mathbf{x}_i\}_{i=1}^{5000}$ states in the region $x_1, x_2 \in [-2, 2]$. For the first-column plots, the resulting LFs from the baseline SOS and NN formulation fail to satisfy the Lyapunov condition for uncertain systems of the form (25), and the violation area is large since neither formulation takes uncertainty into account. For the second-column plots, the resulting LF from the CC-SOS or CC-NN formulation is less sensitive to uncertainty, since both take offline uncertainty samples into account. However, the resulting V still fails to satisfy the Lyapunov condition for (25). The LF resulting from our DRCC-SOS and DRCC-NN formulations in the last column satisfies the Lyapunov conditions for (25), even with out-of-distribution uncertainty. Table 1 shows quantitative results. We report the violation rate and average violation area for each of the 6 formulations: in all cases, the DRCC formulations outperform the CC and baseline formulations (no uncertainty considered) using either the SOS or the neural network approach in terms of violation rate and mean violation area.

Pendulum: Consider a pendulum with angle θ and angular velocity $\dot{\theta}$ following dynamics:

$$\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} \dot{\theta} \\ -\frac{mgl \sin \theta - b\dot{\theta}}{ml^2} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ -\frac{0.05b\dot{\theta}}{ml^2} & -\frac{0.05mgl \sin \theta}{ml^2} \end{bmatrix} \xi, \quad (26)$$

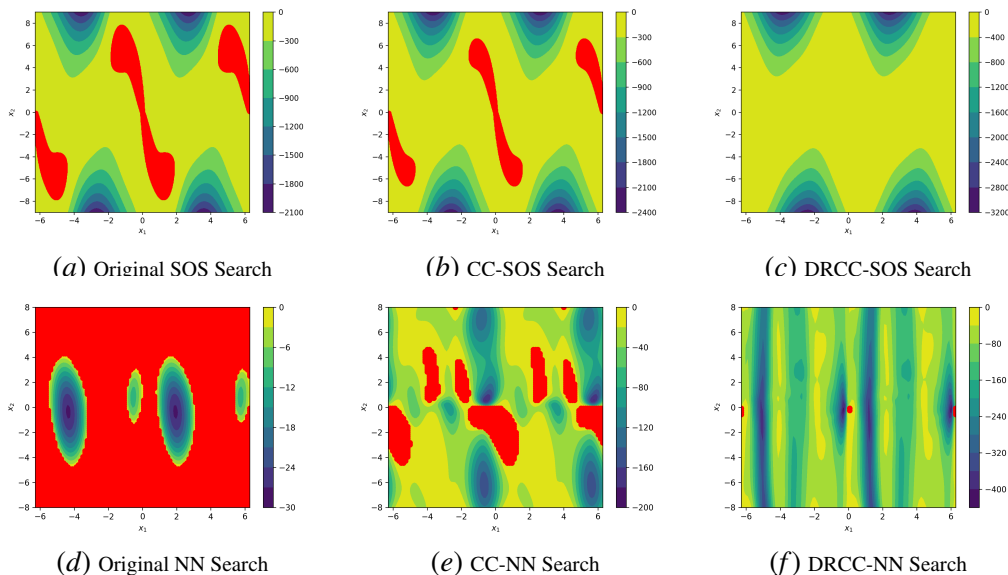


Figure 2: Results from SOS and NN formulations to design LF certificates for a pendulum with perturbation in the damping and length and online uncertainty $\xi^* = [-3.6, 1.4]^\top$. The plots display the value of \dot{V} over the domain, where the red areas indicate positive values (violation of the LF derivative requirement).

where $g = 9.81$ is the gravity acceleration, $m = 1.0$ is the ball mass, $l = 0.5$ is the length, $b = 0.1$ is the damping, and $d(\mathbf{x}) = [d_1(\mathbf{x}), d_2(\mathbf{x})]$ is the perturbation matrix with d_1 and d_2 representing perturbations in damping and length, respectively. We use 3 offline uncertainty samples $\{\xi_i\}_{i=1}^3$ with $\xi_i \sim [\mathcal{N}(0, 1), \mathcal{N}(0, 1)]^\top$, and set the confidence level $\beta = 0.1$. The SOS search polynomial is set to have a degree 4 with Wasserstein radius $r = 0.03$. The neural network ϕ_θ consists of a fully connected four-layer architecture, featuring a 3-D input, three 64-D hidden layers, and a 2-D output. The network employs tanh activations, and the pendulum state θ is rewritten as two states, $\sin \theta$ and $\cos \theta$. The Wasserstein radius is set to $r = 0.12$. We train the network with the ADAM optimizer with learning rate 0.002 and Xavier initializer, and set the parameter $\hat{\alpha} = 0.5$.

