

Neural Contraction Metrics with Formal Guarantees for Discrete-Time Nonlinear Dynamical Systems

Haoyu Li^{1*}

Xiangru Zhong^{2*}

Bin Hu²

Huan Zhang^{1,2}

HAOYULI5@ILLINOIS.EDU

XIANGRU4@ILLINOIS.EDU

BINHU7@ILLINOIS.EDU

HUAN@HUAN-ZHANG.COM

¹*University of Illinois Urbana-Champaign, Siebel School of Data Science and Computing*

²*University of Illinois Urbana-Champaign, Department of Electrical and Computer Engineering*

Editors: N. Ozay, L. Balzano, D. Panagou, A. Abate

Abstract

Contraction metrics are crucial in control theory because they provide a powerful framework for analyzing stability, robustness, and convergence of various dynamical systems. However, identifying these metrics for complex nonlinear systems remains an open challenge due to the lack of scalable and effective tools. This paper explores the approach of learning verifiable contraction metrics parametrized as neural networks (NNs) for discrete-time nonlinear dynamical systems. While prior works on formal verification of contraction metrics for general nonlinear systems have focused on convex optimization methods (e.g. linear matrix inequalities, etc) under the assumption of continuously differentiable dynamics, the growing prevalence of NN-based controllers, often utilizing ReLU activations, introduces challenges due to the non-smooth nature of the resulting closed-loop dynamics. To bridge this gap, we establish a new sufficient condition for establishing formal neural contraction metrics for general discrete-time nonlinear systems assuming only the continuity of the dynamics. We show that from a computational perspective, our sufficient condition can be efficiently verified using the state-of-the-art neural network verifier α, β -CROWN, which scales up non-convex neural network verification via novel integration of symbolic linear bound propagation and branch-and-bound. Built upon our analysis tool, we further develop a learning method for synthesizing neural contraction metrics from sampled data. Finally, our approach is validated through the successful synthesis and verification of NN contraction metrics for various nonlinear examples.

Keywords: Contraction Metrics, Formal Verification, Neural Networks, Nonlinear Systems

1. Introduction

Contraction theory is a prominent framework for analyzing the stability and safety of control systems. It examines how trajectories of a system remain close and converge to each other (Lohmiller and Slotine, 1998; Tran et al., 2018; Giesl et al., 2022; Manchester and Slotine, 2017; Tsukamoto et al., 2021; Bullo, 2022; Dawson et al., 2023). In contrast to Lyapunov theory, which assesses stability with respect to a specific equilibrium point (Lyapunov, 1992), contraction analysis focuses on the existence of a Riemannian-type metric d in which the system is contracting (Simpson-Porco and Bullo, 2014). This approach certifies that the distance between any pair of trajectories diminishes over time, ensuring that all solutions converge toward each other regardless of their initial conditions. Contraction theory has been applied to analyze and synthesize controllers for both continuous-time (Tsukamoto and Chung, 2020; Manchester and Slotine, 2017; Sun et al., 2021;

* Equal Contribution.

Tsukamoto et al., 2021) and discrete-time systems (Wei et al., 2021, 2022). However, many of these approaches are often heuristic-based methods and lack formal verification.

In this paper, we focus on synthesizing and formally verifying a neural network (NN) parametrized contraction metric that is valid over the largest possible domain for a given discrete-time nonlinear dynamical system. Although previous studies such as Giesl and Wendland (2019), Giesl et al. (2024), Pokkakkilath and Giesl (2024), and Fitzsimmons and Liu (2024) employ numerical quadrature or learning-based approaches to synthesize contraction metrics, they share common limitations. First, they all implicitly assume the smoothness of the system dynamics. However, with the advancements of deep learning, NN-based controllers have been popularized (Chang et al., 2019; Dai et al., 2021; Wu et al., 2023; Yang et al., 2024; Wang and Fazlyab, 2024), and the use of ReLU/LeakyReLU activations (Dai et al., 2021; Wu et al., 2023; Yang et al., 2024) just naturally introduce non-smoothness into the closed-loop dynamics. Consequently, traditional synthesis and verification techniques that rely on the Jacobian matrix of the dynamics are no longer applicable. Second, many existing works depend on convex optimization techniques such as linear matrix inequalities for the verification and learning of contraction metrics, which unavoidably increases the conservativeness of the resulting metrics.

To overcome these challenges, we devise a novel verification condition for discrete-time contraction analysis that depends neither on matrix inequalities nor on the Jacobian of the dynamics. Besides handling nonsmooth systems, this novel formulation also significantly reduces the complexity of verification conditions. It allows the reformulation of contraction certification into a verification problem within a bounded set on a computation graph, making it well-suited for advanced neural network verification tools like α, β -CROWN (Zhang et al., 2018; Xu et al., 2020, 2021; Wang et al., 2021), which has demonstrated its exceptional scalability in the verification of neural network controllers (Everett et al., 2023; Yang et al., 2024). In contrast to traditional SMT and MIP solvers, α, β -CROWN is far more efficient for this task, as it leverages the neural network architecture and benefits from GPU acceleration. By combining our novel verification formulation with the efficient neural network verifier, we achieved the first formally verified contraction metrics for a state feedback system with an NN controller. To summarize, our main contributions include:

- We develop a theoretical framework for formal contraction certificates applicable to potentially **non-smooth** discrete-time systems, proposing a **novel verification condition** for contraction metrics that eliminates the need for both matrix inequalities and the Jacobian of the dynamics;
- Unlike previous work with verified contraction metrics (Fitzsimmons and Liu, 2024), which uses expensive verifiers like SMT, our new formulation allows the combination of advanced neural network verification tool α, β -CROWN, enabling us to use **effective learning-based NN contraction metrics** while ensuring efficient **formal verification of contraction**;
- Our experiments on several autonomous and NN-controlled systems demonstrate that our learned NN contraction metric can produce a large verified region of contraction, sometimes close to the upper bound. For the first time in literature, we present verified neural contraction metrics for a **neural network-controlled state-feedback pendulum system**.

