






# Star-Based Reachability Analysis of Deep Neural Networks

Hoang-Dung Tran<sup>1</sup>, Diago Manzanas Lopez<sup>1</sup>, Patrick Musau<sup>1</sup>,  
Xiaodong Yang<sup>1</sup>, Luan Viet Nguyen<sup>2</sup>, Weiming Xiang<sup>1</sup> ,  
and Taylor T. Johnson<sup>1</sup>  

<sup>1</sup> Institute for Software Integrated Systems, Vanderbilt University,  
Nashville, TN, USA

[taylor.johnson@vanderbilt.edu](mailto:taylor.johnson@vanderbilt.edu)

<sup>2</sup> Department of Computer and Information Science, University of Pennsylvania,  
Philadelphia, PA, USA

**Abstract.** This paper proposes novel reachability algorithms for both exact (sound and complete) and over-approximation (sound) analysis of deep neural networks (DNNs). The approach uses star sets as a symbolic representation of sets of states, which are known in short as stars and provide an effective representation of high-dimensional polytopes. Our star-based reachability algorithms can be applied to several problems in analyzing the robustness of machine learning methods, such as safety and robustness verification of DNNs. The star-based reachability algorithms are implemented in a software prototype called the neural network verification (NNV) tool that is publicly available for evaluation and comparison. Our experiments show that when verifying ACAS Xu neural networks on a multi-core platform, our exact reachability algorithm is on average about 19 times faster than Reluplex, a satisfiability modulo theory (SMT)-based approach. Furthermore, our approach can visualize the precise behavior of DNNs because the reachable states are computed in the method. Notably, in the case that a DNN violates a safety property, the exact reachability algorithm can construct a complete set of counterexamples. Our star-based over-approximate reachability algorithm is on average 118 times faster than Reluplex on the verification of properties for ACAS Xu networks, even without exploiting the parallelism that comes naturally in our method. Additionally, our over-approximate reachability is much less conservative than DeepZ and DeepPoly, recent approaches utilizing zonotopes and other abstract domains that fail to verify many properties of ACAS Xu networks due to their conservativeness. Moreover, our star-based over-approximate reachability algorithm obtains better robustness bounds in comparison with DeepZ and DeepPoly when verifying the robustness of image classification DNNs.

## 1 Introduction

Deep neural networks (DNNs) have become one of the most powerful techniques to deal with challenging and complex problems such as image processing [15]

© Springer Nature Switzerland AG 2019

M. H. ter Beek et al. (Eds.): FM 2019, LNCS 11800, pp. 670–686, 2019.

[https://doi.org/10.1007/978-3-030-30942-8\\_39](https://doi.org/10.1007/978-3-030-30942-8_39)

and natural language translation [9, 16] due to its learning ability on large data sets. Recently, the power of DNNs has inspired a new generation of intelligent autonomy which makes use of DNNs-based learning enable components such as autonomous vehicles [5] and air traffic collision avoidance systems [11]. Although utilizing DNNs is a promising approach, assuring the safety of autonomous applications containing neural network components is difficult because DNNs usually have complex characteristics and behavior that are generally unpredictable. Notably, it has been proved that well-trained DNNs may not be robust and are easily to be fooled by a slight change in the input [18]. Several recent incidents in autonomous driving (e.g., Tesla and Uber) raises an urgent need for techniques and tools that can formally verify the safety and robustness of DNNs before utilizing them in safety-critical applications.

Safety verification and robustness certification of DNNs have attracted a huge attention from different communities such as machine learning [1, 2, 13, 17, 20, 25, 26, 31], formal methods [6, 10, 12, 19, 23, 28–30], and security [7, 24, 25], and a recent survey of the area is available [27]. Analyzing the behavior of a DNN can broadly be categorized into exact and over-approximate analyses. For the exact analysis, the SMT-based [12] and polyhedron-based approaches [23, 28] are notable representatives. For the over-approximate analysis, the mixed-integer linear program (MILP) [6], interval arithmetic- [24, 25], zonotope- [20], input partition- [30], linearization- [26], and abstract-domain- [21] based are fast and efficient approaches. While the over-approximate analysis is usually faster and more scalable than the exact analysis, it guarantees only the soundness of the result. In contrast, the exact analysis is usually more time-consuming and less scalable. However, it guarantees both the soundness and completeness of the result [12]. Although the over-approximate analysis is fast and scalable, it is unclear how good the over-approximation is in term of conservativeness since the exact result is not available for comparison. Importantly, if an over-approximation approach is too conservative for neural networks with small or medium sizes, it will potentially produce huge conservative results for DNNs with a large number of layers and thousands of neurons since the over-approximation error is accumulated quickly over layers. Therefore, a scalable, exact reachability analysis is crucial not only for formal verification of DNNs, but also for estimating the conservativeness of current and up-coming over-approximation approaches.

In this paper, we propose a fast and scalable approach for the *exact* and *over-approximate* reachability analysis of DNN with ReLU activation functions using the concept of *star* sets [3], or shortly “star”. Star fits perfectly for the reachability analysis of DNNs due to its following essential characteristics: (1) an efficient (exact) representation of large input sets; (2) fast and cheap affine mapping operations; (3) inexpensive intersections with half-spaces and checking empty. *By utilizing star, we avoid the expensive affine mapping operation in polyhedron-based approach [23] and thus, reduce the verification time significantly.* Our approach performs reachability analysis for feedforward DNNs layer-by-layer. In the case of exact analysis, the output reachable set of each layer is a union of a set of stars. Based on this observation, *the star-based exact*

reachability algorithm naturally can be designed for efficient execution on multi-core platforms where each layer can handle multiple input sets at the same time. In the case of over-approximate analysis, the output reachable set of each layer is a single star which can be constructed by doing *point-wise* over-approximation of the reachable set at all neurons of the layer.

We evaluate the proposed algorithms in comparison with the polyhedron approach [23], Reluplex [12], zonotope [20] and abstract domain [21] approaches on safety verification of the ACAS Xu neural networks [11] and robust certification of image classification DNN. The experimental results show that our exact reachability algorithm can achieve 19 times faster than Reluplex when running on multi-core platform and >70 times faster than the polyhedron approach. Notably, our exact algorithm can visualize the precise behavior of the ACAS Xu networks and can construct the complete set of counter example inputs in the case that a safety property is violated. Our over-approximate reachability algorithm is averagely 118 times faster than Reluplex. It successfully verifies many safety properties of ACAS Xu networks while the zonotope and abstract domain approaches fail due to their large over-approximation errors. Our over-approximate reachability algorithm also provides a better robustness certification for image classification DNN in comparison with the zonotope and abstract domain approaches. In summary, the main contributions of this paper are: (1) propose novel, fast and scalable methods for the exact and over-approximate reachability analysis of DNNs; (2) implement the proposed methods in NNV toolbox that is available online for evaluation and comparison; (3) provide a thorough evaluation of the new methods via real-world case studies.

