POLAR: A Polynomial Arithmetic Framework for Verifying Neural-Network Controlled Systems

Chao Huang[1⋆], Jiameng Fan[2], Xin Chen[3], Wenchao Li[2], and Qi Zhu[4]

[1] University of Liverpool, chao.huang2@liverpool.ac.uk
[2] Boston University, {jmfan, wenchao}@bu.edu
[3] University of Dayton, xchen4@udayton.edu
[4] Northwestern University, qzhu@northwestern.edu

Abstract. We present POLAR[5], a **POL**ynomial **AR**ithmetic-based framework for efficient time-bounded reachability analysis of neural-network controlled systems. Existing approaches leveraging the standard Taylor Model (TM) arithmetic for approximating the neural-network controller cannot deal with non-differentiable activation functions and suffer from rapid explosion of the remainder when propagating TMs. POLAR overcomes these shortcomings by integrating TM arithmetic with *Bernstein polynomial interpolation* and *symbolic remainders*. The former enables TM propagation across non-differentiable activation functions and local refinement of TMs, and the latter reduces error accumulation in the TM remainder for linear mappings in the neural network. Experimental results show POLAR significantly outperforms the state-of-the-art tools on both efficiency and tightness of the reachable set overapproximation.

1 Introduction

Neural networks have been increasingly used as the central decision makers in a variety of control tasks [21, 17].However, the use of neural-network controllers also gives rise to new challenges on verifying the correctness of the resulting closed-loop control systems especially in safety-critical settings [30, 29]. In this paper, we consider the reachability verification problem of neural-network controlled systems (NNCSs). The high-level architecture of a simple NNCS is shown in Figure 1 in which the neural network senses the system state $\vec{x}$ at discrete time steps, and computes the corresponding control values $\vec{u}$ for updating the system dynamics which is defined by an ordinary differential equation (ODE) over $\vec{x}$ and $\vec{u}$. The *time-bounded reachability analysis problem* of an NNCS is to compute a state set that contains all the trajectories of a finite number of control steps from a given initial set. The initial set may represent uncertainties in the

⋆ Corresponding author. Part of the work was done when the author was in Northwestern University, US.

[5] The source code can be found at https://github.com/ChaoHuang2018/POLAR_Tool. The full version of this paper can be found at https://arxiv.org/abs/2106.13867.

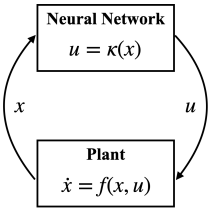

Fig. 1: A typical NNCS model.

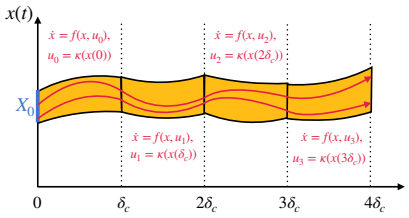

Fig. 2: Executions over 4 control steps.

starting state of the system or error (e.g. localization error) bounds in estimating the current system state during an execution of the system. Figure 2 shows an illustration of reachable sets for 4 steps, where the orange region represents the reachable set, and the two red, arrowed curves are two example trajectories starting from two different initial states in the initial set X_0 (blue).

Reachability analysis of general NNCSs is notoriously difficult due to nonlinearity in both the neural-network controller and the plant. The difficulty is further exacerbated by the coupling of the controller and the plant over multiple control steps. Since exact reachability of general nonlinear systems is undecidable [2], current approaches for reachability analysis largely focus on computing a tight overapproximation of the reachable sets [10, 6, 1]. Verisig [14] leverages properties of the sigmoid activation function and converts an NNCS with these activation functions to an equivalent hybrid system. Thus, existing tools for hybrid system reachability analysis can be directly applied to solve the NNCS reachability problem. However, this approach inherits the efficiency problem of hybrid system reachability analysis and does not scale beyond very small NNCSs. Another line of approach is to draw on techniques for computing the output ranges of neural networks [16, 27, 28, 24, 12, 26] and directly integrating them with reachability analysis tools designed for dynamical systems. NNV [25], for instance, combines star set analysis on the neural network with zonotope-based analysis of the nonlinear plant dynamics from CORA [1]. However, such approaches have been shown to be ineffective for NNCS verification due to the lack of consideration on the interaction between the neural-network controller and the plant dynamics [8, 11, 13]. In particular, since the primary goal of these techniques is to bound the output range of the neural network instead of approximating its input-output function, they *cannot track state dependencies across the closed-loop system and across multiple time steps in reachability analysis.*

More recent advances in NNCS reachability analysis are based on the idea of *function overapproximation* of the neural network controller. A function overapproximation of a neural network κ has two components: an approximated function p and an error term I (e.g. an interval) that bounds the approximation error. Such function overapproximation that produces a *point-wise* approximation of κ with an interval error term (typically called a remainder) is also known as a *Taylor model* (TM). Function-overapproximation approaches can be broadly categorized into two classes: *direct end-to-end approximation* such as Sherlock [8], ReachNN [11] and ReachNN* [9], and *layer-by-layer propagation* such as Verisig

2.0 [13]. The former computes a function overapproximation of the neural network end-to-end by sampling from the input space. The main drawback of this approach is that it does not scale beyond systems with more than a few input dimensions. The latter approach tries to exploit the neural network structure and uses *Taylor model arithmetic* to more efficiently obtain a function overapproximation of κ by propagating the TMs layer by layer through the network (details in Section 3). However, due to limitations of basic TM arithmetic, these approaches *cannot handle non-differentiable activation functions and suffer from rapid growth of the remainder* during propagation. For instance, explosion of the interval remainder would degrade a TM propagation to an interval analysis.

In this paper, we propose a principled POLynomial ARithmetic framework (POLAR) that enables precise layer-by-layer propagation of TMs for general feed-forward neural networks. Basic Taylor model arithmetic cannot handle ReLU that is non-differentiable (cannot produce the polynomial), and also suffers from low approximation precision (large remainder). POLAR addresses the key challenges of applying basic TM arithmetic through a novel use of *univariate Bernstein polynomial interpolation* and *symbolic remainders*. Univariate Bernstein polynomial interpolation enables the handling of non-differentiable activation functions and local refinement of Taylor models (details in Section 3.1). Symbolic remainders can taper the growth of interval remainders by avoiding the so-called wrapping effect [15] in linear mappings. The paper has the following novel contributions: (I) A polynomial arithmetic framework using both Taylor and univariate Bernstein approximations for computing NNCS reachable sets to handle general NN controllers; (II) An adaptation of the symbolic remainder method for ODEs to the layer-by-layer propagation for neural networks; (III) A comprehensive experimental evaluation of our approach on challenging case studies that demonstrates significant improvements of POLAR against SOTA.