2. Preliminaries

Notations Our paper primarily focuses on a discrete time-invariant closed-loop control system

$$x_{k+1} = g(x_k, \pi(x_k)) = f(x_k) \quad (1)$$

where x_k is the system state, g denotes the open-loop dynamics, π denotes the policy, and f denotes the closed-loop nonlinear dynamics. We are interested in two cases: i) f involves neural networks, or ii) f is just a general nonlinear system. We denote the solution of the system starting with initial condition $x = x_0$ to be $\phi(k, x_0)$, i.e. it satisfies $\phi(k+1, x_0) = f(\phi(k, x_0))$. We denote the set of all positive semi-definite/positive definite matrices by $\mathbb{S}_+^{n \times n}$ and $\mathbb{S}_{++}^{n \times n}$ respectively. For two matrices A and B , we write $A \prec B$ or $A \preceq B$ to indicate that $B - A$ is positive definite or positive semi-definite, respectively. A function $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ is said to be a metric on \mathcal{X} if the following conditions holds: (1) $d(x, x) = 0, \forall x \in \mathcal{X}$; (2) $d(x, y) = d(y, x), \forall x, y \in \mathcal{X}$; (3) $d(x, z) \leq d(x, y) + d(y, z), \forall x, y, z \in \mathcal{X}$. We denote the norm induced by a matrix $P \in \mathbb{S}_{++}^{n \times n}$ as $\|\cdot\|_P$ such that $\|x\|_P = \sqrt{x^\top Px}$, and we denote $B(x; \epsilon)$ as the infinity norm ball around x with radius ϵ , i.e. $B(0; \epsilon) = \{y : \|y - x\|_\infty \leq \epsilon\}$.

Contraction Metrics This paper focuses on contraction metrics, which describe the asymptotic property of the differences between two solutions of a system. The formal definition goes as follows.

Definition 1 (Contraction) *The system (1) is said to be **contracting** on a forward invariant set $\mathcal{X} \subset \mathbb{R}^n$ if there exists a metric $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ and a scalar $\rho \in (0, 1)$ such that*

$$d(f(x), f(y)) \leq \rho d(x, y), \forall x, y \in \mathcal{X} \quad (2)$$

We call d the corresponding **contraction metric** and ρ the **contraction rate**.

Clearly, if a system is contracting with some metric d , the distance between any two trajectories of the system will converge to 0 with respect to the metric d . In the simplest setting, one can adopt $d(x, y) = \|x - y\|_2$, and the contraction analysis just reduces to standard Lipschitz constant analysis (Fazlyab et al., 2019; Huang et al., 2021; Araujo et al., 2023; Wang et al., 2024). However, general nonlinear systems typically require using contraction metrics in more general forms. If the closed loop dynamics f on the right hand side of system (1) is continuously differentiable, the existence of a metric d in which the system is contracting can be reduced to existence of a matrix-valued function as covered in Lohmiller and Slotine (1998). The formal statement goes as follows.

Theorem 2 *The system (1) is contracting in \mathcal{X} if there exists a nonsingular matrix-valued function $\Theta : \mathcal{X} \rightarrow \mathbb{R}^{n \times n}$ and positive constants μ, η, ρ such that for all $x \in \mathcal{X}$ we have*

$$\eta I \leq \Theta(x)^\top \Theta(x) \preceq \rho I, \quad F(x)^\top F(x) - I \preceq -\mu I \quad (3)$$

where $F(x) = \Theta(f(x)) \frac{\partial f}{\partial x}(x) \Theta^{-1}(x)$.

For the proof of this theorem, we refer the readers to Tran et al. (2018). This criterion has several drawbacks. First of all, the verification of matrix inequalities can be hard for verifiers. The most common way of verifying matrix inequalities relies on Sylvester's criterion, as used by Fitzsimmons and Liu (2024). As can be seen from theorem 2, we need to verify multiple semi-definite conditions, and each of them requires us to verify a determinant condition for each of the $2^n - 1$ principal minors. Therefore, the verification problem cannot scale when we consider higher-dimension systems. Secondly, this condition requires the system to be continuously differentiable as it requires the Jacobian of the dynamics. However, currently, many state-of-the-art results have controllers synthesized by ReLU networks (Yang et al., 2013; Wu et al., 2023), which makes this theorem inapplicable.

The α, β -CROWN Verifier α, β -CROWN (Zhang et al., 2018; Xu et al., 2020, 2021; Wang et al., 2021; Zhang et al., 2022) is a state-of-the-art neural network verification tool that aims to rigorously verify properties of neural networks. For a box \mathcal{B} and a general function F (which could involve NNs), it aims to prove $F(x) \geq 0$ for all $x \in \mathcal{B}$. The α, β -CROWN verifier implements an efficient linear bound propagation method and branch-and-bound. The verification problem is solved if this linear lower bound is greater than zero on the domain \mathcal{B} . Otherwise, it splits the domain into subdomains and further computes bounds for each subdomain until the condition is verified. More generally, if the problem contains multiple constraints such as $(F_1(x) \geq 0) \vee \dots \vee (F_n(x) \geq 0)$, α, β -CROWN will compute a linear lower bound for each of F_1, \dots, F_n , and claim successful verification of the problem if at least one of the lower bounds is greater than 0. Besides regular neural networks, the α, β -CROWN toolbox supports many operations crucial in control applications such as trigonometric and polynomial functions. Compared to traditional SMT and MIP solvers, α, β -CROWN offers better scalability since its core linear bound propagation procedure exploits the structure of NNs and computation graphs. This procedure does not rely on an external mathematical programming solver and can be efficiently accelerated on GPUs.