## 2 Preliminaries

### 2.1 Machine Learning Models and Symbolic Verification Problem

A feed-forward neural network (FNN) consists of an input layer, an output layer, and multiple hidden layers in which each layer comprises of neurons that are connected to the neurons of preceding layer labeled using weights. Given an input vector, the output of an FNN is determined by three components: the weight matrices  $W_k$ , representing the weighted connection between neurons of two consecutive layers  $k - 1$  and  $k$ , the bias vectors  $b_k$  of each layer, and the activation function  $f$  applied at each layer. Mathematically, the output of a neuron  $i$  is defined by:

$$y_i = f(\sum_{j=1}^n \omega_{ij} x_j + b_i),$$

where  $x_j$  is the  $j^{th}$  input of the  $i^{th}$  neuron,  $\omega_{ij}$  is the weight from the  $j^{th}$  input to the  $i^{th}$  neuron,  $b_i$  is the bias of the  $i^{th}$  neuron. In this paper, we are interested in FNN with ReLU activation functions defined by  $ReLU(x) = \max(0, x)$ .

**Definition 1 (Reachable Set of FNN).** *Given a bounded convex polyhedron input set defined as  $\mathcal{I} \triangleq \{x \mid Ax \leq b, x \in \mathbb{R}^n\}$ , and an  $k$ -layers feed-forward*

neural network  $F \triangleq \{L_1, \dots, L_k\}$ , the reachable set  $F(\mathcal{I}) = \mathcal{R}_{L_k}$  of the neural network  $F$  corresponding to the input set  $\mathcal{I}$  is defined incrementally by:

$$\begin{aligned} \mathcal{R}_{L_1} &\triangleq \{y_1 \mid y_1 = \text{ReLU}(W_1x + b_1), x \in \mathcal{I}\}, \\ \mathcal{R}_{L_2} &\triangleq \{y_2 \mid y_2 = \text{ReLU}(W_2y_1 + b_2), y_1 \in \mathcal{R}_{L_1}\}, \\ &\vdots \\ \mathcal{R}_{L_k} &\triangleq \{y_k \mid y_k = \text{ReLU}(W_ky_{k-1} + b_k) \ y_{k-1} \in \mathcal{R}_{L_{k-1}}\}, \end{aligned}$$

where  $W_k$  and  $b_k$  are the weight matrix and bias vector of the  $k^{\text{th}}$  layer  $L_k$ , respectively. The reachable set  $\mathcal{R}_{L_k}$  contains all outputs of the neural network corresponding to all input vectors  $x$  in the input set  $\mathcal{I}$ .

**Definition 2 (Safety Verification of FNN).** Given a  $k$ -layers feed-forward neural network  $F$ , and a safety specification  $\mathcal{S}$  defined as a set of linear constraints on the neural network outputs  $\mathcal{S} \triangleq \{y_k \mid Cy_k \leq d\}$ , the neural network  $F$  is called to be safe corresponding to the input set  $\mathcal{I}$ , we write  $F(\mathcal{I}) \models \mathcal{S}$ , if and only if  $\mathcal{R}_{L_k} \cap \neg\mathcal{S} = \emptyset$ , where  $\mathcal{R}_{L_k}$  is the reachable set of the neural network with the input set  $\mathcal{I}$ , and  $\neg$  is the symbol for logical negation. Otherwise, the neural network is called to be unsafe  $F(\mathcal{I}) \not\models \mathcal{S}$ .

## 2.2 Generalized Star Sets

**Definition 3 (Generalized Star Set [3]).** A generalized star set (or simply star)  $\Theta$  is a tuple  $\langle c, V, P \rangle$  where  $c \in \mathbb{R}^n$  is the center,  $V = \{v_1, v_2, \dots, v_m\}$  is a set of  $m$  vectors in  $\mathbb{R}^n$  called basis vectors, and  $P : \mathbb{R}^m \rightarrow \{\top, \perp\}$  is a predicate. The basis vectors are arranged to form the star's  $n \times m$  basis matrix. The set of states represented by the star is given as:

$$[\Theta] = \{x \mid x = c + \sum_{i=1}^m (\alpha_i v_i) \text{ such that } P(\alpha_1, \dots, \alpha_m) = \top\}. \tag{1}$$

Sometimes we will refer to both the tuple  $\Theta$  and the set of states  $[\Theta]$  as  $\Theta$ . In this work, we restrict the predicates to be a conjunction of linear constraints,  $P(\alpha) \triangleq C\alpha \leq d$  where, for  $p$  linear constraints,  $C \in \mathbb{R}^{p \times m}$ ,  $\alpha$  is the vector of  $m$ -variables, i.e.,  $\alpha = [\alpha_1, \dots, \alpha_m]^T$ , and  $d \in \mathbb{R}^{p \times 1}$ . A star is an empty set if and only if  $P(\alpha)$  is empty.

**Proposition 1.** Any bounded convex polyhedron  $\mathcal{P} \triangleq \{x \mid Cx \leq d, x \in \mathbb{R}^n\}$  can be represented as a star.<sup>1</sup>

**Proposition 2 [Affine Mapping of a Star].** Given a star set  $\Theta = \langle c, V, P \rangle$ , an affine mapping of the star  $\Theta$  with the affine mapping matrix  $W$  and offset vector  $b$  defined by  $\bar{\Theta} = \{y \mid y = Wx + b, x \in \Theta\}$  is another star with the following characteristics.

$$\bar{\Theta} = \langle \bar{c}, \bar{V}, \bar{P} \rangle, \bar{c} = Wc + b, \bar{v} = \{Wv_1, Wv_2, \dots, Wv_m\}, \bar{P} \equiv P.$$

<sup>1</sup> Proofs appear in the appendix of the extended version of this paper [22].

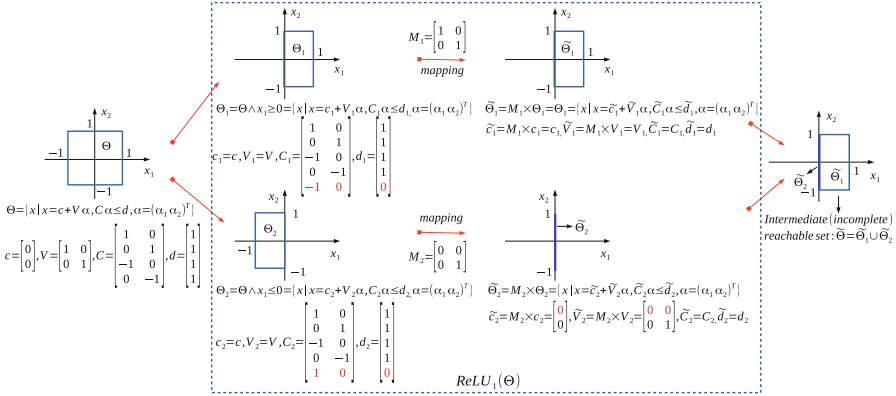


Fig. 1. An example of a stepReLU operation on a layer with two neurons.