2 Preliminaries

A *Neural-Network Controlled System (NNCS)* is a continuous plant governed by a neural network controller. The plant dynamics is defined by an ODE of the form $\dot{\vec{x}} = f(\vec{x}, \vec{u})$ wherein the state variables and control inputs are denoted by the vectors $\vec{x}$ and $\vec{u}$ respectively. We assume the function f is at least locally Lipschitz continuous such that its solution w.r.t. an initial state and constant control inputs is unique [20]. We denote the input-output mapping of the neural network controller as κ. The controller is triggered every δ_c time which is called the *control stepsize*. A system *execution (trajectory)* is produced as follows: starting from an initial state $\vec{x}(0)$, the controller senses the system state at the beginning of every control step $t{=}j\delta_c$ for $j{=}0, 1, \ldots$, and updates the control inputs to $\vec{v}_j{=}\kappa(\vec{x}(j\delta_c))$. The system's dynamics in that control step is governed by the ODE $\dot{\vec{x}}{=}f(\vec{x}, \vec{v}_j)$.

Given an initial state set $X_0 \subset \mathbb{R}^n$, all executions from a state in this set can be formally defined by a *flowmap* function $\varphi_{\mathcal{N}} : X_0 \times \mathbb{R}_{\geq 0} \to \mathbb{R}^n$, such that the system state at any time $t \geq 0$ from any initial state $\vec{x}_0 \in X_0$ is $\varphi_{\mathcal{N}}(\vec{x}_0, t)$. We call a state $\vec{x}' \in \mathbb{R}^n$ *reachable* if there exists $\vec{x}_0 \in X_0$ and $t \geq 0$ such that

$\vec{x}' = \varphi_{\mathcal{N}}(\vec{x}_0, t)$. The *reachability problem* on NNCS is to decide whether a state is reachable in a given NNCS, and it is *undecidable* since NNCS is more expressive than two-counter machines for which the reachability problem is already undecidable [2]. Many formal verification problems can be reduced to the reachability problem. For example, the safety verification problem can be reduced to checking reachability to an unsafe state. In the paper, we focus on computing the reachable set of an NNCS over a bounded number K of control steps. Since flowmap $\varphi_{\mathcal{N}}$ often does not have a closed form due to the nonlinear ODEs, we seek to compute *state-wise overapproximations* for it over multiple time segments, that is, in each control step $[j\delta_c, (j+1)\delta_c]$ for $j = 0, \ldots, K-1$, the reachable set is overapproximated by a group of flowpipes $\mathcal{F}_1(\vec{x}_0, \tau), \ldots, \mathcal{F}_N(\vec{x}_0, \tau)$ over the N uniformly subdivided time segments of the time interval, such that $\mathcal{F}_i(\vec{x}_0, \tau)$ is a *state-wise overapproximation* of $\varphi_{\mathcal{N}}(\vec{x}_0, j\delta_c + (i-1)\delta + \tau)$ for $\tau \in [0, \delta_c/N]$, i.e., $\mathcal{F}_j(\vec{x}_0, \tau)$ contains the exact reachable state from any initial state $\vec{x}_0$ in the i-th time segment of the j-th control step. Here, τ is the local time variable which is independent in each flowpipe. A high-level flowpipe construction algorithm is presented as follows, in which $\hat{X}_0 = X_0$ and $\delta = \delta_c/N$ is called the *time step*.

1: **for** $j = 0$ to $K-1$ **do**
2: Computing an overapproximation $\hat{U}_j$ for the control input range $\kappa(\hat{X}_j)$;
3: Computing the flowpipes $\mathcal{F}_1(\vec{x}_0, \tau), \ldots, \mathcal{F}_N(\vec{x}_0, \tau)$ for the continuous dynamics $\dot{\vec{x}} = f(\vec{x}, \vec{u}), \dot{\vec{u}} = 0$ from the initial set $\vec{x}(0) \in \hat{X}_j$, $\vec{u}(0) \in \hat{U}_j$;
4: $\mathcal{R} \leftarrow \mathcal{R} \cup \{\mathcal{F}_1(\vec{x}_0, \tau), \ldots, \mathcal{F}_N(\vec{x}_0, \tau)\}$;
5: $\hat{X}_{j+1} \leftarrow \mathcal{F}_N(\vec{x}_0, \delta)$;

Notice that $\vec{x}(0)$ denotes the local initial set for the ODE used in the current control step, that is the system reachable set at the time $j\delta_c$, while the variables $\vec{x}_0$ in a flowpipe are the symbolic representation of an initial state in X_0. Intuitively, a flowpipe overapproximates not only the reachable set in a time step, but also the *dependency* from an initial state to its reachable state at a particular time. For settings where the plant dynamics of an NNCS is given as a difference equation in the form of $\vec{x}_{k+1} = f(\vec{x}_k, \vec{u}_k)$, we can obtain *discrete* flowpipes which are the reachable set overapproximations at discrete time points by repeatedly computing the state set at the next step using TM arithmetic.

Dependencies on the initial set. As we mentioned previously, the reachable state of an NNCS at a time $t > 0$ is *uniquely determined* by its initial state if there is no noise or disturbance in the system dynamics or on the state measurements. If we use X_j to denote the exact reachable set $\{\varphi_{\mathcal{N}}(\vec{x}_0, j\delta_c) \,|\, \vec{x}_0 \in X_0\}$ from a given initial set X_0, then the control input range is defined by the set $U_j = \{\kappa(\vec{x}_j) \,|\, \vec{x}_j = \varphi_{\mathcal{N}}(\vec{x}_0, j\delta_c) \text{ and } \vec{x}_0 \in X_0\}$. More intuitively, the set U_j is the image from the initial set X_0 under the mapping $\kappa(\varphi_{\mathcal{N}}(\cdot, j\delta_c))$. *The main challenge in computing NNCS reachable sets is to control the overapproximation, which requires accurately tracking the dependency of a reachable set on the initial set across multiple control steps.* In this paper, we present a polynomial arithmetic framework for tracking such dependencies using Taylor models.

Taylor model arithmetic. Taylor models are originally proposed to compute higher-order overapproximations for the ranges of continuous functions (see [4]).

