Safety Verification of Cyber-Physical Systems with Reinforcement Learning Control

Hoang-Dung Tran
Vanderbilt University
trhoangdung@gmail.com

FeiYang Cai
Vanderbilt University
trhoangdung@gmail.com

Manzanas Lopez Diego
Vanderbilt University
diego.manzanas.lopez@vanderbilt.edu

Patrick Musau
Vanderbilt University
patrick.musau@vanderbilt.edu

Weiming Xiang
Vanderbilt University
xiangwming@gmail.com

Luan Viet Nguyen
University of Pennsylvania
luanvn@seas.upenn.edu

Taylor T. Johnson
Vanderbilt University
taylor.johnson@vanderbilt.edu

Xenofon Koutsoukos
Vanderbilt University
Xenofon.Koutsoukos@vanderbilt.edu

ABSTRACT

This paper proposes a new forward reachability analysis approach to verify the safety of cyber-physical systems (CPS) with reinforcement learning controllers. The foundation of our approach lies on two efficient, exact and over-approximate reachability algorithms for neural network control systems using star set which is an efficient representation of polyhedron. Using these algorithms, we determine the initial conditions for which a safety-critical system with a neural network controller is safe by incrementally searching a critical initial condition where the safety of the system cannot be proved. Our approach produces tight over-approximation error and it is computational efficient which allows the application to practical CPS with learning enable components (LECs). We implement our approach in NNV [20], a recent verification tool for neural network control systems, and evaluate its advantages and applicability by verifying the safety of a practical Advanced Emergency Braking System (AEBS) with a reinforcement learning (RL) controller trained using deep deterministic policy gradient (DDPG) method. The experimental results show that our new reachability algorithms are much less conservative than the polyhedron-based approach [20, 23, 24]. We successfully determine the entire region of the initial conditions of the AEBS with the RL controller such that the safety of the system is guaranteed while the polyhedron-based approach cannot prove the safety properties of the system.

ACM Reference Format:
Hoang-Dung Tran, FeiYang Cai, Manzanas Lopez Diego, Patrick Musau, Weiming Xiang, Luan Viet Nguyen, Taylor T. Johnson, and Xenofon Koutsoukos. 2019. Safety Verification of Cyber-Physical Systems with Reinforcement Learning Control. In *Proceedings of* . ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Deep neural networks have become a popular choice in complex control applications where the control tasks are much more complicated than the traditional control problems. Recently, the power of DNNs has inspired a new generation of intelligent autonomy that makes use of DNNs as learning-based controllers such as autonomous vehicles [4] and air traffic collision avoidance systems [11]. Although utilizing DNNs for intelligent autonomous application is promising, safety verification of autonomy containing neural network components is difficult because DNNs usually have complex characteristics and behavior that are generally unpredictable. Importantly, many pieces of research have proved that well-trained DNNs may not robust and behave unsafely with a slight change in the input [15]. Recent incidents in autonomous driving (e.g., Tesla and Uber) raises an urgent need for techniques and tools that can formally verify the safety of neural network control systems before utilizing them in safety-critical applications.

Safety verification of neural network control systems (NNCS) is a challenging problem because the behaviors of the systems are difficult to estimate or characterize. To explicitly analyze the safety of NNCS, we need to calculate the exact or overapproximate reachable set containing all possible trajectories of the plant that takes the control set from the neural network controller as inputs. The output set of the plant is feedback to the controller to compute the control set for the next control step. Therefore, if the error in the reachable set computation is large, it quickly becomes larger and larger over time which results in too conservative reachable sets that cannot be used for safety verification. In addition, the scalability and efficiency of the reachable set computation are crucial for safety verification of control systems with DNNs controllers. It is required methods that can compute the reachable set of NNCS with large neural network controllers with a reasonable computation time and a small over-approximation error. However, calculating an exact or tight, overapproximate reachable set of a neural network quickly is fundamentally difficult due to the non-linearity of the network. This challenging problem has not addressed well in the existing literature.