3. Contraction for Non-Smooth Dynamics

In this section, we aim to establish a condition that avoids the reliance on the Jacobian of the dynamics f and can be efficiently handled by α, β -CROWN. To achieve this, we propose conditions that guarantee contraction of the closed-loop dynamics under a Riemannian metric only assuming the continuity of f and mild conditions on the domain of interest \mathcal{X} . This approach bridges contraction theory with our practical verification methodology and enables us to efficiently learn and verify neural contraction metrics without the computational limitations imposed by matrix inequalities.

Theorem 3 *Assume that \mathcal{X} is an open, connected, and forward invariant subset in \mathbb{R}^n , and the dynamics f is continuous. Given $\epsilon > 0$ a positive threshold. Suppose there exists uniformly continuous $M(x) : \mathcal{X} \rightarrow \mathbb{S}_{++}^{n \times n}$ with a uniform lower bound $\mu I \preceq M(x)$ that satisfies*

$$\sqrt{(f(x) - f(y))^\top M(f(x))(f(x) - f(y))} \leq \rho \sqrt{(x - y)^\top M(x)(x - y)} \quad (4)$$

for all $x \in \mathcal{X}$ and $y \in B(x; \epsilon) \cap \mathcal{X}$ and some $0 < \rho < 1$. For every pair of nonnegative real numbers $a \leq b$, we define the space of admissible curves with starting time a and end time b as

$$\Gamma_{a,b} := \{\gamma : [a, b] \rightarrow \mathcal{X}, \gamma \text{ is continuous and piecewise regular}\} \quad (5)$$

where piecewise regular means that on each piece we have $\gamma \in C^1$ and $\|\gamma'(t)\| \neq 0$ for all t . The length of an admissible curve $\gamma : [a, b] \rightarrow \mathcal{X}$ under the metric M is defined as

$$L(\gamma) := \int_a^b \sqrt{\gamma'(t)^\top M(\gamma(t))\gamma'(t)} dt = \int_a^b \|\gamma'(t)\|_{M(\gamma(t))} dt \quad (6)$$

For each x, y , we then define $\Gamma_{x,y} := \bigcup_{a,b} \{\gamma \in \Gamma_{a,b}, \gamma(a) = x, \gamma(b) = y\}$. It is then guaranteed that the system (1) is contracting on \mathcal{X} in the metric $d(x, y) = \inf_{\gamma \in \Gamma_{x,y}} L(\gamma)$.

Lemma 4 *For any continuous curve $\gamma : [a, b] \rightarrow \mathcal{X}$ that connects $x, y \in \mathcal{X}$, there exists an admissible curve $\gamma_s : [a, b] \rightarrow \mathcal{X}$ connecting x, y and $\|\gamma(t) - \gamma_s(t)\| \leq \epsilon$ for any $\epsilon > 0$ and any t .*

Proof: Fix $\epsilon > 0$. Since γ is continuous, the image $\gamma([a, b])$ is compact. Since \mathcal{X} is assumed to be open, for each $t \in [a, b]$, there exists an $\epsilon_t > 0$ such that $B(\gamma(t); \epsilon_t) \subset \mathcal{X}$. Now we consider the open cover of $\gamma([a, b])$ by $\{B(\gamma(t); \frac{\epsilon_t}{2})\}_t$. By compactness, it must admit a finite subcover $B(\gamma(t_k); \frac{\epsilon_{t_k}}{2})$. Now pick $r > 0$ such that $r < \min(\min_k \frac{\epsilon_{t_k}}{2}, \epsilon)$. As γ is uniformly continuous on $[a, b]$, there exists a δ such that $|t - s| < \delta$ implies that $|\gamma(t) - \gamma(s)| < \frac{r}{4}$. Now we take a partition of $[a, b]$ into N subintervals $\{[t_i, t_{i+1}]\}_{i=1}^N$ with $\gamma(t_1) = a$ and $\gamma(t_{N+1}) = b$ such that $\frac{\delta}{2} < |t_{i+1} - t_i| < \delta$ for any i . We now consider the curve that is a piecewise straight line on each subinterval $[t_i, t_{i+1}]$. More formally, with the partition at hand, we first construct $N + 1$ points inductively as follows: We first choose $\phi_0 = \gamma(0)$. For $i \leq N$, we choose $\phi_i = \gamma(t_i)$ if $\gamma(t_i) \neq \phi_{i-1}$, and otherwise we choose ϕ_i in $B(\gamma(t_i); \frac{r}{16})$ such that $\phi_i \neq \phi_{i-1}$. For $i = N$ we additionally need to make sure that $\phi_i \neq \gamma(t_{N+1})$. And finally we pick $\phi_{N+1} = \gamma(t_{N+1})$. Now for $t \in [t_i, t_{i+1}]$ we define

$$\gamma_s(t) = \phi_i + \frac{t - t_i}{t_{i+1} - t_i} (\phi_{i+1} - \phi_i) \quad (7)$$

Notice that the above construction makes sure that $\phi_{i+1} \neq \phi_i$ for all i . Therefore, on each piece we have $\gamma'_s(t) = \frac{\phi_{i+1} - \phi_i}{t_{i+1} - t_i} \neq 0$. We can now estimate that on each subinterval, we have

$$\begin{aligned} \|\gamma_s(t) - \gamma(t)\| &\leq \|\gamma_s(t) - \gamma(t_i)\| + \|\gamma(t) - \gamma(t_i)\| \\ &\leq \|\phi_{i+1} - \gamma(t_{i+1})\| + \|\gamma(t_{i+1}) - \gamma(t_i)\| + \|\gamma(t_i) - \phi_i\| + \|\phi_i - \gamma(t_i)\| + \frac{r}{4} \\ &< r \leq \epsilon \end{aligned} \quad (8)$$