**Proposition 3 (Star and Half-space Intersection).** *The intersection of a star  $\Theta \triangleq \langle c, V, P \rangle$  and a half-space  $\mathcal{H} \triangleq \{x \mid Hx \leq g\}$  is another star with following characteristics.*

$$\begin{aligned} \bar{\Theta} &= \Theta \cap \mathcal{H} = \langle \bar{c}, \bar{V}, \bar{P} \rangle, \quad \bar{c} = c, \quad \bar{V} = V, \quad \bar{P} = P \wedge P', \\ P'(\alpha) &\triangleq (H \times V_m)\alpha \leq g - H \times c, \quad V_m = [v_1 \ v_2 \ \dots \ v_m]. \end{aligned}$$

### 3 Star-Based Reachability Analysis of FNNs

#### 3.1 Exact and Complete Analysis

Since any bounded convex polyhedron can be represented as a star (Proposition 1), we assume the input set  $\mathcal{I}$  of an FNN is a star set. From Definition 1, one can see that the reachable set of an FNN is derived layer-by-layer. Since the affine mapping of a star is also a star (Proposition 2), the core step in computing the exact reachable set of a layer with a star input set is applying the ReLU activation function on the star input set, i.e., compute  $ReLU(\Theta)$ ,  $\Theta = \langle c, V, P \rangle$ . For a layer  $L$  with  $n$  neurons, the reachable set of the layer can be computed by executing a sequence of  $n$  stepReLU operations as follows  $\mathcal{R}_L = ReLU_n(ReLU_{n-1}(\dots ReLU_1(\Theta)))$ .

The stepReLU operation on the  $i^{th}$  neuron, i.e.,  $ReLU_i(\cdot)$ , works as follows. First, the input star set  $\Theta$  is decomposed into two subsets  $\Theta_1 = \Theta \wedge x_i \geq 0$  and  $\Theta_2 = \Theta \wedge x_i < 0$ . Note that from Proposition 3,  $\Theta_1$  and  $\Theta_2$  are also stars. Let assume that  $\Theta_1 = \langle c, V, P_1 \rangle$  and  $\Theta_2 = \langle c, V, P_2 \rangle$ . Since the later set has  $x_i < 0$ , applying the ReLU activation function on the element  $x_i$  of the vector  $x = [x_1 \ \dots \ x_i \ x_{i+1} \ \dots \ x_n]^T \in \Theta_2$  will lead to the new vector  $x' = [x_1 \ x_2 \ \dots \ 0 \ x_{i+1} \ \dots \ x_n]^T$ . This procedure is equivalent to mapping  $\Theta_2$  by the mapping matrix  $M = [e_1 \ e_2 \ \dots \ e_{i-1} \ 0 \ e_{i+1} \ \dots \ e_n]$ . Also, applying the ReLU activation function on the element  $x_i$  of the vector  $x \in \Theta_1$  does

---

**Algorithm 3.1.** Star-based exact reachability analysis for one layer.

---

**Input:**  $I = [\Theta_1 \dots \Theta_N]$ ,  $W, b$   $\triangleright$  star input sets, weight matrix, bias vector

**Output:**  $R$   $\triangleright$  exact reachable set

```

1: procedure  $R = \text{LAYERREACH}(I, W, b)$ 
2:    $R = \emptyset$ 
3:   parfor  $i = 1 : N$  do  $\triangleright$  parallel for loop
4:      $I_1 = W * \Theta_i + b = \langle Wc_i + b, WV_i, P_i \rangle$ 
5:      $R_1 = \text{reachReLU}(I_1)$ ,  $R = R \cup R_1$ 
6:   end parfor
7: procedure  $R_1 = \text{REACHReLU}(I_1)$ 
8:    $In = I_1$ 
9:    $[lb, ub] = In.\text{getRange}$   $\triangleright$  get ranges of all input variables
10:   $map = \text{find}(lb < 0)$   $\triangleright$  construct computation map
11:   $m = \text{length}(map)$   $\triangleright$  minimized number of stepReach operations
12:  for  $i = 1 : m$  do
13:     $In = \text{stepReLU}(In, map(i), lb(i), ub(i))$   $\triangleright$  stepReLU operation
14:   $R_1 = In$ 
15: procedure  $\tilde{R} = \text{STEPReLU}(\tilde{I}, i, lb_i, ub_i)$ 
16:   $\tilde{R} = \emptyset$ ,  $\tilde{I} = [\tilde{\Theta}_1 \dots \tilde{\Theta}_k]$   $\triangleright$  intermediate star input and output sets
17:  for  $j = 1 : k$  do
18:     $\mathcal{R}_1 = \emptyset$ ,  $M = [e_1 \ e_2 \ \dots \ e_{i-1} \ 0 \ e_{i+1} \ \dots \ e_n]$ 
19:    if  $lb_i \geq 0$  then  $\mathcal{R}_1 = \tilde{\Theta}_j = \langle \tilde{c}_j, \tilde{V}_j, \tilde{P}_j \rangle$ 
20:    if  $ub_i \leq 0$  then  $R_1 = M * \tilde{\Theta}_j = \langle M\tilde{c}_j, M\tilde{V}_j, \tilde{P}_j \rangle$ 
21:    if  $lb_i < 0$  &  $ub_i > 0$  then
22:       $\tilde{\Theta}'_j = \tilde{\Theta}_j \wedge x[i] \geq 0 = \langle \tilde{c}_j, \tilde{V}_j, \tilde{P}'_j \rangle$ ,  $\tilde{\Theta}''_j = \tilde{\Theta}_j \wedge x[i] < 0 = \langle \tilde{c}_j, \tilde{V}_j, \tilde{P}''_j \rangle$ 
23:       $\mathcal{R}_1 = \tilde{\Theta}'_j \cup M * \tilde{\Theta}''_j$ 
24:   $\tilde{R} = \tilde{R} \cup R_1$ 


```

---

not change the set since we have  $x_i \geq 0$ . Consequently, the result of the stepReLU operation on input set  $\Theta$  at the  $i^{th}$  neuron is a union of two star sets  $ReLU_i(\Theta) = \langle c, V, P1 \rangle \cup \langle Mc, MV, P2 \rangle$ . A concrete example of the first stepReLU operation on a layer with two neurons is depicted in Fig. 1.

The number of stepReLU operation can be reduced if we know beforehand the ranges of all states in the input set. For example, if we know that  $x_i$  is always larger than zero, then we have  $ReLU_i(\Theta) = \Theta$ , or in other words, we do not need to execute the stepReLU operation on the  $i^{th}$  neuron. Therefore, to minimize the number of stepReLU operations and the computation time, we first determine the ranges of all states in the input set which can be done efficiently by solving  $n$ -linear programming problems.

The star-based exact reachability algorithm given in Algorithm 3.1. works as follows. The layer takes the star output sets of the preceding layer as input sets  $I = [\Theta_1, \dots, \Theta_N]$ . The main procedure in the algorithm is *layerReach* which processes the input sets  $I$  in parallel. On each input element  $\Theta_i = \langle c_i, V_i, P_i \rangle$ , the main procedure maps the element with the layer weight matrix  $W$  and bias vector  $b$  which results a new star  $I_1 = \langle Wc_i + b, WV_i, P_i \rangle$ . The reachable set



**Fig. 2.** Over-approximation of ReLU functions with different approaches.

of the layer corresponding to the element  $\Theta_i$  is computed by *reachReLU* sub-procedure which executes a *minimized sequence* of *stepReLU* operations on the new star  $I_1$ , i.e., iteratively calls *stepReLU* sub-procedure. Note that that the *stepReLU* sub-procedure is designed to handle multiple star input sets since the number of star sets may increase after each *stepReLU* operation.