In this paper, we propose a new reachability analysis approach for safety verification of CPS with neural network controllers using the concept of star set. We particularly focus on the safety verification of the Advanced Emergency Braking System (AEBS) in an autonomous car to illustrate and evaluate our approach. The AEBS is controlled by a neural network controller which is trained to stop the vehicle appropriately if it discovers an obstacle on the road. To

, ,
2019. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

guarantee safety, it is required that the time-to-collision (TTC) of the car, *which is a nonlinear function of the car's velocity, acceleration and the distance between the vehicle to the obstacle*, is always larger than a safe threshold defined by the physical characteristics of the vehicle. Our safety verification approach for AEBS works as follows. First, using CARLA, we perform system identification to obtain a discrete, linear state-space model of the car. The car model is then validated via systematic testing. Second, we train a deep neural network controller to perform the emergency braking action using reinforcement learning. Third, we compose the neural network controller with the state-space model to construct a closed-loop Simulink model of the AEBS which is then validated with CARLA results. Fourth, we perform the reachability analysis of the closed-loop model to obtain the reachable set of the AEBS. Finally, we compute the reachable set of the TTC and use it for safety verification.

We limit our reachability analysis approach to feed-forward neural network controllers with the ReLU/Saturation activation functions. Our reachability algorithms can compute both exact and over-approximate reachable sets of the AEBS. Exact reachable set computation is expensive since the number of the reachable sets increases over time steps. In contrast, the over-approximate reachability scheme is much cheaper as it produces a single reachable set at each time step. Importantly, by using star sets, our reachability analysis approach can eliminate or reduce significantly the over-approximation errors which is the main reason that makes the obtained reachable sets more and more conservative over time as shown in the polyhedron approach [20, 23, 24] (and maybe in some existing methods). Our approach successfully verifies the safety of the AEBS and notably, determine the entire region of the initial conditions of the AEBS where safety is guaranteed. This demonstrates the promising applicability of our approach in verifying safety properties of neural network-based autonomous systems at design time. We note that the polyhedron approach fails to prove the safety property of the system due to its over-approximation errors explode quickly over time.

In summary, the main contributions of this study are as follows.

(1) the provision of star-based reachability schemes designed to efficiently compute the reachable set of an discrete, linear neural network control systems with ReLU activation function,
(2) an end-to-end design and implementation of these schemes in a MATLAB®toolbox called NNV [20] which is publicly available for evaluation and comparison,
(3) and a thorough evaluation on the safety verification of the practical automatic emergency braking system.

2 SYSTEM MODEL AND PROBLEM FORMULATION

2.1 System model

In this paper, we are interested in safety verification of CPS with neural network controllers as depicted in Figure 1 in which $x(k)$ and $y(k)$ are the state and the output of the plant at the time step k. The controller is a feedforward neural network (FNN) consisting of an input layer, an output layer, and multiple hidden layers. Each layer is comprised of neurons that are connected to the neurons of the preceding layer labeled using weights [9]. The output of the FNN controller, given a specific input vector is determined by three components: the weight matrices $W_{l,l-1}$, representing the weighted connection between neurons of two consecutive layers $l-1$ and l, the bias vectors b_l of each layer, and the activation function f applied at each layer. Formally, the output of a neuron i is defined by:

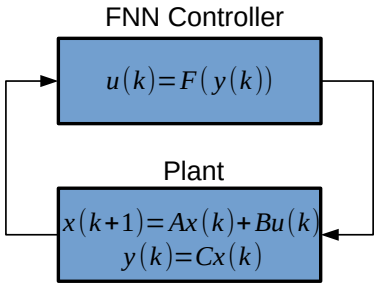

Figure 1: Neural network control system (NNCS).

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

where x_j is the j^{th} input of the i^{th} neuron, ω_{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 consider FNN controller with the ReLU activation functions defined as $ReLU(x) = max(0, x)$.

2.2 Problem formulation

Problem 1 (Safety Verification of NNCS). *Given a CPS with an FNN controller F, and a discrete, linear plant P with the initial states $x(0)$ in an initial set X_0, verify whether or not the state of the plant satisfies a safety property in a bounded time steps k_{max}. Formally, we want to verify if $\forall x(0) \in X_0 \rightarrow g(x(k)) \models S(g(x(k))), \forall 0 \leq k \leq k_{max}$ in which g is a nonlinear transformation function, S is a linear predicate over the transformed state variables $g(x(k))$ defining the safety requirements of the system.*

The core challenges in problem 1 are: 1) given the initial set of states of the plant, how we efficiently compute the reachable set of the plant over time steps which depends on the control input produced by the FNN controller with nonlinear activation functions, 2) how we transform the computed reachable set with a nonlinear transformation function to verify the safety property of the system. It is worth to emphasize that ***a small over-approximation error and timing efficiency in reachable set computation are two crucial metrics that determine the applicability of reachability analysis methods in safety verification of practical NNCS***. Therefore, safety verification of NNCS requires computationally efficient methods that can compute the exact or tight over-approximate reachable sets of NNCS in a reasonable time. However, computing the exact or tight over-approximate reachable sets of an FNN is difficult and usually time-consuming. In addition, simple utilization of the control set from the controller to compute the reachable set of the plant may produce a very coarse reachable set which is useless in safety verification. Overcome the challenges in problem 1 is a fundamental step to tackle the following important problem.