Finally we check that $\gamma_s(t) \in \mathcal{X}$ for any t . As each $\gamma(t)$ lies in some $B(\gamma(t_k); \frac{\epsilon_{t_k}}{2})$, we have

$$\|\gamma_s(t) - \gamma(t_k)\| \leq \|\gamma_s(t) - \gamma(t)\| + \|\gamma(t) - \gamma(t_k)\| \leq r + \frac{\epsilon_{t_k}}{2} \leq \epsilon_{t_k} \quad (9)$$

So clearly by construction we have $\gamma_s(t) \in B(\gamma(t_k); \epsilon_{t_k}) \subset \mathcal{X}$. ■

Proof of Thm 3: For the proof that the d is indeed a metric, we refer the readers to the text [Lee \(2018\)](#). We shall now show that the verification condition (4) implies the contraction in the defined metric. For any $\gamma \in \Gamma_{x,y}$, since γ is assumed to be regular on each piece, there exists a $\gamma_u : [c, d] \rightarrow \mathcal{X}$ such that $\gamma_u \in \Gamma_{x,y}$, $L(\gamma) = L(\gamma')$, and $\|\gamma'_u(t)\| = 1$ for almost every t (this is arclength parametrization, see [Lee \(2018\)](#)). Now we consider the curve $f(\gamma_u)$. As \mathcal{X} is forward invariant, $f(\gamma_u)$ also belongs to \mathcal{X} . Now the difficulty arises since $f(\gamma_u)$ might not be admissible, since we only assume f is continuous. However, a small perturbation on $f(\gamma_u)$ is always possible to make it admissible without affecting its length too much. More formally, we fix a $\rho' \in (\rho, 1)$. By the verification condition (4), we know that for all small enough δ_t we always have

$$\begin{aligned} &(f(\gamma_u(t + \delta_t)) - f(\gamma_u(t)))^\top M(f(\gamma_u(t + \delta_t)))(f(\gamma_u(t + \delta_t)) - f(\gamma_u(t))) \\ &\leq \rho^2 (\gamma_u(t + \delta_t) - \gamma_u(t))^\top M(\gamma_u(t + \delta_t)) (\gamma_u(t + \delta_t) - \gamma_u(t)) \\ &< \rho'^2 (\gamma_u(t + \delta_t) - \gamma_u(t))^\top M(\gamma_u(t + \delta_t)) (\gamma_u(t + \delta_t) - \gamma_u(t)) \end{aligned} \quad (10)$$

Notice that per our assumption, for every fixed t we can bound

$$(\rho'^2 - \rho^2) (\gamma_u(t + \delta_t) - \gamma_u(t))^\top M(\gamma_u(t + \delta_t)) (\gamma_u(t + \delta_t) - \gamma_u(t)) \geq \mu (\rho'^2 - \rho^2) \frac{\delta_t^2}{2} \quad (11)$$

for all small δ_t . Therefore, we see that after dividing by δ_t^2 , the margin of the strict inequality (10) is at least $\frac{\mu}{2}(\rho'^2 - \rho^2)$. Now fix an $\epsilon > 0$, we know that $f(\gamma_u)$ is a continuous curve, so that we can construct an admissible curve $\gamma_p : [c, d] \rightarrow \mathcal{X}$ such that $\|f(\gamma_u(t)) - \gamma_p(t)\| \leq \epsilon$ for any t . Now we see that we have

$$\begin{aligned} & |(f(\gamma_u(t + \delta_t)) - f(\gamma_u(t))) - (\gamma_p(t + \delta_t) - \gamma_p(t))| \\ & \leq |f(\gamma_u(t + \delta_t)) - \gamma_p(t + \delta_t)| + |f(\gamma_u(t)) - \gamma_p(t)| \leq 2\epsilon. \end{aligned} \quad (12)$$

Now by (10), the inequality is strict when the contraction factor is ρ' . This gives us room to wiggle the curve $f(\gamma_u)$ to be γ_p , which clearly belongs to the set $\Gamma_{f(x), f(y)}$. When ϵ is chosen to be small enough, we shall still have the inequality condition 10 for any t and small enough δ_t , i.e.

$$\|\gamma_p(t + \delta_t) - \gamma_p(t)\|_{M(\gamma_p(t + \delta_t))}^2 \leq \rho'^2 \|\gamma_u(t + \delta_t) - \gamma_u(t)\|_{M(\gamma_u(t + \delta_t))}^2, \quad (13)$$

as M is uniformly continuous. Now send $\delta_t^2 \rightarrow 0$ we get $\|\gamma'_p(t)\|_{M(\gamma_p(t))} \leq \rho' \|\gamma'_u(t)\|_{M(\gamma_u(t))}$. This then implies

$$d(f(x), f(y)) \leq \int_c^d \|\gamma'_p(t)\|_{M(\gamma_p(t))} dt \leq \rho' \int_c^d \|\gamma'_u(t)\|_{M(\gamma_u(t))} dt = \rho' L(\gamma) \quad (14)$$

Now, taking infimum over γ on the right-hand side gives the desired contraction result. \blacksquare

4. Methodology

In this section, we present techniques to bridge our novel contraction theory in Section 3 and α, β -CROWN to efficiently learn and verify neural contraction metrics.