**Lemma 1.** *The worst-case complexity of the number of stars in the reachable set of an  $N$ -neurons FNN computed by Algorithm 3.1. is  $\mathcal{O}(2^N)$ .*

**Lemma 2.** *The worst-case complexity of the number of constraints of a star in the reachable set of an  $N$ -neuron FNN computed by Algorithm 3.1. is  $\mathcal{O}(N)$ .*

**Theorem 1 (Verification complexity).** *Let  $F$  be an  $N$ -neuron FNN,  $\Theta$  be a star set with  $p$  linear constraints and  $m$ -variables in the predicate,  $\mathcal{S}$  be a safety specification with  $s$  linear constraints. In the worst case, the safety verification or falsification of the neural network  $F(\Theta) \models \mathcal{S}$ ? is equivalent to solving  $2^N$  feasibility problems in which each has  $N + p + s$  linear constraints and  $m$ -variables.*

*Remark 1.* Although in the worst-case, the number of stars in the reachable set of an FNN is  $2^N$ , in practice, the actual number of stars is usually much smaller than the worst-case result which enhances the applicability of the star-based exact reachability analysis for practical DNNs.

**Theorem 2 (Safety and complete counter input set).** *Let  $F$  be an FNN,  $\Theta = \langle c, V, P \rangle$  be a star input set,  $F(\Theta) = \cup_{i=1}^k \Theta_i$ ,  $\Theta_i = \langle c_i, V_i, P_i \rangle$  be the reachable set of the neural network, and  $\mathcal{S}$  be a safety specification. Denote  $\bar{\Theta}_i = \Theta_i \cap \neg \mathcal{S} = \langle c_i, V_i, \bar{P}_i \rangle$ ,  $i = 1, \dots, k$ . The neural network is safe if and only if  $\bar{P}_i = \emptyset$  for all  $i$ . If the neural network violates its safety property, then the complete counter input set containing all possible inputs in the input set that lead the neural network to unsafe states is  $\mathcal{C}_\Theta = \cup_{i=1}^k \langle c, V, \bar{P}_i \rangle$ ,  $\bar{P}_i \neq \emptyset$ .*

### 3.2 Over-Approximate Analysis

Although the exact reachability analysis can compute the exact behavior of FNN, the number of stars grows exponentially with the number of layers and leads to

---

**Algorithm 3.2.** Star-based over-approximate reachability analysis for one layer.

---

**Input:**  $I = \Theta = \langle c, V, P \rangle$ ,  $W$ ,  $b$  ▷ star input set, weight matrix, bias vector  
**Output:**  $R$  ▷ over-approximate reachable set

```

1: procedure  $R = \text{APPROXLAYERREACH}(I, W, b)$ 
2:    $I_1 = W * \Theta + b = \langle Wc + b, WV, P \rangle$ 
3:    $In = I_1$ 
4:   for  $i = 1 : n$  do ▷  $n$  is the number of neurons of the layer
5:      $In = \text{approxStepReLU}(In, i)$  ▷  $i^{\text{th}}$  approximate-stepReLU operation
6:    $R_1 = In$ 
7: procedure  $\tilde{R} = \text{APPROXSTEPReLU}(\tilde{I}, i)$ 
8:    $\tilde{I} = \tilde{\Theta} = \langle \tilde{c}, \tilde{V}, \tilde{P} \rangle$ 
9:    $[l_i, u_i] = \tilde{\Theta}.\text{getRange}(i)$  ▷ range of  $x[i]$ , i.e.,  $l_i \leq x[i] \leq u_i$ 
10:   $M = [e_1 \ e_2 \ \dots \ e_{i-1} \ 0 \ e_{i+1} \ \dots \ e_n]$ 
11:  if  $l_i \geq 0$  then  $\tilde{R} = \tilde{\Theta} = \langle \tilde{c}, \tilde{V}, \tilde{P} \rangle$ 
12:  if  $u_i \leq 0$  then  $\tilde{R} = M * \tilde{\Theta} = \langle M\tilde{c}, M\tilde{V}, \tilde{P} \rangle$ 
13:  if  $l_i < 0$  &  $u_i > 0$  then
14:     $\tilde{P}(\alpha) \triangleq \tilde{C}\alpha \leq \tilde{d}$ ,  $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_m]^T$  ▷ input set's predicate
15:     $\alpha' = [\alpha_1, \dots, \alpha_m, \alpha_{m+1}]^T$  ▷ new variable  $\alpha_{m+1}$ 
16:     $C_1 = [0 \ 0 \ \dots \ 0 \ -1] \in \mathbb{R}^{1 \times m+1}$ ,  $d_1 = 0$  ▷  $\alpha_{m+1} \geq 0 \Leftrightarrow C_1\alpha' \leq d_1$ 
17:     $C_2 = [V[\tilde{i}, :] \ -1] \in \mathbb{R}^{1 \times m+1}$ ,  $d_2 = -\tilde{c}[\tilde{i}]$  ▷  $\alpha_{m+1} \geq x[\tilde{i}] \Leftrightarrow C_2\alpha' \leq d_2$ 
18:     $C_3 = [\frac{-u_i}{u_i-l_i} \times V[\tilde{i}, :] \ 1]$ ,  $d_3 = \frac{u_i l_i}{u_i-l_i} \times (1-\tilde{c}[\tilde{i}])$  ▷  $\alpha_{m+1} \leq \frac{u_i(x[\tilde{i}]-l_i)}{u_i-l_i} \Leftrightarrow C_3\alpha' \leq d_3$ 
19:     $C_0 = [C \ 0_{m \times 1}]$ ,  $d_0 = \tilde{d}$ 
20:     $C' = [C_0; C_1; C_2; C_3]$ ,  $d' = [d_0; d_1; d_2; d_3]$ 
21:     $P'(\alpha') \triangleq C'\alpha' \leq d'$  ▷ output set's predicate
22:     $\tilde{c}' = M\tilde{c}$ ,  $V' = M\tilde{V}$ ,  $V' = [V' \ e_i]$  ▷  $y[\tilde{i}] = \text{ReLU}(x[\tilde{i}]) = \alpha_{m+1}$ 
23:     $\tilde{R} = \langle \tilde{c}', V', P' \rangle$ 

```

---

an increase in computation cost that limits scalability. In this section, we propose an over-approximation reachability algorithm for FNNs. The main benefit of this approach is that the reachable set of each layer is only a single star that can be constructed efficiently by using an over-approximation of the ReLU activation function at all neurons in the layer. Importantly, our star-based over-approximate reachability algorithm is much less conservative than the zonotope-based [20] and abstract domain [21] based approaches in the way of approximating the ReLU activation function, shown in Fig. 2. The zonotope-based approach [20] over-approximates the ReLU activation function by a minimal parallelogram while the abstract-domain approach [21] over-approximates the ReLU activation function by a triangle. Our star-based approach also over-approximates the ReLU activation function with a triangle as in the abstract-domain approach. However, the abstract-domain approach only uses lower bound and upper bound constraints for the output  $y_i = \text{ReLU}(x_i)$  to avoid the state space explosion [21], for example, in Fig. 2, these constraints are  $y_i \geq 0$ ,  $y_i \leq u_i(x_i - l_i)/(u_i - l_i)$ . To obtain a tighter over-approximation, our approach uses three constraints for the output  $y_i$  instead. The over-approximation rule for a single neuron is given as follows,