Problem 2 (Safety-critical initial condition of NNCS). *Given a CPS in problem 1 with the initial states $x(0) \in X_0$, determine the initial condition of the i^{th} state $x^i(0)$ that "may" make the system*

unsafe while keeping the initial conditions of other states unchanged. We call this initial condition is a "safety-critical initial condition" of the system and assume that the initial conditions of all states are independent.

Problem 2 is even harder than problem 1 since it is almost impossible to perform backward analysis of CPS with neural network controllers to determine an unsafe initial condition (backward analysis is generally intractable in this case). In the following, we first present our core reachability algorithm for neural network control systems (NNCS). Then, we discuss handling the nonlinear transformation on the computed reachable set for checking the safety of the system, i.e., Problem 1 as well as searching safety-critical initial condition, i.e., Problem 2.

3 REACHABILITY ANALYSIS OF NEURAL NETWORK CONTROL SYSTEMS

The reachability analysis of a NNCS depicted in Figure 1 is done as follows. First, from the initial set of states X_0 of the plant P, the controller F takes the output set of the plant Y_0 as an input to compute the control set $U = F(Y_0)$. Note that Y_0 is an affine mapping of the initial set X_0 with the output matrix C, i.e., $Y_0 = CX_0$. The control set U is then applied to the plant to compute the set of the next state $X_1 = AX_0 + BU$. This routine is performed iteratively to obtain the reachable set of the plant $X_0, X_1, \cdots X_k$, $0 \leq k \leq k_{max}$. To obtain tight reachable sets of the NNCS, we compute the exact control set U given the output set Y. Also, we compute the exact reachable set of state X_k given its initial set X_{k-1} and the corresponding control set U_{k-1}.

3.1 Generalized star set

Although computing the exact control set of a FNN controller can be done with a polyhedra approach [20], it is computationally inefficient and not scalable. In addition, the polyhedron-based approach produces a conservative reachable set of the plant because it cannot take advantage of the relationship between U_k and X_k, i.e., $U_k = F(Y_k) = F(CX_k)$. To overcome these challenges, we propose a new reachability analysis approach for NNCS using the concept of star set [2], which is very efficient in affine mapping operations, e.g., $Y_k = CX_k$ and more importantly, it preserves the relationship between U_k and X_k which is crucial to obtain an exact reachable set of the plant. The definition of a star set and its essential properties are given in the following, while further details of the application of star sets as a symbolic state space representation for neural network verification are available [18].

Definition 3.1 (Generalized Star Set [2]). A generalized star set (or simply star) Θ is a tuple $\langle c, V, P\rangle$ where $c \in \mathbb{R}^n$ is the center, $V = \{v_1, v_2, \cdots, v_m\}$ is a set of m vectors in $\mathbb{R}^n$ called basis vectors, and $P : \mathbb{R}^m \rightarrow \{\top, \bot\}$ 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:

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

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

Proposition 3.2. *[Affine Mapping of a Star] Given a star set $\Theta = \langle c, V, P\rangle$, an affine mapping of the star Θ 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.*

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

Proposition 3.3 (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.*

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

We can see that, a star set does not change its predicate over affine mapping operations, and it preserves the center and basis vectors in the intersection with a half-space.