They can be viewed as a higher-order extension of intervals [22], which are sets of real numbers between lower and upper real bounds, e.g., the interval $[a, b]$ wherein $a \leq b$ represents the set of $\{x \,|\, a \leq x \leq b\}$. A *Taylor model (TM)* is a pair (p, I) wherein p is a polynomial of degree k over a finite group of variables $x_1, \ldots, x_n$ ranging in an interval domain $D \subset \mathbb{R}^n$, and I is the remainder interval. The range of a TM is the Minkowski sum of the range of its polynomial and the remainder interval. Thereby we sometimes intuitively denote a TM (p, I) by $p + I$ in the paper. TMs are closed under operations such as addition, multiplication, and integration (see [19]). Given functions f, g that are overapproximated by TMs (p_f, I_f) and (p_g, I_g), respectively, a TM for $f + g$ can be computed as $(p_f + p_g, I_f + I_g)$, and an order k TM for $f \cdot g$ can be computed as $(\, p_f \cdot p_g - r_k \,,\, I_f \cdot B(p_g) + B(p_f) \cdot I_g + I_f \cdot I_g + B(r_k) \,)$, wherein $B(p)$ denotes an interval enclosure of the range of p, and the *truncated part* r_k consists of the terms in $p_f \cdot p_g$ of degrees $> k$. Similar to reals and intervals, TMs can also be organized as vectors and matrices to overapproximate the functions whose ranges are multidimensional. Notice that *a TM is a function overapproximation and not just a range overapproximation like intervals or polyhedra.*

3 Framework of POLAR

In this section, we describe POLAR's approach for computing a TM for the output range of a neural network (NN) when the input range is defined by a TM. POLAR uses the layer-by-layer propagation strategy, and features the following key novelties: (a) A method to compute univariate Bernstein Polynomial **(BP)** overapproximations for activation functions, and selectively uses Taylor or Bernstein polynomials to *limit the overestimation produced when overapproximating the output ranges of individual neurons.* (b) A technique to symbolically represent the intermediate linear transformations of TM interval remainders during the layer-by-layer propagation. The purpose of using Symbolic Remainders **(SR)** is to *reduce the accumulation of overestimation in composing a sequence of TMs.*

3.1 Main Framework

We begin by introducing POLAR's propagation framework that incorporates only (a), and then describe how to extend it by further integrating (b). Although using TMs to represent sets in layer-by-layer propagation is already used in [13], the method only computes Taylor approximations for activation functions, and the TM output of one layer is propagated by the existing arithmetic for TM composition to the next layer. Such a method has the following shortcomings: (1) the activation functions have to be differentiable, (2) standard TM composition is often the source of overestimation even if preconditioning and shrink wrapping are used. Here, we seek to improve the use of TMs in the above two aspects.

Before presenting our layer-by-layer propagation method, we describe how a TM output is computed from a given TM input for a single layer. The idea is illustrated in Fig. 3. The circles in the right column denote the neurons in the

Algorithm 1 Layer-by-layer propagation using polynomial arithmetic and TMs

Input: Input TM $(p_1(\vec{x}_0), I_1)$ with $\vec{x}_0 \in X_0$, the $M+1$ matrices $W_1, \ldots, W_{M+1}$ of the weights on the incoming edges of the hidden and the output layers, the $M+1$ vectors $B_1, \ldots, B_{M+1}$ of the neurons' bias in the hidden and the output layers, the $M+1$ activation functions $\sigma_1, \ldots, \sigma_{M+1}$ of hidden and output layers.
Output: a TM $(p_r(\vec{x}_0), I_r)$ that contains the set $\kappa((p_1(\vec{x}_0), I_1))$.
1: $(p_r, I_r) \leftarrow (p_1, I_1)$;
2: **for** $i = 1$ to $M+1$ **do**
3: $(p_t, I_t) \leftarrow W_i \cdot (p_r, I_r) + B_i$; # Using TM arithmetic
4: Computing a polynomial approximation $p_{\sigma,i}$ for σ w.r.t. the domain (p_t, I_t);
5: Evaluating a conservative remainder $I_{\sigma,i}$ for $p_{\sigma,i}$ w.r.t. the domain (p_t, I_t);
6: $(p_r, I_r) \leftarrow p_{\sigma,i}(p_t + I_t) + I_{\sigma,i}$; # Using TM arithmetic
7: **return** (p_r, I_r).

current layer which is the i-th layer, and those in the left column denotes the neurons in the previous layer. The weights on the incoming edges to the current layer is organized as a matrix W_i, while we use B_i to denote the vector organization of the biases in the current layer. Given that the output range of the neurons in the previous layer is represented as a TM (vector) $(p_i(\vec{x}_0), I_i)$ wherein $\vec{x}_0$ are the variables ranging in the NNCS initial set. Then, the output TM $(p_{i+1}(\vec{x}_0), I_{i+1})$ of the current layer can be obtained as follows. First, we compute the polynomial approximations $p_{\sigma_1,i}, \ldots, p_{\sigma_l,i}$ for the activation functions $\sigma_1, \ldots, \sigma_l$ of the neurons in the current layer. Second, interval remainders $I_{\sigma_1,i}, \ldots, I_{\sigma_l,i}$ are evaluated for those polynomials to ensure that for each $j = 1, \ldots, l$, $(p_{\sigma_j,i}, I_{\sigma_j,i})$ is a TM of the activation function σ_j w.r.t. z_j ranging in the j-th dimension of the set $W_i(p_i(\vec{x}_0) + I_i) + B_i$. Third, $(p_{i+1}(\vec{x}_0, I_{i+1}))$ is computed as the TM composition $p_{\sigma,i}(W_i(p_i(\vec{x}_0) + I_i) + B_i) + I_{\sigma,i}$ wherein $p_{\sigma,i}(\vec{z}) = (p_{\sigma_1,i}(z_1), \ldots, p_{\sigma_l,i}(z_k))^T$ and $I_{\sigma,i} = (I_{\sigma_1,i}, \ldots, I_{\sigma_l,i})^T$. Hence, when there are multiple layers, starting from the first layer, the output TM of a layer is treated as the input TM of the next layer, and the final output TM is computed by composing TMs layer-by-layer.

We give the whole procedure by Algorithm 1, where the polynomial approximation $p_{\sigma,i}$ and its remainder interval $I_{\sigma,i}$ for the vector of activation functions σ in the i-th layer can be computed in the following two ways.

Taylor approximation. When the activation function is differentiable in the range defined by (p_t, I_t). The polynomial $p_{\sigma,i}$ can be computed as the order k Taylor expansion of σ (in each of its dimension) at the center of (p_t, I_t), and the remainder is evaluated using interval arithmetic based on the Lagrange remainder form. More details are described elsewhere [19].

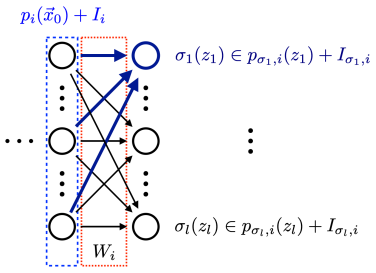

Fig. 3: Single layer propagation

Bernstein interpolation. The use of Bernstein approximation only requires the activation function to be continuous in (p_t, I_t), and can be used not only in more general situations, but also to obtain better polynomial approximations than Taylor expansions (see [18]). Intuitively, an order k Taylor approximation can only guarantee to have the same value as the approximated function at the expansion point, while an order k Bernstein interpolation has the same value as the approximated function at $k+1$ points. We give the details of our Bernstein overapproximation method as follows.

