Certified Neural Approximations of Nonlinear Dynamics

Frederik Baymler Mathiesen*
Delft Center for Systems and Control
Delft University of Technology, The Netherlands
f.b.mathiesen@tudelft.nl

Nikolaus Vertovec*
Department of Computer Science
University of Oxford, UK
nikolaus.vertovec@cs.ox.ac.uk

Francesco Fabiano
Department of Computer Science
University of Oxford, UK
francesco.fabiano@cs.ox.ac.uk

Luca Laurenti
Delft Center for Systems and Control
Delft University of Technology, The Netherlands
L.Laurenti@tudelft.nl

Alessandro Abate
Department of Computer Science
University of Oxford, UK
alessandro.abate@cs.ox.ac.uk

Abstract

Neural networks hold great potential to act as approximate models of nonlinear dynamical systems, with the resulting neural approximations enabling verification and control of such systems. However, in safety-critical contexts, the use of neural approximations requires formal bounds on their closeness to the underlying system. To address this fundamental challenge, we propose a novel, adaptive, and parallelizable verification method based on certified first-order models. Our approach provides formal error bounds on the neural approximations of dynamical systems, allowing them to be safely employed as surrogates by interpreting the error bound as bounded disturbances acting on the approximated dynamics. We demonstrate the effectiveness and scalability of our method on a range of established benchmarks from the literature, showing that it outperforms the state-of-the-art. Furthermore, we highlight the flexibility of our framework by applying it to two novel scenarios not previously explored in this context: neural network compression and an autoencoder-based deep learning architecture for learning Koopman operators, both yielding compelling results.

1 Introduction

Nonlinear dynamical models are ubiquitous across science and engineering and play a central role in describing and designing complex cyber-physical systems [Alur, 2015]. However, to verify and control such systems, it is often necessary to construct an abstraction: an approximation that locally simplifies the model or relaxes its nonlinearities [Derler et al., 2012, Khalil, 2002, Sastry, 1999]. In general, an abstraction translates the *concrete model*—the system under study—into a simpler *abstract model* that is more amenable to analysis [Baier and Katoen, 2008, Clarke et al., 2018]. A common method for synthesizing abstractions is known as *hybridization* and involves partitioning the state space into regions, each representing a state in a finite-state machine [Althoff et al., 2008,

*Authors contributed equally to this article.

Preprint. Under review.

Asarin et al., 2007, Bak et al., 2016, Dang et al., 2010, Frehse, 2005, García Soto and Prabhakar, 2020, Henzinger and Wong-Toi, 1995, Li et al., 2020, Majumdar and Zamani, 2012, Prabhakar et al., 2015, Roohi et al., 2016]. Neural networks with ReLU activations have emerged as a particularly effective approach to *hybridization*, as each network configuration implicitly induces a partition of the input domain into convex polytopes [Abate et al., 2022, Goujon et al., 2024, Villani and Schoots, 2023]. This enables simultaneous learning of both the partitioning and the simplified dynamics.

For neural abstractions to be practically useful, properties inferred from the abstraction—such as those related to reachability or safety—must reliably transfer to the original system. Simulation-based techniques fall short, as they are non-exhaustive and may miss unsafe behaviour. Formal verification provides a sound alternative by exhaustively analyzing all possible inputs and outputs. Prior work has used SMT solvers for this task [Fränzle et al., 2007, Abate et al., 2022, Solanki et al., 2025], but their computational cost severely limits the scale and expressivity of verifiable networks.

To overcome these limitations, we introduce a scalable framework for verification of neural abstractions that avoids reliance on expensive SMT solvers. Our method constructs certified first-order local approximations of the nonlinear system to tightly bound its behaviour. We then use neural network verification tools to certify that the network's output remains sufficiently close to the local approximation. By combining this with a parallelizable input-partitioning strategy, we achieve substantial improvements in verification speeds and can handle neural networks of larger sizes compared to state-of-the-art methods. Finally, to demonstrate the usefulness of our framework, we extend neural abstractions beyond their native developments, demonstrating their effective application to neural network compression (Section 4.2) and to learning Koopman operators [Brunton et al., 2016, 2022] (Section 4.3). Koopman operator theory, in particular, is an intriguing extension, as it is a mathematical framework in which nonlinear dynamical systems are represented as linear operators acting on a higher-dimensional function space.

Our contributions are summarised as follows:

1. We introduce a novel parallelizable method that enables adaptive verification of neural abstractions with nonuniform accuracy across the input domain.
2. We eliminate the reliance on SMT solvers capable of handling nonlinear real arithmetic, thereby removing the primary computational bottleneck in prior approaches.
3. We demonstrate that our approach outperforms existing methods on a variety of benchmarks and highlight the fundamental limitations inherent in the verification of neural abstractions.
4. We demonstrate the effectiveness of our approach across a range of applications, including neural network compression and the discovery of Koopman operators.

We begin in Section 2 by formally introducing neural abstractions, followed by presenting our approach to certification in Section 3. The diverse applications we discuss show that our approach is applicable to a broad range of tasks in the emerging field of Science for AI, and we argue it can be considered a positive contribution to understanding neural network representations [Heil et al., 2025], with no significant negative societal impact.

2 Neural approximations of nonlinear dynamics

We begin by introducing the system dynamics, for which we will synthesise neural network abstractions over a bounded domain. Let $\mathcal{X} \subset \mathbb{R}^n$ denote the bounded input domain of interest, and suppose $f : \mathcal{X} \to \mathbb{R}^m$ is a continuous (nonlinear) function describing the system's dynamics. We focus on two classes of systems:

- **Continuous-time nonlinear systems**, described by differential equations of the form $\frac{dx}{dt} = f(x), \quad x \in \mathcal{X}$;
- **Discrete-time nonlinear systems**, described by difference equations of the form $x_{k+1} = f(x_k), \quad x_k \in \mathcal{X}$.

Given a dynamical system as described above, a neural abstraction is an ϵ-close approximation of the dynamical system, as formally defined in Definition 1 below, which is a generalization of the definition of neural abstractions introduced in Abate et al. [2022] for continuous-time dynamical systems.

Definition 1 (Neural Abstraction)**.** *Consider a dynamical system described by function $f : \mathbb{R}^n \to \mathbb{R}^m$ and let $\mathcal{X} \subset \mathbb{R}^n$ be a region of interest. A feed-forward neural network $\mathrm{N} : \mathbb{R}^n \to \mathbb{R}^m$ defines a neural abstraction, also called a neural approximation, of f with error bound $\epsilon > 0$ over $\mathcal{X}$, if it holds that $\forall x \in \mathcal{X} : \ \|f(x) - \mathrm{N}(x)\| \leq \epsilon$ where $\|\cdot\|$ is the L_∞-norm*[2].

According to Definition 1, a neural abstraction of a dynamical system describes a dynamical system with a bounded additive disturbance d, such that any trajectory (solution to the differential equation) of the original dynamical system f is also a trajectory of the perturbed system. Specifically, in the case where f describes a continuous-time system, we have the following equation for the perturbed system:

$$\frac{dx}{dt} = \mathrm{N}(x) + d, \quad \|d\| \leq \epsilon, \quad x \in \mathcal{X}, \tag{1}$$

where ϵ represents the maximal deviation between the neural network approximation $\mathrm{N}(x)$ and the original function $f(x)$. In the discrete-time case, where f defines an update rule $x^{k+1} = f(x^k)$, any trajectory (solution to the difference equation) of the original system f is also a trajectory of the following discrete-time system:

$$x^{k+1} = \mathrm{N}(x^k) + d, \quad \|d\| \leq \epsilon, \quad x \in \mathcal{X}. \tag{2}$$

As the disturbance d is bounded by ϵ, the original system response, defined by f, is always contained within that of the abstraction, defined by N with disturbance d. Thus, the abstraction is sound, which enables formal guarantees to transfer from the abstraction to the concrete model.