3.2 Exact reachability analysis of the neural network controller

The first step in our reachability analysis is to compute the exact control set $U_k = F(CX_k)$. This computation is done *layer-by-layer* in which the output set of the previous layer is the input set of the next layer. Given a star input set $\bar{\Theta} = \langle \bar{c}, \bar{V}, \bar{P}\rangle$, the reachable set of a layer L can be obtained precisely in two steps. First, an affine map Θ of the input set can be derived quickly with the weight matrix W and bias vector b of the layer, i.e., $\Theta = \langle c = W\bar{c} + b, V = W\bar{V}, P \equiv \bar{P}\rangle$. After calculating the affine map of the input set, the reachable set of the layer $\mathcal{R}_L$ is obtained by applying the ReLU activation function on the affine-mapped set, i.e., $\mathcal{R}_L = ReLU(\Theta)$. Similar to [20], this second step is done by executing a sequence of stepReLU operations $\mathcal{R}_L = ReLU_n(ReLU_{n-1}(\cdots ReLU_1(\Theta)))$. The stepReLU operation on the i^{th} neuron, i.e., $ReLU_i(\cdot)$, works as follows. First, the input star set Θ 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 3.3, Θ_1 and Θ_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 $x = [x_1 \cdots x_i\ x_{i+1} \cdots x_n]^T \in \Theta_2$ will lead to the new vector $x' = [x_1\ x_2 \cdots 0\ x_{i+1} \cdots x_n]^T$. This procedure is equivalent to mapping Θ_2 by the mapping matrix $M = [e_1\ e_2 \cdots e_{i-1}\ 0\ e_{i+1} \cdots e_n]$. Also, applying the ReLU activation function on the element x_i of the vector $x \in \Theta_1$ does not change the set since we have $x_i \geq 0$. Consequently, the result of the stepReLU operation on input set Θ 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 Figure 2.

Similar to [20], we can minimize the number of stepReLU operations. If we know that x_i is always larger than zero, then we have $ReLU_i(\Theta) = \Theta$. 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 overall 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. Furthermore, one can see that a star set can be split into two star sets after a stepReLU operation. Therefore, the exact output set

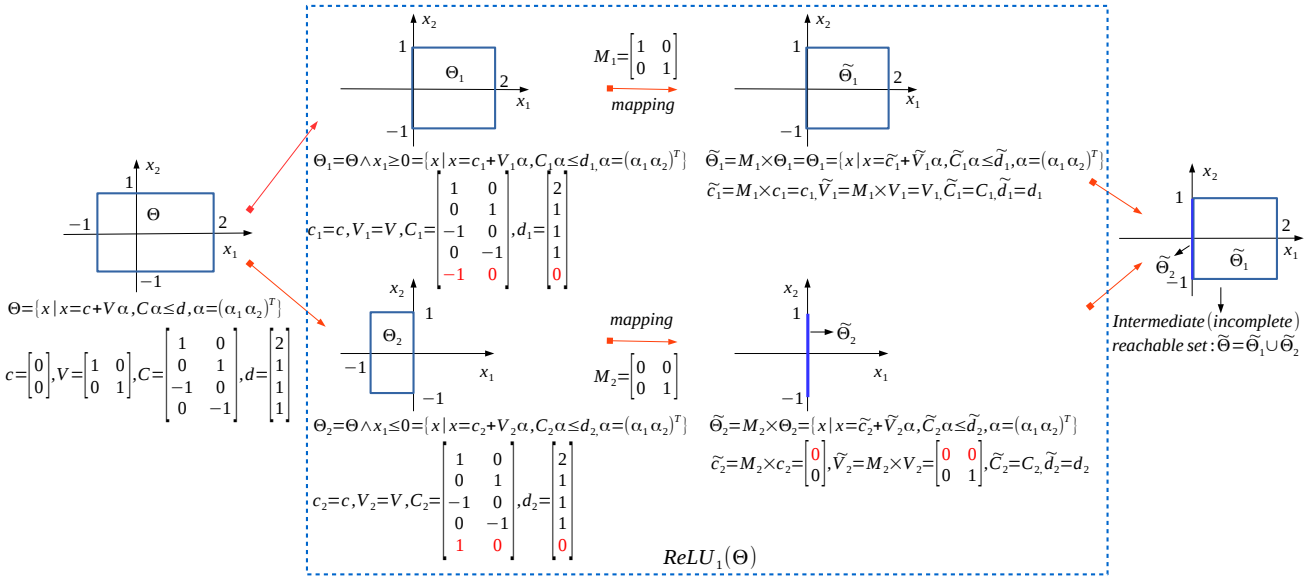

Figure 2: An example of a stepReLU operation on a layer with two neurons.

of a layer is a union of stars which can be handled independently. Based on this observation, the reachability algorithm of an FNN using star set can be designed efficiently to exploit the power of parallel computing as in the polyhedron-based approach [20]. *We emphasize that the exact reachability of an FNN with ReLU activation function can be extended straightforwardly to deal with saturation activation function.*

Although the star set based method is similar to the polyhedron-based approach [20], it is much more efficient and scalable because star set is very fast in affine mapping which is the most expensive step in the polyhedron-based approach, especially for a high dimensional set. More importantly, the computed output set and the input set of the FNN are defined based on the same set of predicate variables, i.e., $\alpha = [\alpha_1, \cdots, \alpha_m]^T$. This property is crucial in eliminating the over-approximation error in computing the reachable set for the plant as addressed in the following.