Bernstein approximation for $\sigma(\vec{z})$ w.r.t. $\vec{z} \in (p_t, I_t)$. Given (p_t, I_t) computed in Line 3, the j-th component of the polynomial vector $p_{\sigma,i}$ is the order k Bernstein interpolation of the activation function σ_j of the j-th neuron. It can be computed as $p_{\sigma_j,i}(z_j) = \sum_{s=0}^{k} \left(\sigma_j(\frac{\overline{Z}_j - \underline{Z}_j}{k} s + \underline{Z}_j) \binom{k}{s} \frac{(z_j - \underline{Z}_j)^s (\overline{Z}_j - z_j)^{k-s}}{(\overline{Z}_j - \underline{Z}_j)^k} \right)$, such that $\overline{Z}_j$ and $\underline{Z}_j$ denote the upper and lower bounds respectively of the range in the j-th dimension of (p_t, I_t), and they can be obtained by interval evaluating TM.

Evaluating the remainder $I_{\sigma,i}$. The j-th component $I_{\sigma_j,i}$ of $I_{\sigma,i}$ is computed as a conservative remainder for the polynomial $p_{\sigma_j,i}$, and it can be obtained as a symmetric interval $[-\epsilon_j, \epsilon_j]$ such that

$$\epsilon_j = \max_{s=1,\cdots,m} \left(\left| p_{\sigma_j,i}(\frac{\overline{Z}_j - \underline{Z}_j}{m}(s - \frac{1}{2}) + \underline{Z}_j) - \sigma_j(\frac{\overline{Z}_j - \underline{Z}_j}{m}(s - \frac{1}{2}) + \underline{Z}_j) \right| + L_j \cdot \frac{\overline{Z}_j - \underline{Z}_j}{m} \right)$$

wherein L_j is a Lipschitz constant of σ_j with the domain (p_t, I_t), and m is the number of samples that are uniformly selected to estimate the remainder. The soundness of the error bound estimation above has been proven in [11] for multivariate Bernstein polynomials. Since univariate Bernstein polynomials is a special case of multivariate Bernstein polynomials, our approach is also sound.

The following theorem states that a TM flowpipe computed by our approach is not only a range overapproximation of a reachable set, but also a function overapproximation for the dependency of a reachable state on its initial state.

Theorem 1. *If $\mathcal{F}(\vec{x}_0, \tau)$ is the i-th TM flowpipe computed in the j-st control step, then for any initial state $\vec{x}_0 \in X_0$, the box $\mathcal{F}(\vec{x}_0, \tau)$ contains the actual reachable state $\varphi_{\mathcal{N}}(\vec{x}_0, (j-1)\delta_c + (i-1)\delta + \tau)$ for all $\tau \in [0, \delta]$.*

3.2 Selection of Polynomial Approximations

Since an activation function is univariate, both of its Taylor and Bernstein approximations have a size which is linear in the order k. Then we investigate the accuracy produced by both approximation forms. Since the main operation in the TM layer-by-layer propagation framework is the composition of TMs, we study the *preservation of accuracy* for both of the forms under the composition with a given TM. We first define the *Accuracy Preservation Problem*.

When a function $f(\vec{x})$ is overapproximated by a TM $(p(\vec{x}), I)$ w.r.t. a bounded domain D, the approximation quality, i.e., size of the overestimation, is directly reflected by the width of I, since $f(\vec{x}) = p(\vec{x})$ for all $\vec{x} \in D$ when I is zero by the

TM definition. Given order k TMs $(p_1(\vec{x}), I_1)$ and $(p_2(\vec{x}), I_2)$ which are overapproximations of the same function $f(\vec{x})$ w.r.t. a bounded domain $D \subset \mathbb{R}^n$, we use $(p_1(\vec{x}), I_1) \prec_k (p_2(\vec{x}), I_2)$ to denote that the width of I_1 is smaller than the width of I_2 in all dimensions, i.e., $(p_1(\vec{x}), I_1)$ is a more accurate overapproximation of $f(\vec{x})$ than $(p_2(\vec{x}), I_2)$.

Accuracy Preservation Problem. Assume $(p_1(\vec{x}), I_1)$ and $(p_2(\vec{x}), I_2)$ are overapproximations of $f(\vec{x})$ with $\vec{x} \in D$, and $(p_1(\vec{x}), I_1) \prec_k (p_2(\vec{x}), I_2)$. Another function $g(\vec{y})$ is overapproximated by a TM $(q(\vec{y}), J)$ whose range is a subset of D, *does* $(p_1(q(\vec{y}) + J), I_1) \prec_k (p_2(q(\vec{y}) + J), I_2)$ *hold by order* k *TM arithmetic?*

We give the following counterexample to show that the answer is **no**, i.e., although $(p_1(\vec{x}), I_1)$ is more accurate than $(p_2(\vec{x}), I_2)$, the composition $p_1(q(\vec{y}) + J) + I_1$ might not be a better order k overapproximation than $p_2(q(\vec{y}) + J) + I_2$ for the composite function $f \circ g$. Given $p_1 = 0.5 + 0.25x - 0.02083x^3$, $I_1 =$ [-7.93e-5, 1.92e-4], and $p_2 = 0.5 + 0.24855x - 0.004583x^3$, $I_2 =$ [-2.42e-4, 2.42e-4], which are both TM overapproximations for the sigmoid function $f(x) = \frac{1}{1+e^{-x}}$ w.r.t. $x \in q(y) + J$ such that $q = 0.1y - 0.1y^2$, $J = [-0.1, 0.1]$, and $y \in [-1, 1]$. We have that $(p_1, I_1) \prec_3 (p_2, I_2)$, however after the compositions using order 3 TM arithmetic, the remainder of $p_1(q(y) + J) + I_1$ is $[-0.0466, 0.0477]$, while the remainder of $p_2(q(y) + J) + I_2$ is $[-0.0253, 0.0253]$, and we do not have $(p_1(q(y) + J), I_1) \prec_3 (p_2(q(y) + J), I_2)$.

Since the accuracy is not preserved under composition, we do not decide which approximation to choose directly based on the their remainders. Instead, we integrate an additional step in Algorithm 1 to replace line 4-6: for each activation function, we compute both Taylor and Bernstein overapproixmations, and choose the one that produces the smaller remainder interval I_r after composition.

3.3 Symbolic Remainders in Layer-by-Layer Propagation

We describe the use of symbolic remainders (SR) in the layer-by-layer propagation of computing an NN output TM. The method was originally proposed in [7] for reducing the overestimation of TM flowpipes in the reachability computation for nonlinear ODEs, we adapt it particularly for reducing the error accumulation in the TM remainders during the layer-by-layer propagation. Unlike the BP technique whose purpose is to obtain tighter TMs for activation functions, the use of SR only aims at reducing the overestimation accumulation in the composition of a sequence of TMs each of which represents the input range of a layer.