---

**Algorithm 3.3.** Reachability analysis for a FNN.

---

**Input:**  $I = \Theta = \langle c, V, P \rangle$ ,  $L = [L_1 \ L_2 \ \dots \ L_k]$ , *scheme*  $\triangleright$  star input set, network’s layers, reachability scheme

**Output:**  $R \triangleright$  reachable set

```

1: procedure  $R = \text{REACH}(I, L, \textit{scheme})$ 
2:    $In = I$ 
3:   for  $i = 1 : k$  do  $\triangleright k$  is the number of layers on the network
4:     if scheme = exact then  $In = L_i.\textit{LayerReach}(In)$ 
5:     else if scheme = approx then  $In = L_i.\textit{ApproxLayerReach}(In)$ 
6:    $R = In$ 

```

---

$$\begin{cases} y_i = x_i & \text{if } l_i \geq 0 \\ y_i = 0 & \text{if } u_i \leq 0 \\ y_i \geq 0, y_i \leq \frac{u_i(x_i - l_i)}{u_i - l_i}, y_i \geq x_i & \text{if } l_i < 0 \text{ and } u_i > 0 \end{cases}$$

where  $l_i$  and  $u_i$  is the lower and upper bounds of  $x_i$ .

Similar to the exact approach, the over-approximate reachable set of a Layer with  $n$  neurons can be computed by executing a sequence of  $n$  *approximate-stepReLU* operations performing the above over-approximation rule. The over-approximate reachability algorithm for a single layer of FNN using star set given in Algorithm 3.2. works as follows. Given a star input set  $\Theta$ , the algorithm computes the affine mapping of the input set using the layer’s weight matrix and bias vector. The resulting star set is the input of a sequence of  $n$  *approximate-stepReLU* operations. An approximate-stepReLU operation first computes the lower and upper bounds of the state variable  $x[i]$  w.r.t the  $i^{th}$  neuron. If the lower bound is not negative (line 11), the approximate-stepReLU operation returns a new intermediate reachable set which is exactly the same as its input set. If the upper bound is not positive (line 12), the approximate-stepReLU operation returns a new intermediate reachable set which is the same as its input set except the  $i^{th}$  state variable is zero. If the lower bound is negative and the upper bound is positive (line 13), the approximate-stepReLU operation introduces a new variable  $\alpha_{m+1}$  to capture the over-approximation of ReLU function at the  $i^{th}$  neuron. As a result, the obtained intermediate reachable set has one more variable and three more linear constraints in the predicate in comparison with the corresponding input set. From this observation, we can see that in the worst case, the over-approximate reachability algorithm will obtain a reachable set with  $N + m_0$  variables and  $3N + n_0$  constraints in the predicate, where  $m_0, n_0$  respectively are the number of variables and linear constraints of the predicate of the input set and  $N$  is the total number of neurons of the FNN.

### 3.3 Reachability Algorithm for FNNs


The reachability analysis of a FNN is done layer-by-layer in which the output set of the previous layer is the input set of the next layer. The reachability algorithm for a FNN is summarized in Algorithm 3.3.

## 4 Evaluation

In this section, we evaluate the proposed star-based reachability algorithms in comparison to existing state-of-the-art approaches including exact (sound and complete) SMT-based (Reluplex [12]) and polyhedron-based [23] approaches, as well as over-approximate approaches, such as those using zonotopes [20] and abstract domains [21]. To clarify intuitively the benefit of our approach, we re-implement the zonotope- and abstract-domain based approaches in our tool called NNv. This allows the visualization of the over-approximate reachable set of these approaches. The evaluation and comparison are done by verifying safety of the ACAS Xu DNNs [11] and the robustness of image classification

**Table 1.** Safety verification results of ACAS Xu networks. Notation: *TO* is ‘Timeout’, **Rel** states for ‘Reluplex’, **Poly** is the Polyhedron method, *UN* states for unknown (due to over-approximation error), *VT* is the verification time in seconds, and  $r_{ES}^1, r_{ES}^2$  are the verification time improvement of the exact, star-based method compared with Reluplex and the polyhedron-based methods;  $r_Z, r_{AD}$  and  $r_{AS}$  respectively are the verification time improvement of the zonotope-, abstract domain- and over-approximate star-based methods compared with Reluplex. The computation time limitation for polyhedron-based method (run on Amazon cloud) was set to be 1 h while for Reluplex, it was set at 24 h.

Prop.	ID	Exact methods						Over-approximation methods								
		Res.	Rel	Poly	Star	$r_{ES}^1$	$r_{ES}^2$	Zonotope		Abstract-domain			Star			
			<i>VT</i>	<i>VT</i>	<i>VT</i>			Res.	<i>VT</i>	$r_Z$	Res.	<i>VT</i>	$r_{AD}$	Res.	<i>VT</i>	$r_{AS}$
$\phi_1$	$N_{1-1}$	Safe	5986	TO	1481.2	<b>4.04x</b>	–	Safe	0.07	<b>85514x</b>	Safe	6.22	<b>962.4x</b>	Safe	13.79	<b>434.1x</b>
	$N_{1-3}$	Safe	1102	TO	77.3	<b>14.3x</b>	> <b>47x</b>	UN	0.06	<b>18367x</b>	UN	6.3	<b>174.9x</b>	UN	6.01	<b>183.4x</b>
	$N_{2-1}$	Safe	1173	TO	51.57	<b>22.8x</b>	> <b>70x</b>	UN	0.062	<b>18919x</b>	UN	5.72	<b>205.1x</b>	Safe	5.91	<b>198.5x</b>
	$N_{2-2}$	Safe	634	TO	35.8	<b>17.7x</b>	> <b>101x</b>	UN	0.084	<b>7548x</b>	UN	5.83	<b>108.8x</b>	Safe	5.8	<b>109.3x</b>
	$N_{2-3}$	Safe	1014	TO	36.1	<b>28.1x</b>	> <b>100x</b>	UN	0.073	<b>13890x</b>	UN	5.89	<b>172.2x</b>	UN	5.79	<b>175.7x</b>
$\phi_3$	$N_{5-1}$	Safe	1097	TO	17.76	<b>61.8x</b>	> <b>202x</b>	UN	0.081	<b>13543x</b>	UN	5.85	<b>187.5x</b>	UN	5.76	<b>190.5x</b>
	$N_{1-5}$	Safe	393	1520.29	33.39	<b>11.8x</b>	<b>45.5x</b>	UN	0.0796	<b>4937x</b>	UN	6.1	<b>64.4x</b>	Safe	5.89	<b>66.7x</b>
	$N_{2-2}$	Safe	451	TO	43.66	<b>10.3x</b>	> <b>83x</b>	UN	0.056	<b>8054x</b>	UN	5.57	<b>81x</b>	Safe	5.66	<b>79.7x</b>
	$N_{2-3}$	Safe	293	2759	37.23	<b>7.8x</b>	<b>74.1x</b>	UN	0.08	<b>3663x</b>	UN	6.04	<b>74.7x</b>	Safe	5.66	<b>79.7x</b>
	$N_{2-8}$	Safe	653	1152.6	31.04	<b>21x</b>	<b>37.1x</b>	Safe	0.102	<b>6401x</b>	Safe	6.36	<b>102.7x</b>	Safe	7.47	<b>87.4x</b>
	$N_{2-9}$	Safe	61	233	6.33	<b>9.63x</b>	<b>36.8x</b>	Safe	0.065	<b>938.5x</b>	Safe	5.76	<b>10.6x</b>	Safe	5.88	<b>10.4x</b>
	$N_{3-7}$	Safe	357	1115.6	14	<b>25.5x</b>	<b>79.7x</b>	UN	0.085	<b>4200x</b>	Safe	5.88	<b>60.7x</b>	Safe	6	<b>59.5x</b>
	$N_{3-8}$	Safe	149	770	15.2	<b>9.8x</b>	<b>50.7x</b>	UN	0.08	<b>1862x</b>	UN	5.53	<b>26.9x</b>	Safe	5.83	<b>25.6x</b>
	$N_{3-9}$	Safe	715	1664.4	40.9	<b>17.5x</b>	<b>40.7x</b>	UN	0.076	<b>9408x</b>	UN	6.25	<b>114.4x</b>	Safe	6.13	<b>116.6x</b>
	$N_{4-9}$	Safe	489	1098.7	22.2	<b>22x</b>	<b>49.49x</b>	Safe	0.049	<b>9980x</b>	Safe	5.3	<b>92.3x</b>	Safe	5.57	<b>87.8x</b>
	$N_{5-1}$	Safe	585	1005.3	18.43	<b>31.74x</b>	<b>54.5x</b>	UN	0.069	<b>8479x</b>	UN	5.8	<b>100.9x</b>	Safe	5.83	<b>100.3x</b>
	$N_{5-7}$	Safe	42	275.1	7.69	<b>5.5x</b>	<b>35.8x</b>	Safe	0.054	<b>778x</b>	Safe	5.5	<b>7.6x</b>	Safe	5.54	<b>7.6x</b>
	<b>Average time improvement</b>						<b>≈ 18.9x</b>	<b>&gt; 70x</b>			<b>≈ 12734x</b>			<b>≈ 150x</b>		