2.1 Training

To obtain a neural network approximation of a dynamical system, we train a neural network N to minimize both the mean and maximum approximation error over a batch of sampled inputs $\{x_1, \ldots, x_M\}$. To this end, we assume for the remainder of the paper that the function f has bounded output and define the following loss function:

$$\mathcal{L} = \frac{1}{M} \sum_{l=1}^{M} \|f(x_l) - \mathrm{N}(x_l)\|_2 + \lambda_{\max} \max_{l \in \{1,\ldots,M\}} \|f(x_l) - \mathrm{N}(x_l)\|_\infty, \tag{3}$$

where the parameter $\lambda_{\max} = 0.001$ balances the trade-off between minimizing the average error and controlling the worst-case error across the sampled domain. We focus on training neural abstractions with ReLU activation functions, although our approach is compatible with more general activation functions, as permitted by the underlying solver (neural network verification tool). For our approach to be sound while also scalable, we utilise a complete solver, specifically Marabou 2.0 [Wu et al., 2024]. Further details regarding the network architecture, training procedure, and hyperparameters are provided in Section 4. Once a neural network N has been trained, the central challenge lies in certifying the accuracy, *i.e.*, formally establishing the relation $\|f(x) - \mathrm{N}(x)\| \leq \epsilon$ between the neural network abstraction and the concrete model. In what follows, we present our primary contribution, a scalable verification approach for this problem based on an adaptive refinement of first-order models. In Section 4, we empirically show how this framework substantially outperforms the state-of-the-art.

3 Certification of ϵ-closeness

We now introduce our verification approach that leverages local first-order models of $f(x)$, thereby enabling the application of modern techniques for neural network verification. To perform the verification, we will seek to prove that no counterexample against ϵ-closeness exists. Thus we seek an assignment of the negation of our desired specification, *i.e.*,

$$\exists x : \ \underbrace{x \in \mathcal{X} \wedge \|f(x) - \mathrm{N}(x)\| > \epsilon}_{\phi}. \tag{4}$$

If we find an assignment for x such that the formula ϕ is *satisfiable*, then we can establish that N is **not** a valid neural abstraction for a given accuracy ϵ. As the search for satisfying assignments is

[2]For the remainder of the paper, unless otherwise specified, all norms are L_∞.

exhaustive, failure to find an assignment constitutes a proof that no such assignment exists, and thus N is a valid neural abstraction for a given accuracy ϵ. In Section 3.1, we will describe the first-order models and how they can be used in the context of verification, followed by certificate refinement in Section 3.2 to combat conservatism introduced by the first-order models.

Remark. *The selection of ϵ can be performed empirically, based on the maximum error observed during training, or predefined according to strict application requirements. Notably, our proposed approach allows for an efficient search for the optimal ϵ within a given computational budget.*

3.1 First-order models

Given that the function f may include nonlinear terms, finding a satisfying assignment of ϕ in (4) typically requires reasoning over quantifier-free nonlinear real arithmetic formulae. This is computationally challenging and does not scale efficiently with problem complexity or the number of optimization variables. To address this, we introduce an over- and under-approximation of the dynamics of f using local first-order Taylor expansions of the vector field f. The choice of first-order models is a balance in the trade-off between expressivity, since we are verifying input-output relations, and maintaining linearity to enable formal verification. In this context, we will adaptively partition the domain of interest $\mathcal{X}$ into hyperrectangles, which are represented as weighted L_∞-balls.

Definition 2. *Given a hyperrectangle with radius $\delta \in \mathbb{R}^n_{\geq 0}$ and center $c \in \mathbb{R}^n$, the weighted L_∞-ball around c, denoted by $\mathcal{H}_\delta(c)$, is defined as*

$$\mathcal{H}_\delta(c) = \{x \in \mathbb{R}^n : |x - c| \leq \delta\},$$

where $|\cdot|$ is the element-wise absolute value and $\leq$ is interpreted element-wise. Similarly, we can define the hyperrectangle by its lower and upper corners, $\mathcal{H}^{\min} = c - \delta$ and $\mathcal{H}^{\max} = c + \delta$, respectively, as $\mathcal{H}_\delta(c) = \{x \in \mathbb{R}^n : \mathcal{H}^{\min} \leq x \leq \mathcal{H}^{\max}\}$.

The choice of partitioning will be discussed in more detail in the following section. For now, let us introduce the local first-order Taylor expansion, including an error bound.

Proposition 1 (Certified first-order Taylor Expansion)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}^m$ be a continuously differentiable function, and let $\mathcal{H}_\delta(c)$ be a hyperrectangle centered at $c \in \mathbb{R}^n$ with radius δ. Then, there exists a hyperrectangle $\mathcal{R} \subseteq \mathbb{R}^m$ such that for all $x \in \mathcal{H}_\delta(c)$, the following relation holds:*

$$f(x) \in (f(c) + \nabla f(c)(x - c)) \oplus \mathcal{R},$$

where $\oplus$ denotes the Minkowski sum.

Computing $\mathcal{R}$ can be done efficiently when f is twice continuously differentiable using the Lagrange error bound (see Appendix C.1 for details). The proof follows directly from Taylor's theorem for multivariate functions, along with the Lagrange error bound for higher-order terms [Joldes, 2011]. For the remainder of the paper, we use the indices $i \in \{1, \ldots, n\}$ when referring to the input dimensions of the function or the neural network, and $j \in \{1, \ldots, m\}$ when referring to the output dimensions. The first-order Taylor expansion in Proposition 1 provides the following sufficient condition for a valid neural abstraction.

Theorem 1. *Let $f : \mathbb{R}^n \to \mathbb{R}^m$, $\mathrm{N} : \mathbb{R}^n \to \mathbb{R}^m$, and let $\mathcal{H}_\delta(c) \subset \mathbb{R}^n$ be a hyperrectangle centered at c with radius δ. Let $f(c) + \nabla f(c)(x - c) \oplus \mathcal{R}$ be a certified Taylor expansion for f in $\mathcal{H}_\delta(c)$. If for each output dimension $j \in \{1, \ldots, m\}$, there does not exist a state x such that either of the following inequalities is satisfied:*

$$x \in \mathcal{H}_\delta(c) \wedge f_j(c) + \nabla f_j(c) \cdot (x - c) + \mathcal{R}_j^{\max} - \mathrm{N}_j(x) \geq \epsilon, \tag{5a}$$

$$x \in \mathcal{H}_\delta(c) \wedge \mathrm{N}_j(x) - f_j(c) - \nabla f_j(c) \cdot (x - c) - \mathcal{R}_j^{\min} \geq \epsilon, \tag{5b}$$

then N is an ϵ-accurate neural abstraction of f over $\mathcal{H}_\delta(c)$, i.e., *$\|f(x) - \mathrm{N}(x)\| \leq \epsilon, \forall x \in \mathcal{H}_\delta(c)$.*

The proof of Theorem 1 is provided in Appendix B. Both f and ∇f are evaluated at the center point, c, of the hyperrectangle $\mathcal{H}_\delta(c)$, ensuring that all terms highlighted in orange in Equations (5a) and (5b) remain fixed for a given hyperrectangle. In contrast, only the terms in blue vary with the specific choice of $x \in \mathcal{H}_\delta(c)$. As a result, the expression $f_j(c) + \nabla f_j(c)(x - c) + \mathcal{R}_j^{\max/\min}$ becomes linear. This allows Theorem 1 to be applied as a relaxation of the nonlinear predicate ϕ from Equation (4), thereby enabling formal verification to proceed without relying on SMT solvers