Consider the TM composition for computing the output TM of a single layer in Fig. 3, the output TM $p_{\sigma,i}(W_i(p_i(\vec{x}_0) + I_i) + B_i) + I_{\sigma,i}$ equals to $Q_i W_i p_i(\vec{x}_0) + Q_i W_i I_i + Q_i B_i + p^R_{\sigma,i}(W_i(p_i(\vec{x}_0) + I_i) + B_i) + I_{\sigma,i}$ such that Q_i is the matrix of the linear coefficients in $p_{\sigma,i}$, and $p^R_{\sigma,i}$ consists of the terms in $p_{\sigma,i}$ of the degrees $\neq 1$. Therefore, the remainder I_i in the second term can be kept symbolically such that we do not compute $Q_i W_i I_i$ out as an interval but keep its transformation matrix $Q_i W_i$ to the computation for the subsequent layers. Given the image S of an interval under a linear mapping, we use $\underline{S}$ to denote that it is kept symbolically, i.e., we keep the interval along with the transformation matrix, and $\overline{S}$ to denote that the image is evaluated as an interval.

Algorithm 2 TM output computation using symbolic remainders, input and output are the same as those in Algorithm 1

Setting $\mathcal{Q}$ as an empty array which can keep $M+1$ matrices;
Setting $\mathcal{J}$ as an empty array which can keep $M+1$ multidimensional intervals;
$\mathbb{J} \leftarrow 0$;
for $i = 1$ to $M+1$ **do**
  Computing the composite TM $(p_{\sigma,i}, I_{\sigma,i})$ using BP;
  Evaluating $q_i(\vec{x}_0) + J_i$ based on $\mathbb{J}$ and $\mathcal{Q}[1]I_1$; # $\mathcal{Q}[1]I_1 = I_1$ when $i = 1$
  $\mathbb{J} \leftarrow J_i$; $\Phi_i \leftarrow Q_i W_i$;
  for $j = 1$ to $i-1$ **do**
    $\mathcal{Q}[j] \leftarrow \Phi_i \cdot \mathcal{Q}[j]$;
  Adding Φ_i to $\mathcal{Q}$ as the last element;
  for $j = 2$ to i **do**
    $\mathbb{J} \leftarrow \mathbb{J} + \mathcal{Q}[j] \cdot \mathcal{J}[j-1]$;
  Adding J_i to $\mathcal{J}$ as the last element;
Computing an interval enclosure I_r for $\mathbb{J} + \mathcal{Q}[1]I_1$; # interval evaluation
return $q_{M+1}(\vec{x}_0) + I_r$.

Then we present the use of SR in layer-by-layer propagation. Starting from the NN input TM $(p_1(\vec{x}_0), I_1)$, the output TM of the first layer is computed as

$$\underbrace{Q_1W_1p_1(\vec{x}_0) + Q_1B_1 + p^R_{\sigma,1}(W_1(p_1(\vec{x}_0) + I_1) + B_1) + I_{\sigma,1}}_{q_1(\vec{x}_0)+J_1} + \underline{Q_1W_1I_1}$$

which can be kept in the form of $q_1(\vec{x}_0) + J_1 + \underline{Q_1W_1I_1}$. Using it as the input TM of the second layer, we have the following TM

$$\begin{aligned} & p_{\sigma,2}(W_2(q_1(\vec{x}_0) + J_1 + \underline{Q_1W_1I_1}) + B_2) + I_{\sigma,2} \\ =& \underbrace{Q_2W_2q_1(\vec{x}_0) + Q_2B_2 + p^R_{\sigma,2}(W_2(q_1(\vec{x}_0) + J_1 + \overline{Q_1W_1I_1}) + B_2) + I_{\sigma,2}}_{q_2(\vec{x}_0)+J_2} \\ & + \underline{Q_2W_2J_1} + \underline{Q_2W_2Q_1W_1I_1} \end{aligned}$$

for the output range of the second layer. Therefore the output TM of the i-th layer can be obtained as $q_i(\vec{x}_0) + \mathbb{J}_i + \underline{Q_iW_i \cdots Q_1W_1I_1}$ such that $\mathbb{J}_i = J_i + \underline{Q_iW_iJ_{i-1}} + \underline{Q_iW_iQ_{i-1}W_{i-1}J_{i-2}} + \cdots + \underline{Q_iW_i \cdots Q_2W_2J_1}$.

We present the SR method by Algorithm 2 in which we use two lists: $\mathcal{Q}[j]$ for $Q_iW_i \cdot \cdots \cdot Q_jW_j$ and $\mathcal{J}[j]$ for $\mathbb{J}_j$ to keep the intervals and their linear transformations. The symbolic remainder representation is replaced by its interval enclosure I_r at the end of the algorithm.

Time and space complexity. Although Algorithm 2 produces TMs with tighter remainders than Algorithm 1 because of the symbolic interval representations under linear mappings, it requires (1) two extra arrays to keep the intermediate matrices and remainder intervals, (2) two extra inner loops which perform $i-1$ and $i-2$ iterations in the i-th outer iteration. The size of $Q_iW_i \cdot \cdots \cdot Q_jW_j$ is determined by the rows in Q_i and the columns in W_j, and hence the maximum number of neurons in a layer determines the maximum size of the matrices in

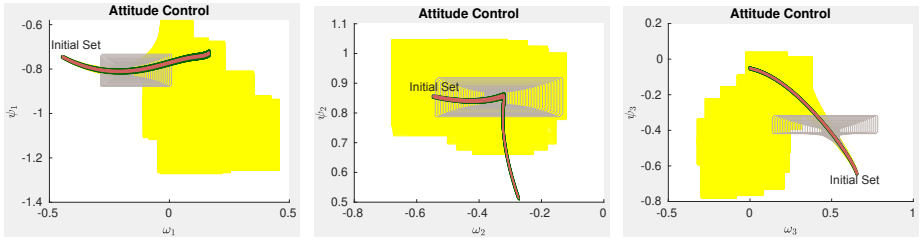

Fig. 4: Comparison between reachable sets of the 6-dimensional attitude control benchmark produced by POLAR (dark green), Verisig 2.0 (gray) and NNV (yellow). The red curves are simulated trajectories.

$\mathcal{Q}$. Similarly, the maximum dimension of J_i is also bounded by the maximum number of neurons in a layer. Because of the two inner loops, time complexity of Algorithm 2 is quadratic in M, whereas Algorithm 1 is linear in M.