4.1. Invariant Set Finding

We first need to identify a forward invariant set \mathcal{X} where the computation of contraction metrics will be performed. We adopt a method that demonstrates state-of-the-art results in certifying the region of exponential attraction to an equilibrium point in discrete-time systems (Yang et al., 2024). We use the following theorem to find a forward invariant set \mathcal{X} inside a box of interest \mathcal{B} :

Theorem 5 *Let $F(x) := V(f(x)) - (1 - \kappa)V(x)$ where $\kappa > 0$. If the condition*

$$(-F(x) \geq 0 \wedge f(x) \in \mathcal{B}) \vee (V(x) \geq \rho) \quad (15)$$

holds for any $x \in \mathcal{B}$, the closed loop dynamics (1) is Lyapunov stable with V being a valid Lyapunov function and the set $\mathcal{X} := \{x : V(x) < \rho\} \cap \mathcal{B}$ is certified forward invariant set as well as an exponential region of attraction.

The proof can be found in Yang et al. (2024). Eq. (15) is then converted into a loss function and trained with a counterexample-guided (CEGIS) approach to obtain the region of attraction.

4.2. Formal Verification of Contraction Metrics

Given a candidate metric $M(x)$ (possibly a NN), following our theory in Sec. 4, we now derive the verification formulation such that a verifier like α, β -CROWN can formally check if $M(x)$ is a contraction metric on the forward invariant domain \mathcal{X} . We introduce the auxiliary set \mathcal{B} as α, β -CROWN cannot handle implicitly defined domains like $\{V(x) < \rho\}$ directly. The verification condition for contraction can be formalized as follows.

Theorem 6 Suppose we have already formally verified the forward invariance of the set $\mathcal{X} = \{x : V(x) < \rho\} \cap \mathcal{B}$. We define

$$G(x, \delta) := (f(x) - f(x + \delta))^{\top} M(f(x))(f(x) - f(x + \delta)) - \rho^2 \delta^{\top} M(x) \delta. \quad (16)$$

If the condition

$$(-G(x, \delta) \geq 0) \vee (x + \delta \notin \mathcal{B}) \vee (V(x) \geq \rho) \vee (V(x + \delta) \geq \rho) \quad (17)$$

holds for any $(x, \delta) \in \mathcal{B} \times B(0; \epsilon)$ where $\epsilon > 0$, the system (1) is certifiably contracting on the domain \mathcal{S} with metric d and contraction rate $\rho' \in (\rho, 1)$.

Proof: Condition (17) enforces equation (16) to hold for all $x \in \mathcal{X}$ and $y \in B(0; \epsilon) \cap \mathcal{X}$. The contraction is then implied by Theorem 3. \blacksquare

This verification condition (4) allows us to verify contraction without relying on any Jacobian type condition or matrix inequality conditions, which enables the efficient verification by α, β -CROWN. Each term of the condition (4) can be represented as constraints on general functions. We then leverage α, β -CROWN’s efficient bound propagation and branch-and-bound method to verify these constraints over the domain \mathcal{B} . While existing solvers like SMT solvers struggle with scalability, our approach is inherently more scalable and benefits from GPU acceleration as a byproduct, further enhancing its efficiency.

4.3. Learning Neural Network based Contraction Metrics

We now present a learning-based approach to synthesize the contraction metric $M(x)$. Given a forward invariant domain \mathcal{X} , we desire to synthesize a uniformly continuous $M : \mathcal{X} \rightarrow \mathbb{R}^{n \times n}$ that has a uniform positive lower bound. To ease the learning and verification, we enforce the definiteness condition and the lower bound condition by parametrizing M to be

$$M(x) := \mu I + R(x)^{\top} R(x) \quad (18)$$

where $\mu > 0$, and R is a learnable function parametrized by a neural network. This ensures that M is symmetric and $M(x) \geq \mu I$ for every $x \in \mathcal{X}$. The learning objective of $M(x)$ should be the satisfaction of the condition (17), as demonstrated in the numerical condition below:

Theorem 7 Except the condition $x + \delta \in \mathcal{B}$, the rest of the verification condition (17) is equivalent to the following numerical condition:

$$L_{\text{violate}}(x, \delta; V, \rho) := \min(G(x, \delta), \rho - V(x), \rho - V(x + \delta)) \leq 0, \quad \forall x \in \mathcal{X} \quad (19)$$

As the condition $x + \delta \in \mathcal{B}$ can be easily verified posthoc, this numerical condition provides a simple training loss for the training of contraction metrics M , which is encoded in the function $G(x, \delta)$. We then adopt a counterexample-guided approach, where counterexamples of condition (19) are found and act as a dataset for the training. Instead of involving a sound verifier to generate the counterexamples as in Chang et al. (2019), we utilized cheap projected gradient descent (PGD) attacks to find counterexamples rather than involving an inefficient verifier at this stage. We found that the cheap attacker will effectively find counterexamples without damaging the learning performance. Also, since it is unknown how large the region where the system can be certified to be contracting is, we

Algorithm 1 Learning Contraction Metrics

Input: metrics parameter θ , dynamics f , region of interest \mathcal{B} , Lyapunov function V , levelset threshold ρ , learning rate η , PGD steps and stepsize n and β , epochs m