(a) Verification times for property  $\phi_3$  on  $N_{2.8}$  net- (b) The (normalized) complete counter input set for work with different number property  $\phi'_3$  on  $N_{2.8}$  network is a part of the normalized cores. (normalized input set (red boxes)).

**Fig. 3.** Verification time reduction with parallel computing and complete counter input set construction. (Color figure online)

DNN against adversarial attacks. All results presented in this section and their corresponding scripts are available online<sup>2</sup>.

### 4.1 Safety Verification for ACAS Xu DNNs

The ACAS Xu networks are DNN-based advisory controllers that map the sensor measurements to advisories in the Airborne Collision Avoidance System X [11]. It consists of 45 DNNs which are trained to replace the traditional memory-consuming lookup table. Each DNN denoted by  $N_{x,y}$  has 5 inputs, 5 outputs, and 6 hidden layers of 50 neurons. The detail about ACAS Xu networks and their safety properties are given in the appendix [22]. The experiments in this case study are done using Amazon Web Services Elastic Computing Cloud (EC2), on a powerful `m5a.24xlarge` instance with 96 cores and 384 GB of memory. The verification results are presented in Table 1. We used 90 cores for the exact reachability analysis of the ACAS Xu networks using the polyhedron- and star-based approaches, and only 1 core for the over-approximate reachability analysis approaches.


**Verification Results and Timing Performance.** Safety verification using star-based reachability algorithms consists of two major steps. The first step constructs the whole reachable set of the networks. The second step checks the intersection of the constructed reachable set with the unsafe region. The verification time (VT) in our approach is the sum of the reachable set computation time (RT) and the safety checking time (ST). The reachable set computation time dominates (averagely 95% of) the verification time in all cases and the verification time varies for different properties. The detail of the reachable set computation time and the safety checking time can be found in the verification results.

<sup>2</sup> <https://github.com/verivital/nnv/tree/fm2019/nnv/examples/Submission/FM2019>.

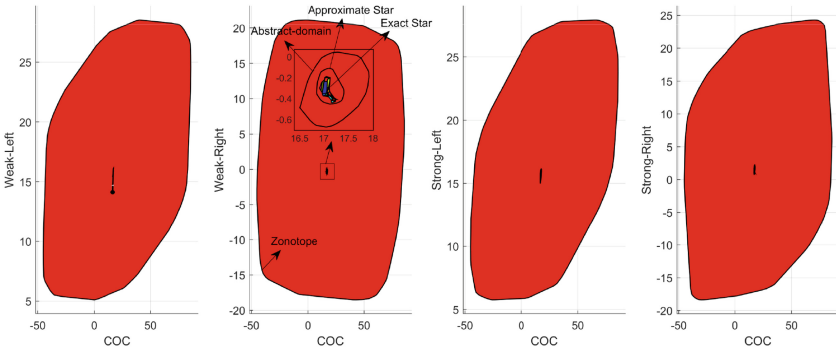
**Exact Star-Based Method.** The experimental results show that the exact star-based approach is on average  $>70$  times faster than the polyhedron-based approach and 18.9 times faster than Reluplex when using parallel computing. Impressively, it can even achieve 61.8 faster than Reluplex when verifying property  $\phi_2$  on  $N_{5,1}$  network. This improvement comes from the fact that star set that is very efficient in affine mapping and intersection with half-spaces which are crucial operations for reachable set computation and safety checking. Therefore, the exact star-based method is much more efficient than the polyhedron-based approach [23]. In addition, the exact star-based algorithm is well-designed and optimized (i.e., minimize the number of stepReLU operations) for efficiently running on multi-core platforms while Reluplex does not exploit the power of parallel-computing. Figure 3a describes the benefits of parallel computing. The figure shows that when a single core is used for a verification task, our approach takes 790.07s which is a little bit slower than Reluplex with 653s. However, our verification time drops quickly to 80.45s, which is 8 times faster than Reluplex, when we use 10 cores for the computation.

**Zonotope-Based Method** [20]. The experimental results show that the over-approximate, zonotope-based method is significantly faster than the exact methods. In some cases, it can verify the safety of the networks with less than 0.1 s, for example, the zonotope-based method successfully verifies property  $\phi_3$  on  $N_{5,7}$  network in 0.054s and the corresponding reachable set is depicted in Fig. 4a. Although the zonotope-based method is time-efficient (on average 12734 times faster than Reluplex), it is unable to verify the safety of many networks due to its huge over-approximation error, i.e., if the over-approximate reachable set reaches an unsafe region, we do not know whether or not the actual reachable set reaches the unsafe region. For example, Fig. 4b describes the reachable set obtained by the zonotope-based method for  $N_{3,8}$  network w.r.t property  $\phi_3$ . As shown in the figure, the obtained reachable set is too conservative and can not be used for safety verification of the network. The main reason that makes the zonotope approach fast is that, to do reachability analysis, we need to compute the lower and upper bounds of each state  $x[i]$  of all neurons in each layer. This information can be obtained straightforwardly in the zonotope method while in the other approaches, i.e., abstract-domain and star-based approaches, this is equivalent to solving  $n$  linear optimization problems where  $n$  is the number of neurons at that layer. The time for solving these optimization problems increase over layers since the number of constraints in the reachable set increases. Therefore, despite the a large over-approximation error, the zonotope-based method is *time-efficient* when dealing with large and deep neural networks.