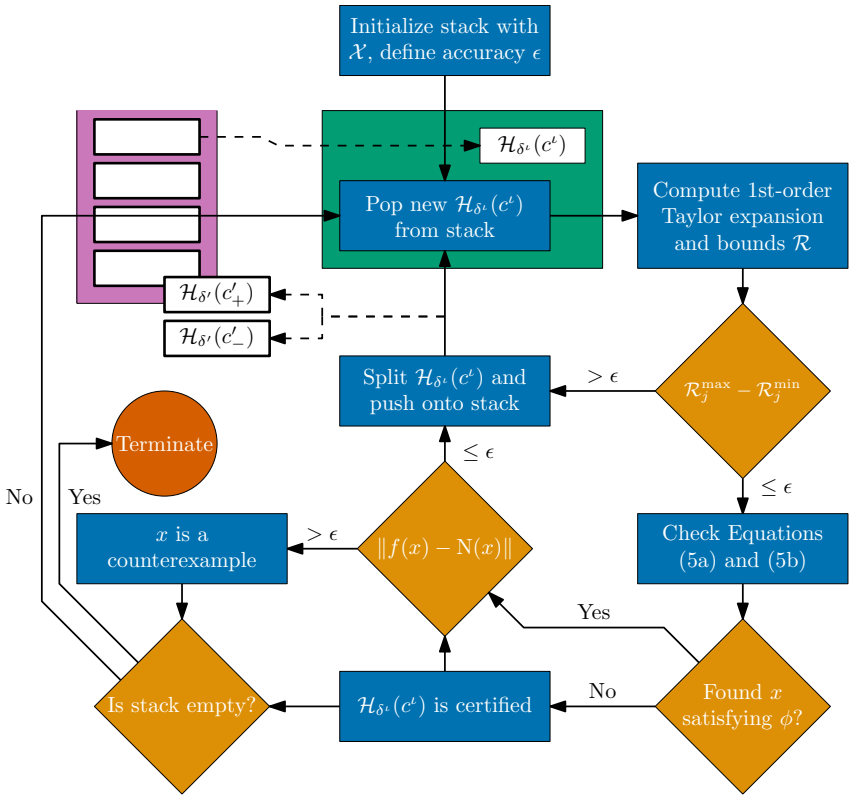

Figure 1: Graphical representation of the neural abstraction verification procedure with certificate refinement.

capable of reasoning over nonlinear real arithmetic. Specifically, we employ Marabou 2.0 [Wu et al., 2024], which implements an extension of the Simplex algorithm that was originally developed to solve linear programs [Dantzig, 2002], to verify the satisfiability of Equations (5a) and (5b).

Since the domain $\mathcal{X}$ can be over-approximated by a finite union of hyperrectangles, *i.e.*, $\mathcal{X} \subseteq \bigcup_{\iota=1}^{I} \mathcal{H}_{\delta^\iota}(c^\iota)$, we can perform verification locally within each hyperrectangle $\mathcal{H}_{\delta^\iota}(c^\iota)$. By applying a local first-order Taylor expansion within each region and bounding the remainder using $\mathcal{R}^{\max}$ and $\mathcal{R}^{\min}$, we obtain tighter bounds on the approximation error compared to approximating over the full domain $\mathcal{X}$. This localized approach effectively reduces the conservatism introduced by the approximation, *i.e.*, omitting higher-order derivative terms, while maintaining soundness of the verification process.

Remark. *We could allow the bound ϵ to vary over the domain $\mathcal{X}$, selecting different values of ϵ for each partition $\mathcal{H}_{\delta^\iota}(c^\iota)$. This would lead to a state-dependent disturbance in Equations* (1) *and* (2)*. Similarly, we could allow different ϵ for each output dimension,* i.e.*, output-weighted ϵ-closeness. However, for the sake of clarity and simplicity in the exposition, we omit this variation of ϵ.*

3.2 Certificate refinement

In the previous section, we introduced first-order Taylor expansions to derive conservative over- and under-approximations of the dynamics f. These approximations, captured in Equations (5a) and (5b), consider the worst-case realizations of the error term $r \in \mathcal{R}$. Consequently, when we compute the bounds $\mathcal{R}_j^{\max}$ and $\mathcal{R}_j^{\min}$, the counterexample x found may not always satisfy the formula ϕ from Equation (4). This happens because the error bounds derived from the Taylor expansion may be overly conservative. To address this issue, we propose a refinement strategy that partitions each hyperrectangle locally, enabling tighter approximations of the dynamics and reducing conservatism. The certification and partitioning strategy is illustrated in Figure 1. The decision to partition a hyperrectangle into two separate hyperrectangles, referred to as a split, results from one of two conditions:

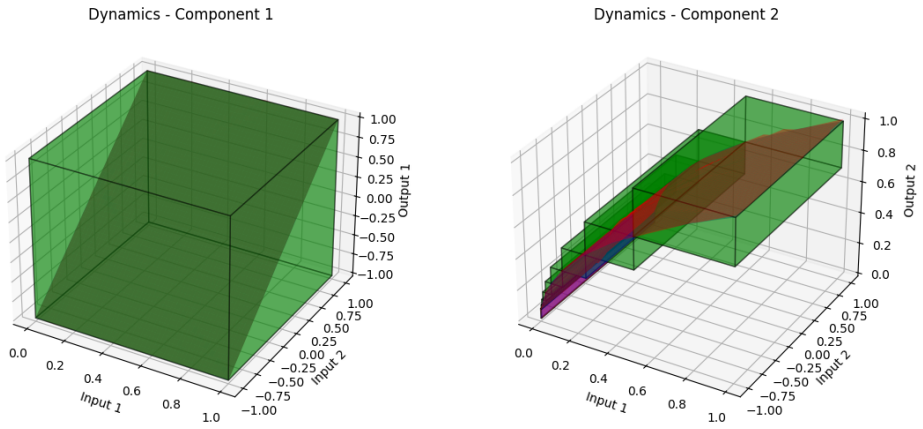

Figure 2: Partitioning of the domain to reduce conservatism. (*Left*) Linear terms do not require partitioning, as they are captured accurately by the first-order model. (*Right*) In regions of high nonlinearity (steeper dark function), finer rectangular partitioning reduces first-order approximation error by adaptively refining the domain.

1. the Taylor remainder term is too conservative, *i.e.*, $\mathcal{R}_j^{\max} - \mathcal{R}_j^{\min} > \epsilon$,
2. a counterexample x does not satisfy $\|f(x) - \mathrm{N}(x)\| > \epsilon$.

In the first case, we split the hyperrectangle based purely on the conservatism of the Taylor approximation, which can be done without reasoning over the neural network, while in the second case, the decision to split depends on the outcome of the network verification. If x satisfies $\|f(x) - \mathrm{N}(x)\| > \epsilon$, no further splitting is necessary and x is returned as a proper counterexample.