We compare the qualitative results between the SOS-based approaches and the NN-based approaches in Fig. 2 with the online uncertainty $\xi^* = [-3.6, 1.4]^\top$. Similar to Fig. 1, only the DRCC-SOS and DRCC-NN formulations meet the Lyapunov conditions within the domain of interest. The derivative violations observed near the small neighborhood of the equilibrium in the DRCC-NN formulation are a common issue in neural network-based Lyapunov functions, as reported in previous studies (Gaby et al., 2022; Chang et al., 2019).

6. Conclusions

We investigated the synthesis of Lyapunov functions for uncertain closed-loop dynamical systems. With only finitely many offline uncertainty samples, we derived novel distributionally robust formulations of sum-of-squares and neural-network approaches. The evaluation shows that LFs learned with our DRCC formulations are valid even for out-of-sample model errors. Future work will consider joint CLF and control policy search under distributional uncertainty and applications to higher-dimensional robotic control systems.

Acknowledgments

The authors gratefully acknowledge support from NSF under grants IIS-2007141 and CCF-2112665.

References

- Amir Ali Ahmadi and Anirudha Majumdar. Some applications of polynomial optimization in operations research and real-time decision making. *Optimization Letters*, 10(4):709–729, April 2016. ISSN 1862-4472. doi: 10.1007/s11590-015-0894-3.
- Zvi Artstein. Stabilization with relaxed controls. *Nonlinear Analysis-theory Methods & Applications*, 7:1163–1173, 1983.
- Nicholas M. Boffi, Stephen Tu, N. Matni, Jean-Jacques E. Slotine, and Vikas Sindhvani. Learning stability certificates from data. In *Conference on Robot Learning*, 2020.
- Dimitris Boskos, Jorge Cortés, and Sonia Martínez. Data-driven ambiguity sets with probabilistic guarantees for dynamic processes. *IEEE Transactions on Automatic Control*, 66(7):2991–3006, 2021. doi: 10.1109/TAC.2020.3014098.
- Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994. doi: 10.1137/1.9781611970777. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611970777>.
- Fernando Castañeda, Jason Choi, Bike Zhang, Claire Tomlin, and Koushil Sreenath. Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects and dynamics. In *2021 American Control Conference (ACC)*, pages 3683–3690, 2021. doi: 10.23919/ACC50511.2021.9483420.
- Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov control. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Zhi Chen, Daniel Kuhn, and Wolfram Wiesemann. Data-driven chance constrained programs over Wasserstein balls. *arXiv: Optimization and Control*, 2018.
- Jason Choi, Fernando Castañeda, Claire Tomlin, and Koushil Sreenath. Reinforcement learning for safety-critical control under model uncertainty, using control Lyapunov functions and control barrier functions. *ArXiv*, abs/2004.07584, 2020.
- Peter Coppens, Mathijs Schuurmans, and Panagiotis Patrinos. Data-driven distributionally robust LQR with multiplicative noise. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 521–530. PMLR, 10–11 Jun 2020. URL <https://proceedings.mlr.press/v120/coppens20a.html>.
- Jeremy Coulson, John Lygeros, and Florian Dörfler. Regularized and distributionally robust data-enabled predictive control. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2696–2701, 2019. doi: 10.1109/CDC40024.2019.9028943.