**Abstract-Domain Based Method** [21]. The over-approximation method using abstract-domain is 150 times faster than Reluplex on average. It is also much less conservative than the zonotope-based method as can be seen from Fig. 4. However, the reachable set computed by the abstract-domain based method is still too conservative, which makes this approach unable to verify the safety properties of many ACAS Xu networks.



(a) Reachable sets of  $N_{5,7}$  network w.r.t property  $\phi_3$  with different methods. All methods successfully verify the property.



(b) Reachable sets of  $N_{3,8}$  network w.r.t property  $\phi_3$  with different methods. The zonotope (used in DeepZ [20]) and abstract-domain (used in DeepPoly [21]) methods cannot verify the property due to large over-approximation error.

**Fig. 4.** Conservativeness of the reachable sets obtained by different methods.

**Over-Approximate Star-Based Method.** The experiments show that our over-approximate star-based approach can obtain tight reachable sets for many networks compared to the exact sets. Therefore, our over-approximate approach successfully verifies safety properties of most of ACAS Xu networks. Notably, it is on average 118.4 times faster than Reluplex. Impressively, it is 434 times faster than Reluplex when verifying property  $\phi_1$  on  $N_{1,1}$  network. In comparison with the zonotope and abstract-domain approaches, our method is timing-comparable with the abstract-domain method and slower than the zonotope method. However, our results is much less conservative than those obtained by the zonotope and abstract-domain methods which are shown in Fig. 4. This makes our approach applicable for safety verification of many ACAS Xu networks where the zonotope and abstract-domain methods cannot verify.

**Benefits of Computing the Reachable Set.** The reachable set computed in our NNV tool are useful for intuitively checking the safety properties of the net-

**Table 2.** Maximum robustness values ( $\delta_{max}$ ) of image classification networks with different methods in which  $k$  is the number of hidden layers of the network,  $N$  is the total number of neurons,  $Tol$  is the tolerance error in searching.

Net	Parameters	Tol	$\delta_{max}$			
			Zonotope	Approximate-Star	Abstract-Domain	Exact-Star
$N_1$	$k = 5, N = 140$	0.0001	0.0046	0.0048	0.0046	$\geq 0.0058$
$N_2$	$k = 5, N = 250$	0.0001	0.0087	0.0101	0.0095	TimeOut
$N_3$	$k = 2, N = 1000$	0.0001	0.0072	0.0089	0.0084	TimeOut
$N_4$	$k = 1, N = 2000$	0.0001	0.0027	0.0027	0.0027	TimeOut
$N_5$	$k = 1, N = 4000$	0.0001	0.0034	0.0034	0.0034	TimeOut

work. For example, Fig. 4a describes the behaviors of  $N_{5,7}$  network corresponding to property  $\phi_3$  requiring that the output  $COC$  is not the minimal score. From the figure, one can see that the  $COC$  is not the minimal score and thus, property  $\phi_3$  holds on  $N_{5,7}$  network. Importantly, as shown in the figure, via visualization, one can intuitively observe the conservativeness of different over-approximation approaches in comparison to the exact ones which is impossible if we use *ERAN*, a *C-Python* implementation of the zonotope and abstract-domain-based methods. Last but not least, the reachable set is useful in the case that we need to verify a set of safety properties corresponding to the same input set. In this case, once the reachable set is obtained, it can be re-used to check different safety properties without rerunning the whole verification procedure as Reluplex does, and thus helps saving a significant amount of time.

**Complete Counter Example Input Set Construction.** Another strong advantage of our approach in comparison with other existing approaches is, in the case that a neural network violates its safety specification, our exact, star-based method can construct a *complete counter input set* that leads the neural network to the unsafe region. The complete counter input set can be used as an adversarial input generator [4, 8] for robust training of the network. We note that finding a single counter input falsifying a safety property of a neural network can be done efficiently using only random simulations. However, constructing a complete counter input set that contains all counter inputs is very challenging because of the non-linearity of a neural network. To the best of our knowledge, our exact star-based approach is the only approach that can solve this problem. For example, assume that we want to check the following property  $\phi'_3 \triangleq \neg(COC \geq 15.8 \wedge StrongRight \leq 15.09)$  on  $N_{2,8}$  network with the same input constraints as in property  $\phi_3$ . Using the available reachable set of  $N_{2,8}$  network, we can verify that the above property  $\phi'_3$  is violated in which 60 stars in 421 stars of the reachable set reach the unsafe region. Using Theorem 2, we can construct a complete counter input set which is a union of 60 stars in 0.9893s. This counter input set depicted in Fig. 3b is a part of the input set that contains all counter inputs that make the neural network unsafe.

## 4.2 Maximum Robustness Certification of Image Classification DNNs

Robustness certification of DNNs becomes more and more important as many safety-critical applications using image classification DNNs can be fooled easily by slightly perturbing a correctly classified input. A network is said to be  $\delta$ -locally-robust at input point  $x$  if for every  $x'$  such that  $\|x - x'\|_\infty \leq \delta$ , the network assigns the same label to  $x$  and  $x'$ . In this case study, instead of proving the robustness of a network corresponding to a given robustness certification  $\delta$ , we focus on finding the maximum robustness certification value  $\delta_{max}$  that a verification method can provide a robustness guarantee for the network. We investigate this interesting problem on a set of image classification DNN with different architectures trained (with an accuracy of 98%) using the well-known MNIST data set consisting of 60000 images of handwritten digits with a resolution of  $28 \times 28$  pixels [14]. The trained networks have 784 inputs and a single output with expected value from 0 to 9. We find the maximum robustness verification value  $\delta_{max}$  for the networks on an image of digit one with the assumption that there is a  $\delta_{max}$ -bounded disturbance modifying the (normalized) values of the input vector  $x$  at all pixels of the image, i.e.,  $|x[i] - x'[i]| \leq \delta_{max}$ . The result are presented in Table 2. We note that the polyhedron and Reluplex approaches are not applicable for these networks because they cannot deal with high dimensional input space. The table shows that our approximate star approach produces larger upper bounds of the robustness values of the networks with many layers. For single layer networks, our approach gives the same results as the zonotope [20] and the abstract domain [21] methods. The exact-star method can prove that the network  $N_1$  is robust with the bounded disturbance  $\delta = 0.0058$ . When  $\delta > 0.0058$ , we ran into the “out of memory” issue in parallel computation since the number of the reachable sets becomes too large. The exact star method reaches timeout (set as 1 h) when finding the maximum robustness value for the other networks.