When a hyperrectangle is split, it is necessary to determine along which axis the split should occur. We choose this axis by identifying the input dimension that contributes most significantly to the Taylor approximation error in the output. Specifically, for each input dimension $i \in \{1, \ldots, n\}$, we perturb the center point c of the hyperrectangle along that dimension. The perturbed point, denoted c', is defined such that $c'_l = c_l$ for all $l \neq i$, and $c'_i = c_i + h_i$, where $h_i \in (0, \delta_i)$ is a small, fixed perturbation magnitude. For each such perturbation, we evaluate the absolute error between the first-order Taylor model and the true dynamics f at c'. The input dimension associated with the largest such error is then selected as the axis along which to split the hyperrectangle. For the selected input dimension i, we split the original hyperrectangle $\mathcal{H}_\delta(c)$ into two smaller hyperrectangles. These are centered at $c + (\delta - \delta')$ and $c - (\delta - \delta')$, respectively, where the new radius vector δ' is defined elementwise as:

$$\delta'_l = \begin{cases} \delta_l & \text{if } l \neq i, \\ \frac{\delta_l}{2} & \text{if } l = i. \end{cases}$$

This effectively halves the size of the hyperrectangle along the selected axis i, while keeping the width in all other dimensions unchanged. Since each input dimension x_i for $i \in \{1, \ldots, n\}$ can influence each output dimension f_j for $j \in \{1, \ldots, m\}$ differently, we perform verification and refinement separately for each output dimension.

The proposed partitioning strategy adapts the size of the hyperrectangles locally in accordance with the nonlinearity of the function $f(x)$. As illustrated in Figure 2, for the component f_1, which is linear, a single hyperrectangle suffices to certify the neural abstraction over the entire domain $\mathcal{X}$. In contrast, the second component, f_2, contains highly nonlinear terms that necessitate finer partitioning in regions where linear approximations are no longer sufficiently tight.

Since the verification of the abstraction can be performed locally over partitions $\mathcal{H}_\delta(c)$, we exploit this structure to parallelize the process across multiple processors, significantly improving performance. To facilitate parallel execution, we employ a shared stack accessed by a pool of worker processes. The domain $\mathcal{X}$ is initially partitioned into a set of hyperrectangles, which are pushed onto the stack. Each process draws a hyperrectangle from the stack, performs the verification procedure described earlier, and either certifies the region, marks it as uncertifiable if a counterexample is found, or splits the region as previously discussed. In the case of a split, the resulting subregions are pushed onto the

stack. The process then retrieves the next hyperrectangle and repeats the procedure, as summarized in Figure 1.

Remark. *The use of a stack (Last-In-First-Out) instead of a queue (First-In-First-Out) corresponds to a depth-first rather than breadth-first exploration of the verification space, consistent with strategies from branch-and-bound algorithms [Morrison et al., 2016]. A queue would be equally valid, though it would require more memory. If early termination with counterexamples is desired rather than verification until full coverage,* e.g.*, in the context of Counter-Example Guided Inductive Synthesis [Abate et al., 2018, 2022], a priority queue can be employed where the hyperrectangles would be weighted by the error relative to their volume (Lebesgue measure).*

4 Experimental Results

We empirically evaluate our approach on the benchmarks introduced in Abate et al. [2022] and described in Table 1, and whose detailed dynamics can be found in Appendix A, as well as on new benchmarks designed to demonstrate the extended capabilities of our method. In particular, in what follows, we first present an empirical comparison with the method introduced in Abate et al. [2022] and based on dReal [Gao et al., 2013], which currently represents the state-of-the-art for neural abstractions; then, to highlight the generality and scalability of our approach, we include two particularly challenging tasks: (i) a neural network compression benchmark, where a network with 5 layers and 1024 neurons per layer is compressed to a network with 5 layers and 128 neurons per layer, achieving a 98.4% reduction in size; and (ii) a verification benchmark based on a trajectory prediction network introduced in Dey and Davis [2023], Lusch et al. [2018], which learns to approximate nonlinear system dynamics through Koopman operator theory. These two benchmarks are discussed in Section 4.2 and 4.3, respectively. All experiments were executed on an Intel i7-6700k CPU (8 cores) with 16GB memory. All experiments can be reproduced using the scripts provided in the accompanying codebase available at `https://anonymous.4open.science/r/certified-neural-approximations-E679`.

4.1 Comparison with `dReal`-based approaches

To benchmark our method against the state-of-the-art, we compare it with the approach of Abate et al. [2022], which is based on `dReal` [Gao et al., 2013]; an SMT-solver over nonlinear real arithmetic[3]. As evident from the results in Table 1, our method scales to larger models more effectively than verification using `dReal`. This improved scalability arises as `dReal` is reasoning over nonlinear real arithmetic, while our method avoids this by reasoning over local linear approximations. Our approach successfully certifies all models with the 3x[64] architecture, while `dReal` exceeds the 1-hour timeout on all large models, with the exception of the *WaterTank* and *NonlinearOscillator*. Our method nevertheless achieves a noticeable speedup on all models (e.g. $\approx$150x faster for the *WaterTank* experiment). For small networks and simple systems (*WaterTank*, *NonLipschitzVectorField1*, and *NonLipschitzVectorField2*), our approach is slighty slower compared to the benchmark. This is due to the overhead introduced by our use of the Marabou solver, which involves translating the neural network into a linear program. It is important to emphasize that this overhead is not intrinsic to our approach but rather a consequence of the current implementation and the solver choice.

4.2 Compression of learned dynamics

To demonstrate the capabilities of our approach beyond constructing abstractions of known analytical dynamics, we apply the verification procedure to a neural network compression benchmark. The motivation for this benchmark is that learning a large over-parameterized neural network from data is common with state-of-the-art techniques, where the resulting neural network is typically very accurate and implicitly regularized [Martin and Mahoney, 2021, Belkin et al., 2019, Jacot et al., 2018]. The network is, however, computationally expensive, which is important in contexts such as embedded systems, and less amenable to analysis, *e.g.*, via the Piece-Wise Affine (PWA) representation induced by the ReLU network structure. Neural network compression seeks to reduce the size of a model—typically by pruning—while preserving its input-output behaviour [Luo et al., 2017, Memmel et al.,

[3]Note that our approach allows one to establish a certified subset of the domain, while the method in Abate et al. [2022] provides only a single counterexample. Our approach can be further extended with a variable ϵ over the domain of interest, allowing for a tighter certification in general.

Table 1: Verification Results for Learning Dynamical Systems[3]

Model	Network	Our approach Certified (%)	Our approach Time (s)	dReal Result	dReal Time (s)
WaterTank	[12]	100.0	1.02	✓	0.017
	3x[64]	100.0	3.08	✓	458.91
JetEngine	[10, 16]	100.0	9.30	✓	27.18
	3x[64]	100.0	47.01	Timeout (1h)	
SteamGovernor	[12]	100.0	12.59	✓	39.37
	3x[64]	100.0	160.03	Timeout (1h)	
Exponential	2x[14]	100.0	9.24	✓	9.99
	3x[64]	100.0	13.97	Timeout (1h)	
NonLipschitzVectorField1	[10]	100.0	2.65	✓	0.028
	3x[64]	100.0	9.52	Timeout (1h)	
NonLipschitzVectorField2	[12, 10]	100.0	7.75	✓	4.55
	3x[64]	100.0	40.09	Timeout (1h)	
VanDerPolOscillator	3x[64]	100.0	134.45	Timeout (1h)	
Sine2D	3x[64]	100.0	193.04	Timeout (1h)	
NonlinearOscillator	3x[64]	100.0	4.77	✓	234.52

2024]. However, the challenge in neural network compression is to guarantee that the compressed neural network is ϵ-close to the original network.

For this compression benchmark, we first train a 5-layer ReLU network with 1024 neurons per layer to learn the dynamics of the Lorenz attractor (described in Appendix A.10) from observed trajectories. We then seek to reduce the model by independently training a smaller ReLU network, consisting of 5 layers with 128 neurons per layer, to replicate the input-output behaviour of the larger model—without access to its training trajectories or the underlying system dynamics. The larger model comprises $4,205,571$ parameters, while the compressed model contains only $66,947$, corresponding to a 98.4% reduction in size. From the perspective of verification, the larger network serves as the reference dynamics, while the smaller network acts as an abstraction of those learned dynamics. This setup allows us to evaluate our method's ability to handle large-scale networks and non-analytical dynamics.