Sizes of the TMs. All the TMs computed in the layer-by-layer propagation are over the same variables $\vec{x}_0$ which are symbolic representation for the NNCS initial set, i.e., $\vec{x}_0 \in X_0$. Therefore, the maximum size of an order k TM over n variables is bounded by $\binom{n+k}{n}$, and *hence the TM sizes are independent from the total number of neurons in the hidden layers of the NN controller.*

4 Experiments

We perform a comprehensive empirical study of POLAR against state-of-the-art (SOTA) techniques. We first demonstrate the performance of POLAR on two examples with high dimensional states and multiple inputs, which are far beyond the ability of SOTAs (Section 4.1). A comprehensive comparison over the full benchmarks in [11, 13] is then given (Section 4.2). Finally, we present ablation studies, scalability analysis, and the ability to handle discrete-time systems (Section 4.3). More detailed results can be found in the full version of the paper.

All our experiments were run on a machine with 6-core 2.90 GHz Intel Core i5 and 8GB of RAM. POLAR is implemented in C++. We present the results for POLAR, Verisig 2.0 and Sherlock using a single core without parallelization. The results of ReachNN* were computed on the same machine with the aid of GPU acceleration on an Nvidia GeForce RTX 2060 GPU.

State-of-the-art tools. We compare with SOTA tools in the NNCS reachability analysis literature, including Sherlock [8] (only works for ReLU), Verisig 2.0 [13] (only works for sigmoid and tanh), NNV [25], and ReachNN*[9][6].

4.1 High Dimensional Case Studies: Attitude Control & QUAD.

We consider an attitude control of a rigid body with 6 states and 3 control inputs [23], and quadrotor (QUAD) with 12 states and 3 control inputs [3] to

[6] The results of ReachNN* are based on GPU acceleration.

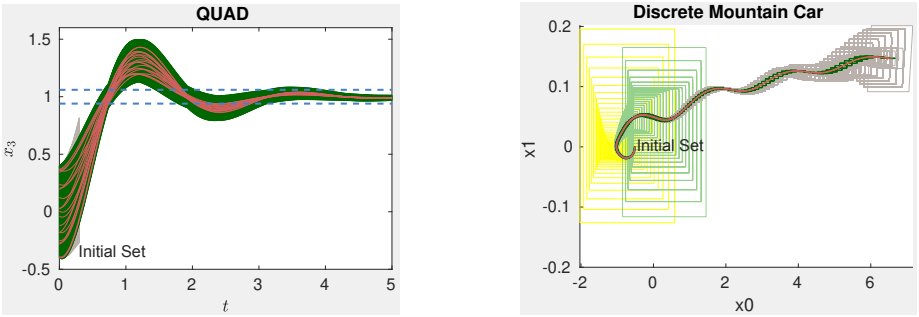

(a) QUAD (b) Mountain Car

Fig. 5: (a) Results of QUAD. POLAR for 50 steps (dark green sets), Verisig 2.0 for 3 steps (grey sets), and simulation traces for 50 steps (red curves). It took POLAR 1533 seconds to compute the flowpipes for 50 steps. On the other hand, it took Verisig 2.0 more than 5 hours to compute the flowpipes for the first 3 steps, and the TM remainders computed in the 4^{th} step are already of the size 10^{15}. NNV crashed with out-of-memory errors when computing the 1^{st} step. (b) Results of Mountain Car. POLAR for 150 steps (dark green sets), Verisig 2.0 for 150 steps (grey sets), ReachNN* for 90 steps (light green sets), NNV for 65 steps, and simulation traces for 150 steps (red curves).

evaluate the performance of POLAR on difficult problems. The complexity of these two example lies in the combination of the numbers of the state variables and control inputs. For each example, we trained a sigmoid neural-network controller and compare POLAR with Verisig 2.0 and NNV. The detailed setting of these two examples can be found in the full version of the paper.

The result for the attitude control benchmark is shown in Figure 4, and the result for the QUAD benchmark is shown in Figure 5a. In the attitude control benchmark, POLAR computed the TM flowpipes for 30 control steps in 201 seconds. From Figure 4, We can observe that the flowpipes computed by POLAR are tight w.r.t. the simulated traces. As a comparison, although Verisig 2.0 [13] can handle this system in theory, its remainder exploded very quickly and the tool crashed after only a few steps. NNV computed flowpipes for 25 steps by doing extensive splittings on the state space and crashed with out-of-memory errors. In the QUAD benchmark, POLAR computed the TM flowpipes for 50 control steps in 1533 seconds, while Verisig 2.0 and NNV took hours to compute flowpipes just for the first few steps.

4.2 Comparison over A Full Set of Benchmarks

We compare POLAR with the SOTA tools mentioned previously, including Sherlock, Verisig 2.0, NNV, and ReachNN* over the full benchmarks in [11, 13]. We refer to [11, 13] for more details of these benchmarks. The results are presented

Table 1: V: number of state variables, σ: activation functions, M: number of hidden layers, n: number of neurons in each hidden layer. For each approach, we give the runtime in seconds if it verifies the property. 'Unknown': the property is not verified. '–': the approach cannot be applied due to the type of σ.

#	V	NN Controller			**POLAR**	ReachNN*	Sherlock	Verisig 2.0
		σ	M	n		[9]	[8]	[13]
1	2	ReLU	2	20	**12**	26	42	–
		sigmoid	2	20	**17**	75	–	47
		tanh	2	20	**20**	Unknown	–	46
		ReLU+tanh	2	20	**13**	71	–	–
2	2	ReLU	2	20	**2**	5	3	–
		sigmoid	2	20	9	13	–	**7**
		tanh	2	20	**3**	73	–	Unknown
		ReLU+tanh	2	20	**2**	Unknown	–	–
3	2	ReLU	2	20	**16**	94	143	–
		sigmoid	2	20	**36**	146	–	44
		tanh	2	20	**26**	137	–	38
		ReLU+sigmoid	2	20	**15**	150	–	–
4	3	ReLU	2	20	**2**	8	21	–
		sigmoid	2	20	**3**	22	–	11
		tanh	2	20	**3**	21	–	10
		ReLU+tanh	2	20	**2**	12	–	–
5	3	ReLU	3	100	**13**	103	15	–
		sigmoid	3	100	76	**27**	–	190
		tanh	3	100	**76**	Unknown	–	179
		ReLU+tanh	3	100	**10**	Unknown	–	–
6	4	ReLU	3	20	**16**	1130	35	–
		sigmoid	3	20	**21**	13350	–	83
		tanh	3	20	**19**	2416	–	70
		ReLU+tanh	3	20	**15**	1413	–	–
ACC	6	tanh	3	20	**343**	Unknown	–	3344
QMPC	6	tanh	2	20	**61**	–[1]	–	652
Attitude Control	6	sigmoid	3	64	**201**	–[1]	–	Unknown
QUAD	12	sigmoid	3	64	**1533**	–[1]	–	Unknown

[1] This example has multi-dimensional control inputs. ReachNN* only supports NN controllers that produce single-dimensional control inputs.

in Table 1 where NNV is not included since we did not successfully use it to prove any of the benchmarks likely because it is designed for linear systems. Similar results for NNV are also observed in [13]. We can see POLAR successfully verifies all the cases and the runtime is **on average 8x and up to 94x faster**[7] compared with the tool with the second best efficiency. The "Unknown" verification results either indicate the overapproximation was too large for verifying the safety property or the tool terminated early due to an explosion of the overapproximation. POLAR achieves the best performance among all the tools.