Output: A learned contraction metric, parameterized by θ

```

1: for  $scale = 0.1, 0.2, \dots, 1.0$  do
2:    $\rho' = scale \times \rho$                                 // Scale the threshold to adjust the contraction region
3:   Training dataset  $\mathcal{D} = \emptyset$                       // Initialize new empty dataset
4:   for  $iter = 1, 2, \dots$  do
5:      $(\xi_x^i, \xi_\delta^i) \leftarrow$  a batch of random PGD initializations
6:     for  $descent = 1, 2, \dots, n$  do
7:        $L \leftarrow L_{\text{violate}}(\xi_x^i, \xi_\delta^i; V, \rho')$ 
8:        $\xi_x^i \leftarrow \text{Project}_{\mathcal{B}}(\xi_x^i + \beta \cdot \frac{\partial L}{\partial \xi_x^i})$           // PGD update for x
9:        $\xi_\delta^i \leftarrow \text{Project}_{\mathcal{B}(0, \epsilon)}(\xi_\delta^i + \beta \cdot \frac{\partial L}{\partial \xi_\delta^i})$           // PGD update for \delta
10:       $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\xi_x^i, \xi_\delta^i)\}$                                 // Update Dataset with counterexamples
11:      for  $epoch = 1, 2, \dots, m$  do
12:         $L_{\text{train}} \leftarrow \sum_{\mathcal{D}} L_{\text{violate}}(\xi_x^j, \xi_\delta^j; V, \rho)$           // Summing over all the violations in the dataset
13:         $\theta \leftarrow \theta - \eta \nabla_{\theta} L_{\text{train}}$                                 // Gradient Descent to optimize metric parameters

```

gradually increase the size of the forward invariant set \mathcal{X} during training. The training algorithm is shown in algorithm 1.

Remark: We note that it is better not to include a strong regularization term on the parameters of the NN metric. Intuitively, suppose in addition that the forward invariant set \mathcal{X} is convex and bounded, the contraction metric admits an upper bound $M(x) \preceq \eta I$, in addition to the lower bound $\mu I \preceq M(x)$. One can then prove that $\sqrt{\mu} \|x - y\| \leq d(x, y) \leq \sqrt{\eta} \|x - y\|$ and $\|\phi(k, x_0) - \phi(k, x'_0)\| \leq \rho \sqrt{\frac{\eta}{\mu}} \|\phi(k-1, x_0) - \phi(k-1, x'_0)\|$. The regularization encourages $|\eta - \mu|$ to be smaller, potentially enforcing a contraction condition with respect to the Euclidean norm. This condition is much stricter than contraction in metric d , and thus could be harmful to the learning.

5. Experiments

To demonstrate the effectiveness of our method, we evaluate four different discrete-time systems that contain both smooth autonomous systems and a nonsmooth neural network state feedback system. All four systems are discretized with the explicit Euler method with a step size of 0.05. For each system, we first identify a Lyapunov function and a verified region of attraction using the method described in section 4.1. Next, we train a contraction metric $M(x)$ following the procedure described in section 4.3. The local contraction condition (17) is verified using α, β -CROWN with the specification described in 4.2. For comparison, we also compute a constant contraction metric for each system by directly using $M(x) = P$ provided by a quadratic Lyapunov function $V(x) = x^\top P x$ of the system. We shall demonstrate that by combining our training and verification, a large certified region of contraction can be achieved in all cases. In table 1, we summarize the verified region of the learned NN contraction metric and its runtime.

5.1. Nonlinear Autonomous Dynamical System

Van der Pol Equation. We consider the Euler discretization of the reverse Van der Pol equation:

$$\dot{x}_1 = -x_2, \quad \dot{x}_2 = x_1 - \mu (1 - x_1^2) x_2 \quad (20)$$

System	ROA	NN metric (ours)		Constant metric (baseline)		Verification Runtime
	ρ	$\rho (\uparrow)$	$r (\uparrow)$	$\rho (\uparrow)$	$r (\uparrow)$	
Van der Pol polynomial	0.119	0.119	100%	0.04	33.7%	109s
two-machine power	10.8	10.0	92.8%	1.5	13.8%	290s
inverted pendulum	0.086	0.086	100%	0.027	31.3%	62s
	672	580	88.6%	invalid	invalid	5.8hrs

Table 1: Verified regions of contraction of neural contraction metrics and verification runtime for different systems. These results demonstrate the effectiveness of the neural network contraction metrics, as well as the efficiency in terms of verification. The ρ of ROA is the largest region we consider, and thus serves as an **upper bound** of what we can achieve. The ratio of the volume of each verified region to the volume of ROA is also presented, denoted as r . We can see that the NN metric achieves a significantly larger verified region compared to the baseline (illustrated in Fig. 1). “Invalid” means that the synthesized metric does not provide a nontrivial region of contraction.

with a step size of 0.05. Here we set $\mu = 3$ following [Fitzsimmons and Liu \(2024\)](#). On the domain $\mathcal{B} = [-1.2, 1.2] \times [-2.3, 2.3]$, a quadratic Lyapunov function is employed to identify a forward invariant subset of the form $V_\rho = \{x \mid V(x) < \rho\}$ with $\rho = 0.119$. A metric that certifies contraction over the entire set V_ρ is found. Figure 1 shows the verified invariant sets where the contraction condition given by the learned NN contraction metric is satisfied (blue ellipse), which occupies the full forward invariant set (dashed black ellipse). In comparison, the largest verified invariant set satisfying the contraction condition given by the constant metric (depicted in the bottom row) is significantly smaller. Figure 1 further illustrates the regions where the contraction condition given by the learned metric is satisfied (white) and unsatisfied (grey).

Polynomial System. We consider the Euler discretization of the polynomial system from [Fitzsimmons and Liu \(2024\)](#):

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = -2x_1 + \frac{1}{3}x_1^3 - x_2 \quad (21)$$

with a step size of 0.05. On a domain $\mathcal{B} = [-4, 4] \times [-4, 4]$, we first compute a quadratic Lyapunov function with verified ROA given by sublevel set V_ρ with $\rho = 10.8$. With a NN contraction metric, we’re able to verify contraction on a sublevel set with $\rho = 10.0$, which occupies over 90 percent of the ROA. In contrast, with a constant metric, the largest sublevel set satisfying the contraction condition is only $\rho = 1.5$, as shown in Figure 1.