3.3 Exact reachability analysis of the discrete linear plant

As shown in previous subsection, the exact control set $U_k = F(CX_k)$ is a union of stars, $U_k = \cup_{j=1}^{L} \tilde{\Theta}_j$. Therefore, the exact reachable set of the plant for the next step is also a union of stars, $X_{k+1} = AX_k + BU_k$. Interestingly, the state set $X_k = \langle c, V, P \rangle$ and the control set U_k are defined based on a unique predicate variable vector α and for any star in the control set, its predicate contains all linear constraints of the state set X_k as can be seen in Figure 2. This leads to an important fact that, only a subset of X_k can lead to an individual control set $\tilde{\Theta}_j \in U$ and the predicate of this subset is exactly the predicate of the individual control set. Therefore, the next state set corresponding to the individual control set $\tilde{\Theta}_j = \langle \tilde{c}_j, \tilde{V}_j, \tilde{P}_j \rangle$ is $X_{k+1}^j = \langle Ac + B\tilde{c}_j, AV + B\tilde{V}_j, \tilde{P}_j \rangle$. Consequently, the exact next state set of the plant is $X_{k+1} = \cup_{j=1}^{L} X_{k+1}^j$.

3.4 Reachability algorithm for NNCS

As shown previously, we can compute the exact reachable set of NNCS depicted in Figure 1 by computing the exact control set and the exact state set of the plant. For a single initial state set, after one time step, it may produce many other state sets. Therefore, the number of state sets increases quickly over time which makes the exact analysis time-consuming even using parallel computing. To handle this state sets explosion, we can obtain a single convex hull of the state sets after every step and use it for the next step computation. Computing the convex hull for a set of stars is essentially computing the convex hull of a set of convex polyhedrons which is computationally expensive. To overcome this challenge, we instead compute the interval hull of a set of stars for the next step computation which can be done efficiently by solving a set of linear programming optimization problems. The experimental results show that, by using only the interval hull of the star state sets, we still can obtain a tight over-approximation of the exact reachable set for the NNCS and more importantly, the over-approximation error does not explode over time. The reachability algorithm for a NNCS is summarized in Algorithm 1 in which the user can choose to compute the exact or the over-approximate reachable sets of the NNCS.

Lemma 3.4. *The exact scheme in Algorithm 1 produces the exact reachable sets of the NNCS depicted in Figure 1.*

Proof. The proof can be derived inductively based on the exact computation of the reachable set of the plant and the neural network controller in every step. □

Algorithm 1 Reachability Algorithm for NNCS

% F: neural network controller
% A, B, C: plant's matrices $x_{k+1} = Ax + Bu,\ y_k = Cx_k$
% I: initial set of states of the plant
% k_{max}: number of steps
% *scheme*: reachability analysis scheme, "exact" or "approx"
% R: reachable set
procedure R = Reach($F, A, B, C, I, k_{max}, scheme$)
  $R = cell(1, k_{max} + 1)$
  $R\{1, 1\} = I$
  for $k = 1 : k_{max}$ **do**
    $X_k = R\{1, k\}, M = length(X_k)$
    for $i = 1 : M$ **do**
      $X_k^i = X_k(i) = \langle c, V, P\rangle$
      $U_k = F(CX_k^i) = \cup_{j=1}^{L} \tilde{\Theta}_j = \cup_{j=1}^{L} \langle \tilde{c}_j, \tilde{V}_j, \tilde{P}_j\rangle$
      $X_{k+1} = []$
      for $j = 1 : L$ **do**
        $X_{k+1}^j = \langle Ac + B\tilde{c}_j, AV + B\tilde{V}_j, \tilde{P}_j\rangle$
        $X_{k+1} = [X_{k+1}\ X_{k+1}^j]$
    if *scheme* == *exact* **then** $R\{1, k + 1\} = X_{k+1}$
    else $R\{1, k + 1\} = IntervalHull(X_{k+1})$