We remark that the hyperparameter settings used by all of the three tools for the benchmarks in Table 1 were set to be the same for a fair, lateral comparison. However, they are not the best settings for POLAR. For example, POLAR finishes in **0.5 seconds** for the benchmark #1 with an integration stepsize that is same as the control stepsize and a TM order of 4.

[7] These are lower bounds on the improvements since other tools terminated early for certain settings due to explosion of their computed flowpipes.

4.3 Discussion

POLAR demonstrates substantial performance improvement over existing tools. In this section, we seek to further explore the capability of POLAR. We conduct several experiments for the QUAD benchmark to better understand the limitation and scalability of POLAR. We also include a mountain car example to show that POLAR is able to handle discrete-time systems.

Ablation Studies. To explore the impact of the two proposed techniques, namely Bernstein polynomial interpolation (BP) and symbolic remainder (SR) on the overall performance, we conduct a series of experiments on the QUAD benchmark with different configurations. Table 2 shows the performance of POLAR with and without the proposed techniques SR and BP in the NN propagation: 1) TM: only TM arithmetic is used; 2) TM+SR: SR is used with TM arithmetic; 3) BP is used with TM arithmetic; and 4) Both BP and SR are used with TM arithmetic. Based on the results, we can observe that SR significantly improves the accuracy of the reachable set overapproximation. Finally, the combination of basic TM with BP and SR not only achieves the best accuracy, but also is the most efficient. While the additional BP and SR operations can incur runtime overhead compared with basic TM, they help to produce a tighter overestimation and thus reduce the state space being explored during reachability analysis. As a result, the overall performance including runtime is better.

The following further observations can be obtained from Table 2. (i) Both of the independent use of BP and SR techniques significantly improves the performance of reachable set overapproximations. (ii) When the BP technique is used, Bernstein approximation is often not used on activation functions, but the few times for which they are used significantly improve the accuracy. The reason of having this phenomenon is that Taylor and Bernstein approximations are similarly accurate in approximating activation functions with small domain. However, the Lagrange form-based remainder evaluation in Taylor polynomials performs better than the sample-based remainder evaluation in Bernstein polynomials in those cases. It can also be seen that for each X_0, the use of Bernstein approximation becomes more frequent when the TMs has larger remainders. (iii) When both BP and SR techniques are used, the approach produces the tightest TMs compared with the other columns in the table even though the use Bernstein approximation is less often. The reason is that the remainders of the TMs are already well-limited and most of the activation functions handled in the reachability computation are with a "small" TM domain.

Scalability Analysis. Table 1 shows that POLAR can handle much larger NNCSs compared with the current SOTA. To better understand the scalability of POLAR, we further conduct scalability analysis on the size of the NN controller and the width of the initial set using the QUAD benchmark. The experiment results in Figure 6 for the neural networks with different widths and depths show that POLAR scales well on the number of layers and the number of neurons in each layer in the NN controller. On the other hand, the time cost grows rapidly when the width of the initial set becomes larger. Such a phenomenon already exists in the literature for reachability analysis of ODE systems [5]. The reason

Table 2: Ablation Studies for POLAR on the QUAD benchmark. We compare the width of TM remainder on x_3 at the 50^{th} step under different settings. For settings with BP, we also list the percentage of times where BP is used among 9600 neurons. If a setting cannot compute flowpipes for all 50 steps, it is marked as *Unknown*. X_0 is the radius of the initial set. k is the order of the TM.

X_0	k	TM		TM+SR		TM+BP			TM+BP+SR		
		Width	Time (s)	Width	Time (s)	Width	Time (s)	BP %	Width	Time (s)	BP %
0.05	2	7.5e-04	229	1.3e-04	233	6.8e-04	228	5.79%	1.2e-04	231	1.34%
	3	5.2e-04	273	6.5e-05	251	5.0e-04	274	3.62%	6.5e-05	251	0%
	4	4.9.e-04	332	6.2e-05	270	4.7e-04	336	3.57%	6.2e-05	270	0%
0.1	2	*Unknown*	–	2.3e-03	319	1.0e-02	325	9.68%	1.1e-03	289	4.80%
	3	1.8e-03	352	2.2e-04	287	1.7e-03	349	6.85%	2.2e-04	287	0%
	4	1.6e-03	431	1.9e-04	304	1.5e-03	427	6.70%	1.9e-04	304	0%
0.2	2	*Unknown*	–	*Unknown*	–	*Unknown*	–	–	*Unknown*	–	–
	3	9.0e-03	721	1.9e-03	412	7.8e-03	670	4.03%	1.6e-03	394	0.77%
	4	5.0e-03	761	9.2e-04	403	4.7e-03	728	4.38%	8.1e-04	396	0.07%
0.4	2	*Unknown*	–	*Unknown*	–	*Unknown*	–	–	*Unknown*	–	–
	3	*Unknown*	–	*Unknown*	–	*Unknown*	–	–	*Unknown*	–	–
	4	*Unknown*	–	*Unknown*	–	*Unknown*	–	–	3.7e-02	1533	3.25%

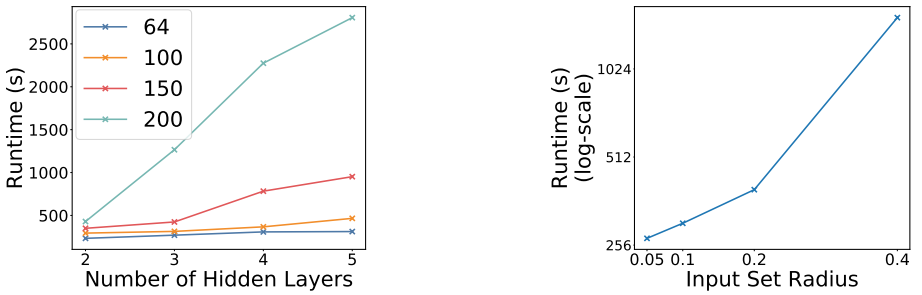

Fig. 6: Scalability analysis for POLAR on the QUAD benchmark. We present the runtime of QUAD for 50 steps reachability analysis. Under all settings, POLAR can verify that the system reaches the target set at the 50^{th} step. Left figure: Runtime on different neural network architectures with the input set radius as 0.05. We study neural-network controllers with different number of layers (2, 3, 4, 5) and neurons (64, 100, 150, 200). Right figure: Runtime on the different input set radius of the QUAD benchmark. We use the same network in Figure 5 which has 3 hidden layers with 64 neurons in each layer.

for this is that when the initial set is larger, it is more difficult to track the state dependencies and requires keeping more terms in a TM flowpipe.