To fit within our verification framework, we construct certified first-order Taylor expansions of the neural network dynamic model using CROWN [Zhang et al., 2022]. This is necessary since the network is not twice continuously differentiable, which is required to calculate Lagrange error bounds. CROWN computes (local) linear relaxations of a nonlinear function, particularly neural networks, by recursively relaxing nonlinearities on the computation graph corresponding to the function (see Appendix C.2 for details). The verification procedure was executed with $\epsilon = 0.6$ and completed in 89 hours and 13 minutes. In total, $3,504,327$ hyperrectangles were checked and certified or further split according to the algorithm in Figure 1.

4.3 Verification of Koopman embeddings

We now shift to an operator-theoretic viewpoint of dynamical systems, wherein the evolution of a system is described through the action of a (linear) operator on measurement functions. This framework, known as *Koopman theory*, offers a powerful lens for analysing complex, nonlinear systems [Brunton et al., 2022]. Notably, Koopman theory provides a route to uncovering intrinsic coordinate systems in which the nonlinear dynamics manifest as linear. Originally introduced in Koopman [1931], the Koopman operator represents a nonlinear dynamical system via an infinite-dimensional linear operator acting on a Hilbert space of measurement functions. Despite the underlying system's nonlinearity, the Koopman operator is linear, and its spectral decomposition fully characterizes the system's behaviour [Brunton et al., 2016, 2022].

In general, however, the Koopman operator is infinite-dimensional, making its exact computation intractable. Thus, the aim is commonly to construct finite-dimensional approximations that capture the dominant behaviour of the system. This entails identifying a low-dimensional invariant subspace spanned by eigenfunctions of the Koopman operator, within which the dynamics evolve linearly. Despite the promise of Koopman embeddings, obtaining tractable representations has remained a central challenge in control theory. Utilizing neural networks to discover and represent Koopman eigenfunctions has emerged as a promising approach in recent years [Lusch et al., 2018, Dey and Davis, 2023]. While Koopman operators are commonly learned from data and a typical analysis of learned Koopman embeddings would verify structural properties—such as the orthonormality of eigenfunctions in Hamiltonian systems—such a treatment lies beyond the scope of this work. We instead demonstrate that our approach to verification can be applied to neural architectures deployed for learning Koopman embeddings, by verifying that the learned system evolution remains ϵ-close to the true system dynamics.

We adopt a standard setup using the `dlkoopman` library [Dey and Davis, 2023]. The network architecture comprises an autoencoder that learns the encoding and decoding of states into a Koopman-invariant subspace, along with a linear transformation within that subspace (Figure 3). This network architecture can be interpreted as a discrete-time neural abstraction that advances the system ahead one-time step and outputs the state at time step $k+1$ (Equation (2)). The notable difference to the discrete-time neural abstraction is that we consider the output of the network to be a trajectory, *i.e.*, a sequence of H-subsequent states. Applications of such trajectory tasks are commonly found throughout control theory, notably in model predictive control [Rawlings et al., 2017].

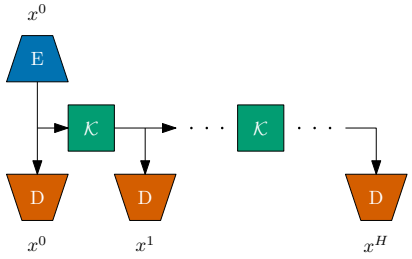

Figure 3: The network architecture used to learn a Koopman operator. The Encoder (blue) lifts the input into a higher dimensional space where linear multiplication with $\mathcal{K}$ (green) propagates the system. To obtain trajectory points, $[x^0, x^1, \ldots, x^{50}]$, each propagated state in the lifted space is decoded (orange).

We consider the Quadratic System provided in Appendix A.11, a benchmark problem frequently studied in the literature [H. Tu et al., 2014, Brunton et al., 2016, Lusch et al., 2018, Dey and Davis, 2023]. To improve reproducibility, we leverage the dataset provided by Dey and Davis [2023] to learn the evolution of the system from data. The final trained model takes an initial state x^0 and produces the sequence $[x^0, x^1, \ldots, x^{50}]$, which represents the evolution of the dynamics. The network architecture is summarized in Figure 3. Verification of the abstraction completed in 63.8min with 1955 counterexamples found and 97.20% of the domain certified. An interesting observation is that, although the network achieved a low prediction loss—specifically, a mean squared error of 0.002 on the validation set—our verification framework was still able to identify counterexamples where the prediction error exceeds the specified tolerance of $\epsilon = 0.1$. At the same time, the method certifies that the network ϵ-accurately predicts the system evolution over 97.20% of the input domain. Recall that in the presence of counterexamples, previous verification methods fail to identify regions where the model remains ϵ-accurate. Meanwhile, our approach offers valuable insight by localizing the regions in which the model can still be trusted, even when global verification fails.

5 Limitations

The primary focus of our approach has been to scale the certification of neural abstractions to accommodate larger, more expressive networks and to broaden their applicability. A key limitation that remains, however, is the scalability with respect to the input domain. Our certificate refinement strategy addresses this issue by verifying each output dimension independently, enabling local refinement of inputs only when they significantly affect a given output in a highly nonlinear manner. Nevertheless, for systems that are highly coupled and nonlinear, this approach can still lead to an exponential increase in the number of hyperrectangles that must be verified as the input dimensionality grows. This exponential blow-up—commonly referred to as the *curse of dimensionality*—is a pervasive challenge in safety-critical control and formal verification. Mitigating this issue through a judicious choice of coordinate representation that reduces coupling remains a focus for future research.

6 Conclusion

We presented a method for certifying neural abstractions of dynamical systems using local first-order models. By leveraging Taylor expansions and local partitioning, we provide a way to efficiently verify neural networks as abstractions of complex, nonlinear dynamical systems without relying on expensive SMT solvers for nonlinear arithmetic. Our approach reduces conservatism by adaptively refining the approximation based on the local nonlinearity of the system, leading to more accurate and scalable verification. We demonstrated the effectiveness of our approach on several challenging benchmarks, including network compression tasks and the verification of a trajectory prediction task based on Koopman linearization.

References

Alessandro Abate, Cristina David, Pascal Kesseli, Daniel Kroening, and Elizabeth Polgreen. *Counterexample Guided Inductive Synthesis Modulo Theories*, page 270–288. Springer International Publishing, 2018. ISBN 9783319961453. doi: 10.1007/978-3-319-96145-3_15. URL http://dx.doi.org/10.1007/978-3-319-96145-3_15.

Alessandro Abate, Alec Edwards, and Mirco Giacobbe. Neural abstractions. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

Matthias Althoff, Olaf Stursberg, and Martin Buss. Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization. In *Proceedings of the 47th IEEE Conference on Decision and Control (CDC)*, pages 4042–4048. IEEE, 2008.

Rajeev Alur. *Principles of Cyber-Physical Systems*. MIT Press, 2015. ISBN 9780262548922.

Eugene Asarin, Thao Dang, and Antoine Girard. Hybridization methods for the analysis of nonlinear systems. *Acta Informatica*, 43(7):451–476, 2007.

Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008. ISBN 9780262026499.