- Hongkai Dai, Benoit Landry, Lujie Yang, Marco Pavone, and Russ Tedrake. Lyapunov-stable neural-network control. *arXiv preprint arXiv:2109.14152*, 2021.
- Charles Dawson, Sicun Gao, and Chuchu Fan. Safe control with learned certificates: A survey of neural Lyapunov, barrier, and contraction methods. *ArXiv*, abs/2202.11762, 2022a.
- Charles Dawson, Zengyi Qin, Sicun Gao, and Chuchu Fan. Safe nonlinear control using robust neural Lyapunov-barrier functions. In *Proceedings of the 5th Conference on Robot Learning*, volume 164, pages 1724–1735. PMLR, 08–11 Nov 2022b.
- Vikas Dhiman*, Mohammad Javad Khojasteh*, Massimo Franceschetti, and Nikolay Atanasov. Control barriers in bayesian learning of system dynamics. *IEEE Transactions on Automatic Control*, 2021. doi: 10.1109/TAC.2021.3137059.
- Thai Duong and Nikolay Atanasov. Adaptive control of se(3) hamiltonian dynamics with learned disturbance features. *IEEE Control Systems Letters*, 6:2773–2778, 2022. doi: 10.1109/LCSYS.2022.3177156.
- Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171:115–166, 2018.
- Nathan Gaby, Fumin Zhang, and Xiaojing Ye. Lyapunov-net: A deep neural network architecture for lyapunov function approximation. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 2091–2096, 2022. doi: 10.1109/CDC51059.2022.9993006.
- Kevin Galloway, Koushil Sreenath, Aaron Ames, and Jessy Grizzle. Torque saturation in bipedal robotic walking through control Lyapunov function-based quadratic programs. *IEEE Access*, 3: 323–332, 2015.
- James Harrison, Apoorva Sharma, and Marco Pavone. Meta-learning priors for efficient online bayesian regression. In Marco Morales, Lydia Tapia, Gildardo Sánchez-Ante, and Seth Hutchinson, editors, *Algorithmic Foundations of Robotics XIII*, pages 318–337, Cham, 2020. Springer International Publishing.
- João Hespanha, Daniel Liberzon, and Andrew Teel. Lyapunov conditions for input-to-state stability of impulsive systems. *Automatica*, 44(11):2735–2744, 2008. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automatica.2008.03.021>. URL <https://www.sciencedirect.com/science/article/pii/S0005109808002689>.
- David Hilbert. Ueber die darstellung definiter formen als summe von formenquadraten. *Mathematische Annalen*, 32:342–350, September 1888.
- Ashish Ranjan Hota, Ashish Kumar Cherukuri, and John Lygeros. Data-driven chance constrained optimization under Wasserstein ambiguity sets. *2019 American Control Conference (ACC)*, pages 1501–1506, 2019.
- Ashkan Jasour. Risk aware and robust nonlinear planning. <https://rarnop.mit.edu/risk-aware-and-robust-nonlinear-planning>, 2019.
- Hassan Khalil. *Nonlinear Systems*. Prentice Hall, 1996.

- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint: 1412.6980*, 2015.
- Jean-Bernard Lasserre. Tractable approximations of sets defined with quantifiers. *Mathematical Programming*, 151:507–527, 2015.
- Monique Laurent. *Sums of Squares, Moment Matrices and Optimization Over Polynomials*, pages 157–270. Springer New York, New York, NY, 2009. ISBN 978-0-387-09686-5. doi: 10.1007/978-0-387-09686-5_7. URL https://doi.org/10.1007/978-0-387-09686-5_7.
- Kehan Long, Vikas Dhiman, Melvin Leok, Jorge Cortés, and Nikolay Atanasov. Safe control synthesis with uncertain dynamics and constraints. *IEEE Robotics and Automation Letters*, 7(3): 7295–7302, 2022a. doi: 10.1109/LRA.2022.3182544.
- Kehan Long, Yinzhuang Yi, Jorge Cortes, and Nikolay Atanasov. Safe and stable control synthesis for uncertain system models via distributionally robust optimization, 2022b. URL <https://arxiv.org/abs/2210.01341>.
- Arkadi Nemirovski and Alexander Shapiro. Convex approximations of chance constrained programs. *SIAM J. Optim.*, 17:969–996, 2006.
- Antonis Papachristodoulou and Stephen 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.*, volume 3, pages 3482–3487 vol.3, 2002. doi: 10.1109/CDC.2002.1184414.
- Pablo Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.
- Spencer M. Richards, Felix Berkenkamp, and Andreas Krause. The Lyapunov neural network: Adaptive stability certification for safe learning of dynamic systems. In *Proceedings of The 2nd Conference on Robot Learning*, volume 87, Oct 2018.
- Ralph Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.
- Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming*, volume 16. SIAM, Philadelphia, PA, 2014.
- Eduardo Daniel Sontag. A ‘universal’ construction of Artstein’s theorem on nonlinear stabilization. *Systems & Control Letters*, 13(2):117–123, 1989a. ISSN 0167-6911. doi: [https://doi.org/10.1016/0167-6911\(89\)90028-5](https://doi.org/10.1016/0167-6911(89)90028-5).
- Eduardo Daniel Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, 1989b. doi: 10.1109/9.28018.
- Eduardo Daniel Sontag and Yuan Wang. On characterizations of the input-to-state stability property. *Systems & Control Letters*, 24(5):351–359, 1995.
- Andrew Taylor, Victor Dorobantu, Hoang Le, Yisong Yue, and Aaron 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, 2019. doi: 10.1109/IROS40897.2019.8967820.

Rick Voßwinkel and Klaus Röbenack. Determining input-to-state and incremental input-to-state stability of nonpolynomial systems. *International Journal of Robust and Nonlinear Control*, 30(12): 4676–4689, 2020. doi: <https://doi.org/10.1002/rnc.5012>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rnc.5012>.

Weijun Xie. On distributionally robust chance constrained programs with Wasserstein distance. *Math. Program.*, 186:115–155, 2021.