Two-machine Power System. We consider the Euler discretization of the two-machine power system from [Fitzsimmons and Liu \(2024\)](#), using a step size of 0.05. The system is described by

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = -0.5x_2 - (\sin(x_1 + \delta) - \sin(\delta)) \quad (22)$$

with $\delta = \pi/3$. Similar to the examples above, on $\mathcal{B} = [-1, 1] \times [-1, 1]$, we identify a quadratic Lyapunov function with verified ROA V_ρ given by $\rho = 0.086$. With a NN contraction metric, we’re able to verify contraction on the whole ROA. In comparison with a constant metric, the largest sublevel V_ρ satisfying the contraction is the one with $\rho = 0.027$, which occupies only 30 percent of the ROA as shown in Figure 1.

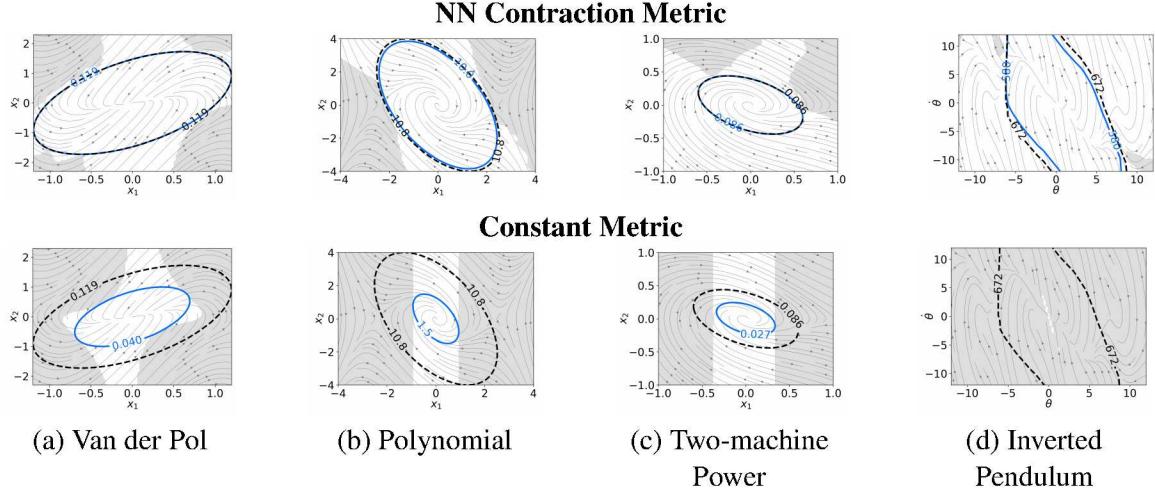


Figure 1: The figures visualize the certified region of contraction for our neural contraction metrics and the baseline constant metrics. For each $x \in \mathcal{B}$, we simulate the minimum of $G(x, \delta)$ for a small neighborhood of δ . The grey region indicates $G(x, \delta) < 0$, and the white region indicates the converse. In principle, any forward invariant set inside the white region can be certified as contraction. The top rows visualize the behavior of our learned neural contraction metrics, whereas the bottom row visualizes the baseline of using a constant contraction metric.

5.2. Non-Smooth Neural Network Controlled System

Inverted Pendulum We consider the Euler discretization of the controlled inverted pendulum

$$\ddot{\theta} = -\frac{\beta}{ml^2}\dot{\theta} + \frac{g}{l}\sin(\theta) + \frac{u}{ml^2} \quad (23)$$

with a step size of 0.05. We took the pretrained NN-based state feedback controller and an NN Lyapunov function given in [Yang et al. \(2024\)](#), which achieved state-of-the-art performance in discrete time ROA synthesis. Thanks to the flexibility of our novel formulation and the efficiency of α, β -CROWN, we can also find contraction metrics for NN-controlled state feedback systems, which are often highly non-linear and *non-smooth*. The largest ROA is given by V_ρ with $\rho = 672$. Using our pipeline, we successfully learn a contraction metric that certifies the contraction of the system in the sublevel set with $\rho = 580$, which occupies around 90% percent of the ROA as shown in Figure 1. A constant quadratic metric doesn't provide any reasonable region that satisfies the contraction condition. To the best of our knowledge, this provides the first certified region of contraction for an NN-controlled state feedback nonlinear system.

6. Conclusion

To summarize, we examine the stability of controlled discrete-time systems using contraction metrics, proposing a novel formulation that certifies contraction under the weaker assumption of continuity in the dynamics. By integrating the advanced neural network verifier α, β -CROWN, we enhance the training and verification process and certify contraction over nearly the entire forward-invariant set. Notably, this approach provides the first formally certified contraction metrics for a discrete-time neural network controlled state-feedback system in the literature.¹

1. Supplementary material can be found at <https://arxiv.org/abs/2504.17102>.

Acknowledgments

B. Hu is generously supported by the AFOSR award FA9550-23-1-0732. H. Zhang is supported in part by the AI2050 program at Schmidt Sciences (AI2050 Early Career Fellowship) and NSF (IIS-2331967).

References

Alexandre Araujo, Aaron J Havens, Blaise Delattre, Alexandre Allauzen, and Bin Hu. A unified algebraic perspective on Lipschitz neural networks. In *International Conference on Learning Representations*, 2023.

Francesco Bullo. *Contraction theory for dynamical systems*. Francesco Bullo, 2022.

Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. *Advances in neural information processing systems*, 32, 2019.

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 for robotics and control. *IEEE Transactions on Robotics*, 39(3):1749–1767, 2023.