Discrete-time NNCS. Finally, we use Mountain car, a common benchmark in Reinforcement Learning literature, to show that POLAR also works on discrete-time systems. The comparison with Verisig 2.0, ReachNN* and NNV is shown in Figure 5b. POLAR also outperforms these tools substantially for this example.

5 Conclusion

In this paper, we propose POLAR, a polynomial arithmetic framework, which integrates TM flowpipe construction, Bernstein overapproximation, and symbolic remainder method to efficiently compute reachable set overapproximations for NNCS. Empirical comparison shows POLAR performs significantly better than SOTAs on both computation efficiency and tightness of reachable set estimation. Our future work includes parallelization of POLAR on GPUs to further improve computation efficiency.

Acknowledgement

We gratefully acknowledge the support from the National Science Foundation awards CCF-1646497, CCF-1834324, CNS-1834701, CNS-1839511, IIS-1724341, CNS-2038853, ONR grant N00014-19-1-2496, and the US Air Force Research Laboratory (AFRL) under contract number FA8650-16-C-2642.

References

1. Althoff, M.: An introduction to CORA 2015. In: International Workshop on Applied veRification for Continuous and Hybrid Systems (ARCH). EPiC Series in Computing, vol. 34, pp. 120–151 (2015)
2. Alur, R., Dill, D.L.: A theory of timed automata. Theoretical computer science **126**(2), 183–235 (1994)
3. Beard, R.: Quadrotor dynamics and control rev 0.1 (2008)
4. Berz, M., Makino, K.: Verified integration of ODEs and flows using differential algebraic methods on high-order Taylor models. Reliable computing **4**, 361–369 (1998)
5. Chen, X.: Reachability Analysis of Non-Linear Hybrid Systems Using Taylor Models. Ph.D. thesis, RWTH Aachen University (2015)
6. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Flow*: An analyzer for non-linear hybrid systems. In: Proc. of CAV'13. LNCS, vol. 8044, pp. 258–263 (2013)
7. Chen, X., Sankaranarayanan, S.: Decomposed reachability analysis for nonlinear systems. In: Proc. of RTSS'16. pp. 13–24 (2016)
8. Dutta, S., Chen, X., Sankaranarayanan, S.: Reachability analysis for neural feedback systems using regressive polynomial rule inference. In: Proc. of HSCC'19. pp. 157–168. ACM (2019)
9. Fan, J., Huang, C., Chen, X., Li, W., Zhu, Q.: ReachNN*: A tool for reachability analysis of neural-network controlled systems. In: Proceedings of International Symposium on Automated Technology for Verification and Analysis (ATVA). LNCS, vol. 12302, pp. 537–542. Springer (2020)
10. Frehse, G., Guernic, C.L., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: Spaceex: Scalable verification of hybrid systems. In: Proceedings of International Conference on Computer Aided Verification (CAV). Lecture Notes in Computer Science, vol. 6806, pp. 379–395 (2011)
11. Huang, C., Fan, J., Li, W., Chen, X., Zhu, Q.: ReachNN: Reachability analysis of neural-network controlled systems. ACM Trans. Embed. Comput. Syst. **18**(5s), 106:1–106:22 (2019)

12. Huang, C., Fan, J., Li, W., Chen, X., Zhu, Q.: Divide and slide: Layer-wise refinement for output range analysis of deep neural networks. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD) (2020)
13. Ivanov, R., Carpenter, T., Weimer, J., Alur, R., Pappas, G.J., Lee, I.: Verisig 2.0: Verification of neural network controllers using taylor model preconditioning. In: Proc. of CAV'21. LNCS, vol. 12759, pp. 249–262. Springer (2021)
14. Ivanov, R., Weimer, J., Alur, R., Pappas, G.J., Lee, I.: Verisig: verifying safety properties of hybrid systems with neural network controllers. In: Proc. of HSCC'18. pp. 169–178. ACM (2019)
15. Jaulin, L., Kieffer, M., Didrit, O., Walter, É.: Interval analysis. In: Applied Interval Analysis. Springer (2001)
16. Katz, G., Barrett, C.W., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: An efficient SMT solver for verifying deep neural networks. In: Proc. of CAV'17. LNCS, vol. 10426, pp. 97–117. Springer (2017)
17. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. The Journal of Machine Learning Research **17**(1), 1334–1373 (2016)
18. Lorentz, G.G.: Bernstein Polynomials. American Mathematical Society (2013)
19. Makino, K., Berz, M.: Taylor models and other validated functional inclusion methods. International Journal of Pure and Applied Mathematics **4**(4), 379–456 (2003)
20. Meiss, J.D.: Differential Dynamical Systems. SIAM publishers (2007)
21. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529–533 (2015)
22. Moore, R.E., Kearfott, R.B., Cloud, M.J.: Introduction to Interval Analysis. SIAM (2009)
23. Prajna, S., Parrilo, P.A., Rantzer, A.: Nonlinear control synthesis by convex optimization. IEEE Transactions on Automatic Control **49**(2), 310–314 (2004)
24. Singh, G., Ganvir, R., Püschel, M., Vechev, M.T.: Beyond the single neuron convex barrier for neural network certification. In: Proc. of NeurIPS'19. pp. 15072–15083 (2019)
25. Tran, H., Yang, X., Lopez, D.M., Musau, P., Nguyen, L.V., Xiang, W., Bak, S., Johnson, T.T.: NNV: the neural network verification tool for deep neural networks and learning-enabled cyber-physical systems. In: Proc. of CAV'20. LNCS, vol. 12224, pp. 3–17. Springer (2020)
26. Wang, Z., Huang, C., Zhu, Q.: Efficient global robustness certification of neural networks via interleaving twin-network encoding. In: Proceedings of DAT2'22 (2022)
27. Weng, T.W., Zhang, H., Chen, H., Song, Z., Hsieh, C.J., Daniel, L., Dhillon, I.: Towards fast computation of certified robustness for relu networks. In: Proceedings of ICML'18 (2018)
28. Zhang, H., Weng, T.W., Chen, P.Y., Hsieh, C.J., Daniel, L.: Efficient neural network robustness certification with general activation functions. In: Proceedings of NeurIPS'18. pp. 4944–4953 (2018)
29. Zhu, Q., Huang, C., Jiao, R., Lan, S., Liang, H., Liu, X., Wang, Y., Wang, Z., Xu, S.: Safety-assured design and adaptation of learning-enabled autonomous systems. In: Proceedings of ASPDAC'21 (2021)
30. Zhu, Q., Li, W., Kim, H., Xiang, Y., Wardega, K., Wang, Z., Wang, Y., Liang, H., Huang, C., Fan, J., Choi, H.: Know the unknowns: Addressing disturbances and uncertainties in autonomous systems. In: Proceedings of ICCAD'20 (2020)