Stanley Bak, Sergiy Bogomolov, Thomas A. Henzinger, Taylor T. Johnson, and Pradyot Prakash. Scalable static hybridization methods for analysis of nonlinear systems. In *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control (HSCC)*, pages 155–164. ACM, 2016.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019. doi: 10.1073/pnas.1903070116. URL https://www.pnas.org/doi/abs/10.1073/pnas.1903070116.

Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE*, 11(2):e0150171, February 2016. ISSN 1932-6203. doi: 10.1371/journal.pone.0150171. URL http://dx.doi.org/10.1371/journal.pone.0150171.

Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern koopman theory for dynamical systems. *SIAM Review*, 64(2):229–340, 2022. doi: 10.1137/21M1401243.

Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. *Model Checking, 2nd Edition*. MIT Press, 2018. ISBN 9780262038836.

Thao Dang, Oded Maler, and Romain Testylier. Accurate hybridization of nonlinear systems. In *Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, pages 11–20. ACM, 2010.

George B Dantzig. Linear programming. *Operations research*, 50(1):42–47, 2002.

Patricia Derler, Edward A. Lee, and Alberto Sangiovanni Vincentelli. Modeling cyber–physical systems. *Proceedings of the IEEE*, 100(1):13–28, 2012. doi: 10.1109/JPROC.2011.2160929.

Sourya Dey and Eric William Davis. Dlkoopman: A deep learning software package for koopman theory. In Nikolai Matni, Manfred Morari, and George J. Pappas, editors, *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pages 1467–1479. PMLR, 15–16 Jun 2023.

Goran Frehse. Phaver: Algorithmic verification of hybrid systems past hytech. In *Hybrid Systems: Computation and Control (HSCC)*, volume 3414 of *Lecture Notes in Computer Science*, pages 258–273. Springer, 2005.

Martin Fränzle, Christian Herde, Tino Teige, Stefan Ratschan, and Tobias Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure1. *Journal on Satisfiability, Boolean Modeling and Computation*, 1(3–4):209–236, May 2007. ISSN 1574-0617. doi: 10.3233/sat190012.

Sicun Gao, Soonho Kong, and Edmund M. Clarke. dreal: An smt solver for nonlinear theories over the reals. In Maria Paola Bonacina, editor, *Automated Deduction – CADE-24*, pages 208–214, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-38574-2.

Miriam García Soto and Pavithra Prabhakar. Hybridization for stability verification of nonlinear switched systems. In *Proceedings of the 41st IEEE Real-Time Systems Symposium (RTSS)*, pages 244–256. IEEE, 2020.

Alexis Goujon, Arian Etemadi, and Michael Unser. On the number of regions of piecewise linear neural networks. *Journal of Computational and Applied Mathematics*, 441:115667, 2024. doi: 10.1016/j.cam.2023.115667.

Jonathan H. Tu, Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, 2014. ISSN 2158-2505. doi: 10.3934/jcd.2014.1.391. URL `http://dx.doi.org/10.3934/jcd.2014.1.391`.

Tilmann Heil, Fran Meissner, and Nikolaus Vertovec. Techno-material entanglements and the social organisation of difference. *Ethnic and Racial Studies*, 48(9):1859–1875, 2025. doi: 10.1080/01419870.2025.2469694.

Thomas A. Henzinger and Howard Wong-Toi. Linear phase-portrait approximations for nonlinear hybrid systems. In *Hybrid Systems*, volume 1066 of *Lecture Notes in Computer Science*, pages 377–388. Springer, 1995.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 8580–8589, Red Hook, NY, USA, 2018. Curran Associates Inc.

Mioara Maria Joldes. *Rigorous Polynomial Approximations and Applications*. Theses, Ecole normale supérieure de lyon - ENS LYON, September 2011.

Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 2002. ISBN 9780130673893.

B. O. Koopman. Hamiltonian systems and transformations in hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, 17(5):315–318, 1931. ISSN 00278424, 10916490. URL `http://www.jstor.org/stable/86114`.

Ding Li, Stanley Bak, and Sergiy Bogomolov. Reachability analysis of nonlinear systems using hybridization and dynamics scaling. In *Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 12288 of *Lecture Notes in Computer Science*, pages 265–282. Springer, 2020.

Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pages 5058–5066, 2017.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018. doi: 10.1038/s41467-018-07210-0.

Rupak Majumdar and Majid Zamani. Approximately bisimilar symbolic models for digital control systems. In *Computer Aided Verification (CAV)*, volume 7358 of *Lecture Notes in Computer Science*, pages 362–377. Springer, 2012.

Charles H. Martin and Michael W. Mahoney. Implicit self-regularization in deep neural networks: evidence from random matrix theory and implications for learning. *J. Mach. Learn. Res.*, 22(1), January 2021. ISSN 1532-4435.

Eva Memmel, Clara Menzen, Jetze Schuurmans, Frederiek Wesel, and Kim Batselier. Position: Tensor networks are a valuable asset for green AI. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 35340–35353. PMLR, 21–27 Jul 2024.

David R Morrison, Sheldon H Jacobson, and Joshua J Sauppe. Branch-and-bound algorithms: A survey of recent advances in searching, branching, and pruning. *Discrete Optimization*, 19:79–102, 2016. doi: 10.1016/j.disopt.2016.01.005.

Pavithra Prabhakar, Geir E. Dullerud, and Mahesh Viswanathan. Stability preserving simulations and bisimulations for hybrid systems. *IEEE Transactions on Automatic Control*, 60(12):3210–3225, 2015.

James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, Madison, WI, 2nd edition, 2017.

Nima Roohi, Pavithra Prabhakar, and Mahesh Viswanathan. Hybridization based cegar for hybrid automata with affine dynamics. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 9636 of *Lecture Notes in Computer Science*, pages 752–769. Springer, 2016.

Shankar Sastry. *Nonlinear Systems*, volume 10 of *Interdisciplinary Applied Mathematics*. Springer, New York, NY, 1999. ISBN 9780387985138.

Prashant Solanki, Nikolaus Vertovec, Yannik Schnitzer, Jasper Van Beers, Coen de Visser, and Alessandro Abate. Certified approximate reachability (care): Formal error bounds on deep learning of reachable sets. *arXiv preprint arXiv:2503.23912*, 2025. URL `https://arxiv.org/abs/2503.23912`.

Mattia Jacopo Villani and Nandi Schoots. Any deep relu network is shallow. *arXiv preprint arXiv:2306.11827*, 2023. URL `https://arxiv.org/abs/2306.11827`.

Haoze Wu, Omri Isac, Aleksandar Zeljić, Teruhiro Tagomori, Matthew Daggitt, Wen Kokke, Idan Refaeli, Guy Amir, Kyle Julian, Shahaf Bassan, Pei Huang, Ori Lahav, Min Wu, Min Zhang, Ekaterina Komendantskaya, Guy Katz, and Clark Barrett. *Marabou 2.0: A Versatile Formal Analyzer of Neural Networks*, page 249–264. Springer Nature Switzerland, 2024. ISBN 9783031656309. doi: 10.1007/978-3-031-65630-9_13.

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*, 33:1129–1141, 2020.

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*, 35:1656–1670, 2022.

A Dynamical systems

A.1 Water Tank

A simple first-order nonlinear dynamical system modelling the water level in a tank with constant inflow and outflow dependent on the water pressure (proportional to the square root of height).

$$\dot{x} = 1.5 - \sqrt{x}$$

where $x > 0$ represents the water level. For certification, we use $\epsilon = 0.097$ for the small network and a tighter $\epsilon = 0.007$ for the larger network. The input domain for certification is $\mathcal{X} = [0.1, 10.0]$.