## 5 Conclusion and Future Work

We have proposed two reachability analysis algorithms for DNNs using star sets, one that is exact (sound and complete) but has scalability challenges and one that over-approximates (sound) with better scalability. The exact algorithm can compute and visualize the exact behaviors of DNNs. The exact method is more efficient than standard polyhedra approaches, and faster than SMT-based approaches when running on multi-core platforms. The over-approximate algorithm is much faster than the exact one, and notably, it is much less conservative than recent zonotope and abstract-domain based approaches. Our algorithms are applicable for real world applications as shown in the safety verification of ACAS Xu DNNs and robustness certification of image classification DNNs. In future work, we are extending the proposed methods for convolutional neural networks (CNN) and recurrent neural networks (RNN), as well as improving scalability for other types of activation functions such as tanh and sigmoid.

**Acknowledgments.** We thank Gagandeep Singh from ETH Zurich for his help on explaining DeepZ and DeepPoly methods as well as running his tool ERAN. We also thank Shiqi Wang from Columbia University for his explanation about his interval propagation method, and Guy Katz from The Hebrew University of Jerusalem for his explanation about ACAS Xu networks and Reluplex. The discussions with Gagandeep Singh, Shiqi Wang, and Guy Katz is the main inspiration of our work in this paper. The material presented in this paper is based upon work supported by the Air Force Office of Scientific Research (AFOSR) through contract number FA9550-18-1-0122, and the Defense Advanced Research Projects Agency (DARPA) through contract number FA8750-18-C-0089. The U.S. government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of AFOSR or DARPA.

## References

1. Akintunde, M., Lomuscio, A., Maganti, L., Pirovano, E.: Reachability analysis for neural agent-environment systems. In: Sixteenth International Conference on Principles of Knowledge Representation and Reasoning (2018)
2. Akintunde, M.E., Kevorchian, A., Lomuscio, A., Pirovano, E.: Verification of RNN-based neural agent-environment systems. In: Proceedings of the 33th AAAI Conference on Artificial Intelligence (AAAI19), Honolulu, HI, USA. AAAI Press (2019, to appear )
3. Bak, S., Duggirala, P.S.: Simulation-equivalent reachability of large linear systems with inputs. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 401–420. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63387-9\\_20](https://doi.org/10.1007/978-3-319-63387-9_20)
4. Bastani, O., Ioannou, Y., Lampropoulos, L., Vytiniotis, D., Nori, A., Criminisi, A.: Measuring neural net robustness with constraints. In: Advances in Neural Information Processing Systems, pp. 2613–2621 (2016)
5. Bojarski, M., et al.: End to end learning for self-driving cars. arXiv preprint [arXiv:1604.07316](https://arxiv.org/abs/1604.07316) (2016)
6. Dutta, S., Jha, S., Sanakaranarayanan, S., Tiwari, A.: Output range analysis for deep neural networks. arXiv preprint [arXiv:1709.09130](https://arxiv.org/abs/1709.09130) (2017)
7. Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M.: Ai 2: safety and robustness certification of neural networks with abstract interpretation. In: 2018 IEEE Symposium on Security and Privacy (SP) (2018)
8. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint [arXiv:1412.6572](https://arxiv.org/abs/1412.6572) (2014)
9. Hinton, G., et al.: Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* **29**(6), 82–97 (2012)
10. Huang, X., Kwiatkowska, M., Wang, S., Wu, M.: Safety verification of deep neural networks. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 3–29. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63387-9\\_1](https://doi.org/10.1007/978-3-319-63387-9_1)
11. Julian, K.D., Kochenderfer, M.J., Owen, M.P.: Deep neural network compression for aircraft collision avoidance systems. arXiv preprint [arXiv:1810.04240](https://arxiv.org/abs/1810.04240) (2018)
12. Katz, G., Barrett, C., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: an efficient SMT solver for verifying deep neural networks. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 97–117. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63387-9\\_5](https://doi.org/10.1007/978-3-319-63387-9_5)



13. Kouvaros, P., Lomuscio, A.: Formal verification of CNN-based perception systems. arXiv preprint [arXiv:1811.11373](https://arxiv.org/abs/1811.11373) (2018)
14. LeCun, Y.: The MNIST database of handwritten digits (1998). <http://yann.lecun.com/exdb/mnist/>
15. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017)
16. Liu, W., Wang, Z., Liu, X., Zeng, N., Liu, Y., Alsaadi, F.E.: A survey of deep neural network architectures and their applications. *Neurocomputing* **234**, 11–26 (2017)
17. Lomuscio, A., Maganti, L.: An approach to reachability analysis for feed-forward ReLU neural networks. arXiv preprint [arXiv:1706.07351](https://arxiv.org/abs/1706.07351) (2017)
18. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: DeepFool: a simple and accurate method to fool deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2574–2582 (2016)
19. Pulina, L., Tacchella, A.: An abstraction-refinement approach to verification of artificial neural networks. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 243–257. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-14295-6\\_24](https://doi.org/10.1007/978-3-642-14295-6_24)
20. Singh, G., Gehr, T., Mirman, M., Püschel, M., Vechev, M.: Fast and effective robustness certification. In: Advances in Neural Information Processing Systems, pp. 10825–10836 (2018)
21. Singh, G., Gehr, T., Püschel, M., Vechev, M.: An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.* **3**(POPL), 41 (2019)
22. Tran, H.D., et al: Star-based reachability analysis of deep neural networks: extended version. In: 23rd International Symposium on Formal Methods (2019). [http://www.taylortjohnson.com/research/tran2019fm\\_extended.pdf](http://www.taylortjohnson.com/research/tran2019fm_extended.pdf)
23. Tran, H.D., et al.: Parallelizable reachability analysis algorithms for feed-forward neural networks. In: 7th International Conference on Formal Methods in Software Engineering (FormalISE 2019), Montreal, Canada (2019)
24. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Efficient formal safety analysis of neural networks. In: Advances in Neural Information Processing Systems, pp. 6369–6379 (2018)
25. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Formal security analysis of neural networks using symbolic intervals. arXiv preprint [arXiv:1804.10829](https://arxiv.org/abs/1804.10829) (2018)
26. Weng, T.W., et al.: Towards fast computation of certified robustness for ReLU networks. arXiv preprint [arXiv:1804.09699](https://arxiv.org/abs/1804.09699) (2018)
27. Xiang, W., et al.: Verification for machine learning, autonomy, and neural networks survey. *CoRR* abs/1810.01989 (2018)
28. Xiang, W., Tran, H.D., Johnson, T.T.: Reachable set computation and safety verification for neural networks with ReLU activations. arXiv preprint [arXiv:1712.08163](https://arxiv.org/abs/1712.08163) (2017)
29. Xiang, W., Tran, H.D., Johnson, T.T.: Output reachable set estimation and verification for multilayer neural networks. *IEEE Trans. Neural Netw. Learn. Syst.*, 1–7 (2018)
30. Xiang, W., Tran, H.D., Johnson, T.T.: Specification-guided safety verification for feedforward neural networks. In: AAAI Spring Symposium on Verification of Neural Networks (2019)
31. Zhang, H., Weng, T.W., Chen, P.Y., Hsieh, C.J., Daniel, L.: Efficient neural network robustness certification with general activation functions. In: Advances in Neural Information Processing Systems, pp. 4944–4953 (2018)