Algorithm 2 Safety Verification for NNCS

Input: R, g, U: Reachable set of the NNCS, transformation function, unsafe region
Output: $safe = true$ or $safe = uncertain$
procedure $safe$ = Verify(R, g, U)
  $k_{max} = length(R)$
  for $k = 1 : k_{max}$ **do**
    $X_k = R\{1, k\}$
    $\underline{z}_k = min(g(x_k)),\ \bar{z}_k = max(g(x_k)),\ x_k \in X_k$
    $\tilde{Z}_k = [\underline{z}_k,\ \bar{z}_k]$
    if $\tilde{Z}_k \cap U = \emptyset$ **then** $safe = true$
    else $safe = uncertain$, break

4 VERIFICATION OF NEURAL NETWORK CONTROL SYSTEMS

4.1 Safety verification

Although a safety property of CPS is usually represented as a linear predicate over the system's states x_k, there are many cases where the safety property is defined as a linear predicate over a state variable z_k that is a nonlinear transformation of the system's states, i.e., $z_k = g(x_k)$, where g is a nonlinear transformation function. Let $\mathcal{U}(z_k) \triangleq Hz_k \leq h$ be the unsafe region of a NNCS, the safety verification of the NNCS, i.e., Problem 1, is equivalent to checking $Z_k \cap \mathcal{U}(z_k) = \emptyset?\ \forall 0 \leq k \leq k_{max}$, where $Z_k = \{z_k |\ z_k = g(x_k),\ x_k \in X_k\}$ is the transformed reachable set of the system. Since computing the exact transformed reachable set is computationally expensive and may be even infeasible, we compute an over-approximation of the exact transformed reachable set $\tilde{Z}_k$ and use it for safety verification. The system is safe if $\tilde{Z}_k \cap \mathcal{U}(z_k) = \emptyset,\ \forall 0 \leq k \leq k_{max}$. Particularly, we compute the tightest interval bounding the exact transformed reachable set by solving the following nonlinear optimization problem.

$$\tilde{Z}_k = [\underline{z}_k,\ \bar{z}_k],\ \underline{z}_k = min(g(x_k)),\ \bar{z}_k = max(g(x_k)),\ x_k \in X_k.$$

Safety verification of the NNCS is summarized in Algorithm 2 which solves the above nonlinear optimization problem to obtain the tightest interval of the transformed reachable set and uses it to verify the safety of the system at each time step.

4.2 Characterization of safe initial condition

Safety verification of a NNCS can reason about the safety of the system w.r.t a specific initial condition. In some cases, we are interested in the upper bound of a particular state $x^i(0)$ in the initial condition where the safety of the system is still guaranteed. For example, if a car detects an obstacle and applies the brake to stop, it is important to know what is the maximum velocity of the vehicle such that the braking action can guarantee the safety of the car. To search for that maximum velocity, we start from the initial condition that the system is safe, then we increase the upper bound of the speed by some δ, i.e., $x^i(0) = x^i(0) + \delta$, and check the safety of the system with the new initial condition. We continue to increase the upper bound until the safety is uncertain. We can obtain the maximum allowable velocity with the error of $[-\delta, \delta]$.

5 CASE STUDY

Our approach is implemented in *NNV* [20], a Matlab toolbox for safety verification of DNNs. The proposed approach is evaluated on a practical automatic emergency braking system (AEBS) for an autonomous car. The architecture of the AEBS is described in Figure 3 in which the car is equipped with a perception component to detect automatically the obstacle on the road and a reinforcement learning (RL) based controller to control the brake of the car. All results presented in this paper and their corresponding scripts are available online at https://www.dropbox.com/s/912184bij8yqz37/EMSOFT2019.zip?dl=0, or in the GitHub repository for our NNV tool at https://github.com/verivital/nnv, as well as in a CodeOcean capsule that can be executed without a Matlab license [19].

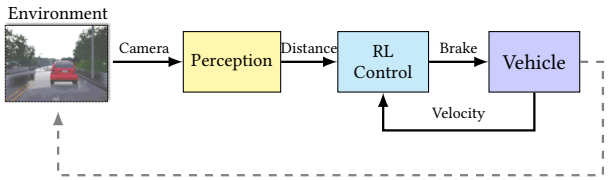

Figure 3: Emergency Braking System Architecture

5.1 Scenario of Interest

In our system, we consider the scenario that the host car automatically detects another static vehicle and applies a brake to decelerate and stop to avoid the potential collision as shown in Figure 4.