A.2 Jet Engine

A two-dimensional nonlinear system with polynomial dynamics that models the behaviour of a simplified jet engine:

$$\begin{aligned}\dot{x} &= -y - 1.5x^2 - 0.5x^3 - 0.1\\ \dot{y} &= 3x - y\end{aligned}$$

For certification, we use $\epsilon = 0.039$ for the small network and a tighter $\epsilon = 0.012$ for the larger network. The input domain for certification is $\mathcal{X} = [-1.0, 1.0] \times [-1.0, 1.0]$.

A.3 Steam Governor

A mechanical governor used in steam engines, formulated as a three-dimensional nonlinear system with trigonometric terms:

$$\begin{aligned}\dot{x} &= y\\ \dot{y} &= z^2 \sin(x)\cos(x) - \sin(x) - 3y\\ \dot{z} &= -(\cos(x) - 1)\end{aligned}$$

For the implementation we use the trigonometric identity $\sin(x)\cos(x) = \frac{1}{2}\sin(2x)$ to rewrite $\dot{y}$ as $\frac{1}{2}z^2\sin(2x) - \sin(x) - 3y$. For certification, we use $\epsilon = 0.105$ for the small network and a tighter $\epsilon = 0.06$ for the larger network. The input domain for certification is $\mathcal{X} = [-1.0, 1.0] \times [-1.0, 1.0] \times [-1.0, 1.0]$.

A.4 Exponential System

The exponential system features highly nonlinear dynamics with nested nonlinearities combining trigonometric, exponential, and polynomial terms:

$$\begin{aligned}\dot{x} &= -\sin(e^{y^3+1}) - y^2\\ \dot{y} &= -x\end{aligned}$$

For certification, we use $\epsilon = 0.112$ for the small network and a tighter $\epsilon = 0.04$ for the larger network. The input domain for certification is $\mathcal{X} = [-1.0, 1.0] \times [-1.0, 1.0]$.

A.5 Non-Lipschitz Vector Field 1 (NL1)

A non-Lipschitz continuous vector field:

$$\begin{aligned}\dot{x} &= y\\ \dot{y} &= \sqrt{x}\end{aligned}$$

where $x \geq 0$. For certification, we use $\epsilon = 0.11$ for the small network and a tighter $\epsilon = 0.03$ for the larger network. The input domain for certification is $\mathcal{X} = [0.0, 1.0] \times [-1.0, 1.0]$.

A.6 Non-Lipschitz Vector Field 2 (NL2)

A more challenging non-Lipschitz continuous vector field:

$$\dot{x} = x^2 + y$$
$$\dot{y} = (x^2)^{1/3} - x$$

For certification, we use $\epsilon = 0.081$ for the small network and a tighter $\epsilon = 0.02$ for the larger network. The input domain for certification is $\mathcal{X} = [-1.0, 1.0] \times [-1.0, 1.0]$.

A.7 Van der Pol Oscillator

The classical Van der Pol oscillator with nonlinear damping:

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

where $\mu > 0$. For certification, we use $\epsilon = 0.25$. The input domain for certification is $\mathcal{X} = [-3.0, 3.0] \times [-3.0, 3.0]$.

A.8 Sine 2D System

The Sine 2D system represents a two-dimensional nonlinear oscillator with sinusoidal coupling:

$$\dot{x} = \sin(\omega_y \cdot y)$$
$$\dot{y} = -\sin(\omega_x \cdot x)$$

with parameter values $\omega_x = 1.0, \omega_y = 0.5$. For certification, we use $\epsilon = 0.02$. The input domain for certification is $\mathcal{X} = [-\pi, \pi] \times [-\pi, \pi]$. We utilize a LeakyReLU activation function for networks learning the Sine 2D System, both to improve learning accuracy and to demonstrate the applicability of our approach beyond standard ReLU activation functions.

A.9 Nonlinear Oscillator

The nonlinear oscillator combines linear, cubic, and sinusoidal terms:

$$\dot{x} = -ax - bx^3 + c\sin(x)$$

with parameter values $a = 1.0, b = 1/2, c = 0.3$. For certification, we use $\epsilon = 0.165$. The input domain for certification is $\mathcal{X} = [-3.0, 3.0]$.

A.10 Lorenz Attractor

The three-dimensional Lorenz Attractor, famous for exhibiting chaotic behaviour:

$$\dot{x} = \sigma(y - x)$$
$$\dot{y} = x(\rho - z) - y$$
$$\dot{z} = xy - \beta z$$

with parameter values $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. The input domain for certification is $\mathcal{X} = [-30.0, 30.0] \times [-30.0, 30.0] \times [0.0, 60.0]$

A.11 Quadratic System Dynamics and Discrete-Time Solution Derivation

A simple system governed by the continuous time dynamics:

$$\dot{x}_1 = \mu x_1$$
$$\dot{x}_2 = \lambda(x_2 - x_1^2)$$

where x_1 and x_2 are the state variables, and μ and λ are system parameters. The system includes a linear term for x_1 and a quadratic term involving x_1^2 in the equation for x_2. For training, the initial conditions there chosen at random and the trajectory is computed over the time horizon $[0, 1]$,

with parameters $\mu = -0.05, \lambda = -1$ and timestep of 0.02. The input domain for certification is $\mathcal{X} = [-0.5, 0.5] \times [-0.5, 0.5]$.

To generate trajectories of the system, we integrate the system numerically and sample the trajectory evenly across the time horizon. For the purpose of verification, we derive the analytic expression of the discrete-time system. The differential equation for $x_1(t)$ yields the solution

$$x_1(t) = x_1(0)e^{\mu t}$$

Solving the differential equation for $x_2(t)$ with $2\mu \neq \lambda$ yields

$$x_2(t) = \left(x_2(0) + \frac{\lambda x_1(0)^2}{2\mu - \lambda}\right) e^{\lambda t} - \frac{\lambda x_1(0)^2}{2\mu - \lambda} e^{2\mu t}$$

For a discrete time step Δt, we obtain the discrete-time system is:

$$\begin{aligned} x_1^{n+1} &= x_1^n e^{\mu \Delta t} \\ x_2^{n+1} &= \left(x_2^n + \frac{\lambda (x_1^n)^2}{2\mu - \lambda}\right) e^{\lambda \Delta t} - \frac{\lambda (x_1^n)^2}{2\mu - \lambda} e^{2\mu \Delta t} \end{aligned}$$

where $\Delta = 0.02$ in the experiments.

B Proof of Theorem 1

Proof. We can express $f(x)$ as:

$$f(x) = f(c) + \nabla f(c)(x - c) + r,$$

where $r \in [\mathcal{R}^{\min}, \mathcal{R}^{\max}]$. If no satisfying assignment exists for Equation (5a), it follows that:

$$f_j(c) + \nabla f_j(c)(x - c) + \mathcal{R}_j^{\max} - \mathrm{N}_j(x) < \epsilon.$$

Since $f_j(c) + \nabla f_j(c)(x - c) + \mathcal{R}_j^{\max}$ provides an upper bound for $f_j(x)$, we have:

$$f_j(x) - \mathrm{N}_j(x) < \epsilon.$$

Similarly, for the lower bound, Equation (5b):

$$\mathrm{N}_j(x) - f_j(x) < \epsilon.$$

These bounds hold for all $j \in \{1, \ldots, m\}$. Therefore, when no satisfying assignment is found for all $j \in \{1, \ldots, m\}$, it follows that:

$$\|f(x) - \mathrm{N}(x)\| \leq \epsilon, \quad \forall x \in \mathcal{H}_\delta(c).$$