Michael Everett, Rudy Bunel, and Shayegan Omidshafiei. Drip: Domain refinement iteration with polytopes for backward reachability analysis of neural feedback loops. *IEEE Control Systems Letters*, 7:1622–1627, 2023.

Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of Lipschitz constants for deep neural networks. *Advances in neural information processing systems*, 32, 2019.

Maxwell Fitzsimmons and Jun Liu. Computation and formal verification of neural network contraction metrics. *IEEE Control Systems Letters*, pages 1–1, 2024. doi: 10.1109/LCSYS.2024.3478272.

Peter Giesl and Holger Wendland. Construction of a contraction metric by meshless collocation. *Discrete and Continuous Dynamical Systems-B*, 24(8):3843–3863, 2019.

Peter Giesl, Sigurdur Hafstein, and Christoph Kawan. Review on contraction analysis and computation of contraction metrics. *arXiv preprint arXiv:2203.01367*, 2022.

Peter Giesl, Sigurdur Hafstein, and Iman Mehrabinezhad. Contraction metric computation using numerical integration and quadrature. *Discrete and Continuous Dynamical Systems-B*, 29(6):2610–2632, 2024.

Yujia Huang, Huan Zhang, Yuanyuan Shi, J Zico Kolter, and Anima Anandkumar. Training certifiably robust neural networks with efficient local Lipschitz bounds. *Advances in Neural Information Processing Systems*, 34:22745–22757, 2021.

John M Lee. *Introduction to Riemannian manifolds*, volume 2. Springer, 2018.

Winfried Lohmiller and Jean-Jacques E Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.

Aleksandr Mikhailovich Lyapunov. The general problem of the stability of motion. *International journal of control*, 55(3):531–534, 1992.

Ian R Manchester and Jean-Jacques E Slotine. Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design. *IEEE Transactions on Automatic Control*, 62(6):3046–3053, 2017.

Sareena Pokkakkillath and Peter Giesl. Construction of contraction metrics for discrete-time dynamical systems using meshfree collocation. *Discrete and Continuous Dynamical Systems-B*, 29(4):2043–2071, 2024.

John W Simpson-Porco and Francesco Bullo. Contraction theory on riemannian manifolds. *Systems & Control Letters*, 65:74–80, 2014.

Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *Conference on Robot Learning*, pages 1519–1539. PMLR, 2021.

Duc N Tran, Björn S Rüffer, and Christopher M Kellett. Convergence properties for discrete-time nonlinear systems. *IEEE Transactions on Automatic Control*, 64(8):3415–3422, 2018.

Hiroyasu Tsukamoto and Soon-Jo Chung. Neural contraction metrics for robust estimation and control: A convex optimization approach. *IEEE Control Systems Letters*, 5(1):211–216, 2020.

Hiroyasu Tsukamoto, Soon-Jo Chung, and Jean-Jacques E Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview. *Annual Reviews in Control*, 52:135–169, 2021.

Jiarui Wang and Mahyar Fazlyab. Actor-critic physics-informed neural lyapunov control. *arXiv preprint arXiv:2403.08448*, 2024.

Shiqi Wang, Huan Zhang, Kaidi Xu, Xue Lin, Suman Jana, Cho-Jui Hsieh, and J Zico Kolter. Beta-CROWN: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network verification. *Advances in Neural Information Processing Systems*, 34, 2021.

Zi Wang, Bin Hu, Aaron J Havens, Alexandre Araujo, Yang Zheng, Yudong Chen, and Somesh Jha. On the scalability and memory efficiency of semidefinite programs for Lipschitz constant estimation of neural networks. In *International Conference on Learning Representations*, 2024.

Lai Wei, Ryan McCloy, and Jie Bao. Control contraction metric synthesis for discrete-time nonlinear systems. *IFAC-PapersOnLine*, 54(3):661–666, 2021.

Lai Wei, Ryan McCloy, and Jie Bao. Discrete-time contraction-based control of nonlinear systems with parametric uncertainties using neural networks. *Computers & Chemical Engineering*, 166:107962, 2022.

Junlin Wu, Andrew Clark, Yiannis Kantaros, and Yevgeniy Vorobeychik. Neural lyapunov control for discrete-time systems. *Advances in neural information processing systems*, 36:2939–2955, 2023.

Kaidi Xu, Zhouxing Shi, Huan Zhang, Yihan Wang, Kai-Wei Chang, Minlie Huang, Bhavya Kailkhura, Xue Lin, and Cho-Jui Hsieh. Automatic perturbation analysis for scalable certified robustness and beyond. *Advances in Neural Information Processing Systems*, 33, 2020.

Kaidi Xu, Huan Zhang, Shiqi Wang, Yihan Wang, Suman Jana, Xue Lin, and Cho-Jui Hsieh. Fast and Complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In *International Conference on Learning Representations*, 2021.

Lujie Yang, Hongkai Dai, Zhouxing Shi, Cho-Jui Hsieh, Russ Tedrake, and Huan Zhang. Lyapunov-stable neural control for state and output feedback: A novel formulation. In *Forty-first International Conference on Machine Learning*, 2024.

Xiong Yang, Derong Liu, and Yuzhu Huang. Neural-network-based online optimal control for uncertain non-linear continuous-time systems with control constraints. *IET Control Theory & Applications*, 7(17):2037–2047, 2013.

Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. *Advances in Neural Information Processing Systems*, 31:4939–4948, 2018.

Huan Zhang, Shiqi Wang, Kaidi Xu, Linyi Li, Bo Li, Suman Jana, Cho-Jui Hsieh, and J Zico Kolter. General cutting planes for bound-propagation-based neural network verification. *Advances in Neural Information Processing Systems*, 2022.