The host car starts from rest and accelerates to a random initial velocity v_0, which introduces the uncertainty to the system. Then, the car keeps this velocity v_0 till an obstacle is detected at distance d_0 from the perception module and switches to the reinforcement learning braking controller. The goal of the controller is to stop the car to avoid the collision and also not too far from the obstacle,

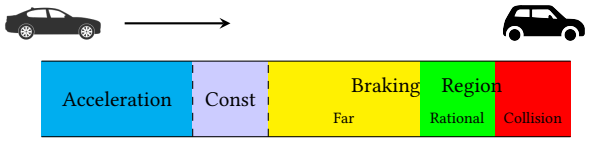

Figure 4: Illustration of Emergency Braking System

which means the car should stop within the safety and rational region.

5.2 Safety Specification

The safety property of the AEBS is defined based on the concept of time to collision (TTC)[12, 13]. TTC measures the time it wold take to collide if the vehicle continues traveling based on the current acceleration of $a_k = u_k$ and velocity v_k. Smaller TTC means a higher collision risk. The safety specification of the AEBS can be written by

$$(\mathrm{TTC}_k(d_k, v_k, a_k) > \tau(v_k))\ \mathcal{U}\ (k = k_{max})$$

where $\tau(v_k)$ is the time to stop when applying the full brake for velocity v_k, shown in Figure 5, d_k is the current distance from the car to the obstacle, k_{max} is the maximum number of steps we want to verify the safety of the system, and $\mathcal{U}$ is the until operator. Generally, the safety specification means that the car is safe if it still has enough time for a full braking action, i.e., full braking action can successfully stop the car before a collision occurs.

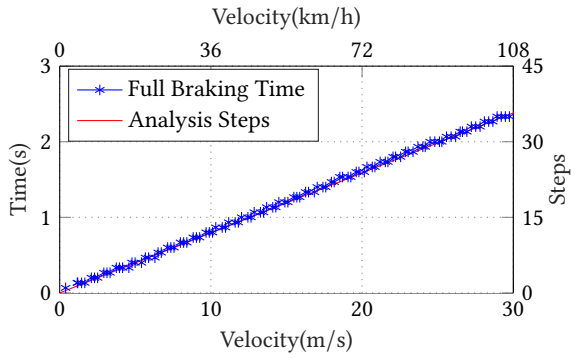

Figure 5: Full Braking Time and Required Reachability Analysis Steps vs. velocity

Because of the discontinuity caused by the denominator when velocity or acceleration equals zero, it is more efficient to evaluate the collision risk using the inverse TTC introduced in [3]. The inverse TTC is proportional to the collision risk: the higher it is, the higher the collision risk is. The safety specification using the inverse TTC is given below,

$$(\mathrm{TTC}_k^{-1}(d_k, v_k, a_k) < \tau^{-1}(v_k))\ \mathcal{U}\ (k = k_{max}),$$

where, the inverse TTC is defined by:

$$\mathrm{TTC}_k^{-1}(d_k, v_k, a_k) = \begin{cases} \frac{v_k}{d_k} & \text{for } a_k = 0 \\ \frac{-a_k}{v_k - \sqrt{v_k^2 + 2a_k d_k}} & \text{for } v_k^2 + 2a_k d_k \geq 0 \wedge a_k \neq 0 \\ 0 & \text{for } v_k^2 + 2a_k d_k < 0 \wedge a_k \neq 0. \end{cases}$$

5.3 RL-based Controller

We train the RL-based controller for the host car using Deep Deterministic Policy Gradient (DDPG)[14], which is a popular reinforcement learning method that combines the value-based and the policy-based method. There are two parts in this approach including actor and critic. Critic uses the off-policy data to learn the Q-function, which evaluates how good the action a taken is in given state s. The actor can learn the continuous action policy by using the Q-function. In practice, it is difficult to obtain the exact Q-function and policy function. Therefore, two neural networks are introduced to solve this problem, which is critic network $Q(s, a|\theta^Q)$ and actor network $\mu(s|\theta^\mu)$ with weights θ^Q and θ^μ. Coming back to our braking system, the reinforcement learning controller consumes the state s, consisting of distance to the leader vehicle d and host car's velocity v, and computes the action – brake T.

For a reinforcement learning system, the reward function should be appropriately designed to achieve the goal. In our case, the task is to stop the car in a safe and rational region. Thus, we define the reward function as

$$r = -\alpha \times I \times \mathbf{1}(collision) - [(d_t - B) \times \beta + \lambda] \times \mathbf{1}(v_t = 0 \wedge d_t > B)$$

where d_t and v_t indicates the distance to the leader car and velocity at time step t, α, β and λ are coefficients greater than zero, $\mathbf{1}(\cdot)$ returns a value of 1 if the statement inside is true and 0 otherwise.