□

C Certified Taylor expansions

C.1 Using elementary functions

Suppose that $f : \mathbb{R}^n \to \mathbb{R}^m$ is composed of smooth elementary functions and is at least twice continuously differentiable. We consider the first-order Taylor approximation of f around a point $x_0 \in \mathbb{R}^n$:

$$f(x) \approx f(x_0) + \nabla f(x_0)(x - x_0),$$

and define the remainder $\mathcal{R}(x) \in \mathbb{R}^m$ componentwise as:

$$\mathcal{R}_j(x) = f_j(x) - \left[f_j(x_0) + \nabla f_j(x_0)^\top (x - x_0)\right].$$

This remainder captures the contribution of second and higher-order terms. By the *Lagrange form* of the Taylor remainder, we have:

$$\mathcal{R}_j(x) = \frac{1}{2}(x - x_0)^\top \nabla^2 f_j(\xi)(x - x_0),$$

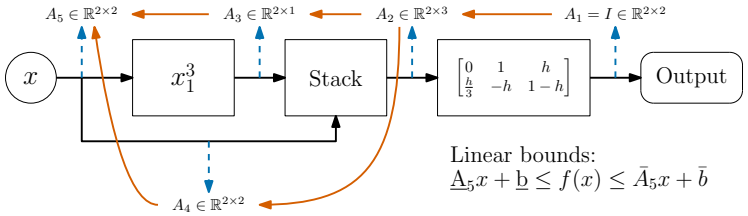

Figure 4: The computation graph with CROWN annotation for function $f(x) = \begin{bmatrix} x_1 + hx_2 \\ x_2 + h(\frac{1}{3}x_1^3 - x_1 - x_2) \end{bmatrix}$.

for some ξ on the line segment between x_0 and x. To bound the magnitude of the remainder, we find a constant M_j such that

$$\|\nabla^2 f_j(x)\|_2 \leq M_j \quad \text{for all } x \in \mathcal{H}_\delta(c),$$

where $\mathcal{H}_\delta(c)$ is the compact, convex hyperrectangle containing x_0 and x. Then,

$$|\mathcal{R}_j(x)| \leq \frac{1}{2} M_j \|x - x_0\|^2.$$

We use precomputed second derivatives of standard elementary functions, *e.g.*, for $\sin$ we obtain a tight upper bound by considering the bound on $|\sin(x)|$ over a given interval $[a, b]$. We evaluate:

- The absolute values at the endpoints: $|\sin(a)|$ and $|\sin(b)|$
- The values at critical points where the derivative $\cos(x) = 0$, i.e., $x = \frac{\pi}{2} + k\pi$

We identify all such critical points within $[a, b]$, evaluate $|\sin(x)|$ at each, and take the maximum of these values and the endpoint values:

$$\max_{x\in[a,b]} |\sin(x)| = \max \left(|\sin(a)|, |\sin(b)|, \max_{\substack{x\in[a,b] \\ \cos(x)=0}} |\sin(x)| \right)$$

Since $|\sin(x)| = 1$ at all such critical points, we simply check whether any $x = \frac{\pi}{2} + k\pi$ lies in the interval. If so, the maximum is 1. Otherwise, the bound is the maximum of the endpoints.

If the second derivative of a function is monotonic over the domain $\mathcal{H}_\delta(c)$, we can obtain a tight bound by evaluating at the endpoints:

- If $\nabla^2 f_j$ is increasing on $\mathcal{H}_\delta(c)$, then $M_j = \max_{x\in\mathcal{H}_\delta(c)} \nabla^2 f_j(x) = \nabla^2 f_j(x_{\max})$
- If $\nabla^2 f_j$ is decreasing on $\mathcal{H}_\delta(c)$, then $M_j = \nabla^2 f_j(x_{\min})$

For composite functions $f = h \circ g$, we compute second-derivative bounds using the chain and product rules:

$$\nabla^2 f_j(x) = J_g(x)^\top \nabla^2 h_j(g(x)) J_g(x) + \sum_i \frac{\partial h_j}{\partial z_i}(g(x)) \nabla^2 g_i(x),$$

where $J_g(x)$ is the Jacobian of g. Each term is bounded using known bounds on the elementary derivatives of h and g, and the ranges of intermediate variables over $\mathcal{H}_\delta(c)$.

C.2 Using CROWN

If the function f is not twice continuously differentiable, e.g. for a ReLU network, then the Lagrange bound is not valid to compute the Taylor remainder. Instead, we employ CROWN [Zhang et al., 2022], also known as Linear Bound Propagation, which was originally developed to verify neural networks via linear relaxations of the network. CROWN comes in several flavours including forward-mode (similar to the Taylor bound propagation above), backward-mode, forward-backward-mode (relaxations via forward mode), CROWN-IBP (relaxations via Interval Bound Propagation), α-CROWN (optimization of bounds), β-CROWN (neuron splitting branch-and-bound), and GCP-CROWN (general cutting planes). We only employ backward mode, which is the original version. In the remainder, when we refer to CROWN we mean backward-mode CROWN.

The idea of CROWN is to operate on the computation graph and relax nonlinearities based on node local input intervals, which can be computed using CROWN itself recursively. Figure 4 exemplifies this process on a composition of polynomial functions, for ease of exposition. First, the nonlinear term x^3 is locally relaxed to upper and lower linear functions based on the input range. Then, starting from the output with the linear functions $\underline{A}y + \underline{b} = Iy + 0$ and $\overline{A}y + \overline{b} = Iy + 0$, the bounding functions are propagated backward through the computation graph to the input. If the computation graph contains multiple nonlinearities, using the backward propagation from each nonlinearity to the input, the input range to each node is calculated to locally relax it. We refer to [Xu et al., 2020] for details on how to compute local relaxations based on input bounds and how to propagate through linear and locally relaxed nonlinear operations.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes] In Section 5 we discuss the limitation or our approach in detail.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes] The proof of Theorem 1 is provided in the appendix while remaining theoretical claims are discussed in detail in the main paper.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes] All the information needed to reproduce the main results is given in the appendix. We have also provided the code and scripts to reproduce the experimental results. In addition, to ease reproducibility of the verification results, we have made the ONNX files of all trained models available.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes] All code as well as the ONNX models used for verification have been made available. We have followed the Neurips submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`).

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (*e.g.*, data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes] In the appendix we have listed the dynamics and parameters of all dynamical systems used in our benchmarks. We believe this reflects a level of detail that is necessary to appreciate the results. Full details on all hyperparameters used for training and verification are given in the code repository.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [NA] Care was taken to explain that our results are not of a statistical nature as we utilise formal verification. As such error bars, confidence intervals, and statistical significance tests were not used in our experiments.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes] We specify the memory, runtime and CPU cores utilised for the experiments. We have ensured that all comparative benchmarks were run on similar hardware so as to accurately reflect the effects of computer resources.

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

Answer: [Yes] We have reviewed the NeurIPS Code of Ethics, and have addressed the Societal Impact of our work in the introduction, Section 1.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes] Our work is foundational research and not tied to particular applications. We have nevertheless addressed potential societal impacts in the introduction, Section 1.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (*e.g.*, pretrained language models, image generators, or scraped datasets)?

    Answer: [NA] We do not believe the data or models we have made available have a high risk for misuse.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (*e.g.*, code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes] We are the authors of the code and assets used in this work, with the exception of the benchmark dataset from Section 4.2, for which we have credited the original owners.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes] The code we have provided is well documented and we provide the scripts to generate the ONNX models, for full reproducibility.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA] The paper does not involve crowdsourcing nor research with human subjects.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA] The paper does not involve crowdsourcing nor research with human subjects.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA] This research does not involve LLMs as any important, original, or non-standard components.