The term of the reward function, $-\alpha \times I \times \mathbf{1}(collision)$ penalizes a collision event based on the collision impulse I. The other term $-[(d_t - B) \times \beta + \lambda] \times \mathbf{1}(v_t = 0 \wedge d_t > B)$ penalizes a too early stop based on B, the final distance to the boundary line between rational and far region. During the braking process (before the car comes to a stop), there is no penalty or reward. Intuitively, this reward function will guide the car to stop within the rational region.

We use CARLA [6] to generate the scenario and to train the reinforcement learning controller. The time step used in the simulation is $\Delta t = 1/15$ s. In the simulations, the vehicle firstly accelerates to a velocity of v_0, and keeps the speed untill it detects an obstacle at a distance d_0. The d_0 and v_0 are the initial states of the braking system. To simulate a more realistic scenario, we introduce some uncertainty to the initial states of the system. The initial velocity of the vehicle is uniformly sampled between 90 km/h and 100 km/h, and the initial distance depends on the range of the perception module, which is approximately 100 m. After initial state, the car switches to the reinforcement learning controller which consists of two neural networks trained with DDPG algorithm with the hyper-parameters in Table 1 is presented below:

- Actor NN architecture[1]:

  2(State) × 50(ReLU) × 30(ReLU) × 1(SatReLU, Action)

[1] SatReLU is the ReLU function with max value 1.

Table 1: Hyper-parameters for DDPG algorithm.

	Actor	Critic
Optimizer	Adam	Adam
Learning rate	10^{-4}	10^{-3}
Target update rate	0.9	0.9
Reply buffer size	10^5	
Reply batch size	32	
Discount factor	0.99	
Reward function	$\alpha = 0.01, B = 5, \beta = 1.6, \lambda = 20$	

- Critic NN architecture[2]:

$$\left.\begin{matrix} 2(\text{State}) \times 50(\text{ReLU}) \times 30 \\ 1(\text{Action}) \times 30 \end{matrix}\right\} \times 30(\text{ReLU}) \times 1(\text{Q Value})$$

We trained the reinforcement learning for 1000 episodes, and the neural network converges, showing an attractive performance. Also, one of the experiment trajectories is plotted in Figure 7. At the beginning of involving the reinforcement learning controller, the distance is 97.3 m, and the velocity is 91.98 km/h (= 25.55 m/s). After 128 steps, about 8.53 s, the ego vehicle stops at about 1.88 m far from the obstacle vehicle.

5.4 System Identification and Validation

We transfer the braking system from CARLA to MATLAB & Simulink to perform reachability analysis and safety verification for the system. The diagram of the Simulink model of the AEBS is shown in Figure 6. For simulation and verification, only the actor is needed.

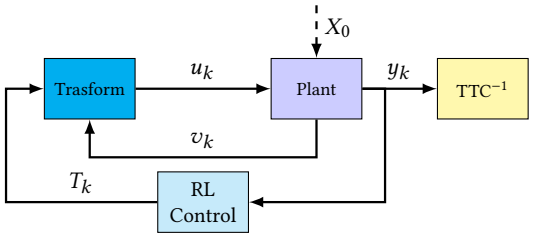

Figure 6: Emergency Braking System Simulink Diagram

The plant of the braking system is described by following discrete state-space equation

$$\begin{cases} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k \end{cases}$$

where $x_k = [d_k \; v_k]^T$ is the state vector including the distance d_k and the velocity v_k of the car at step k, u_k is the input, which is the acceleration applied to the plant, y_k is the output, and A, B, C, D are the coefficient matrices given below,

$$A = \begin{bmatrix} 1 & -\Delta t \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 0 \\ \Delta t \end{bmatrix}, C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, D = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

where $\Delta t = 1/15$ is simulation time step.

[2]The empty activation function means no activation is applied.

It is important to emphasize that the input of the plant u_k does not match with the output of the reinforcement learning controller T_k. The u_k is the acceleration applied to the car, but the T_k is the braking force. Thus, a neural network transformation with 80 neurons is trained to bridge this gap between u_k and T_k.

To validate the Simulink model of AEBS, we run experiments in Simulink and CARLA with the same initial states and compare them as shown in Figure 7. From the plot, we can see that the Simulink model captures very well the behaviors of the (actual) AEBS in CARLA.

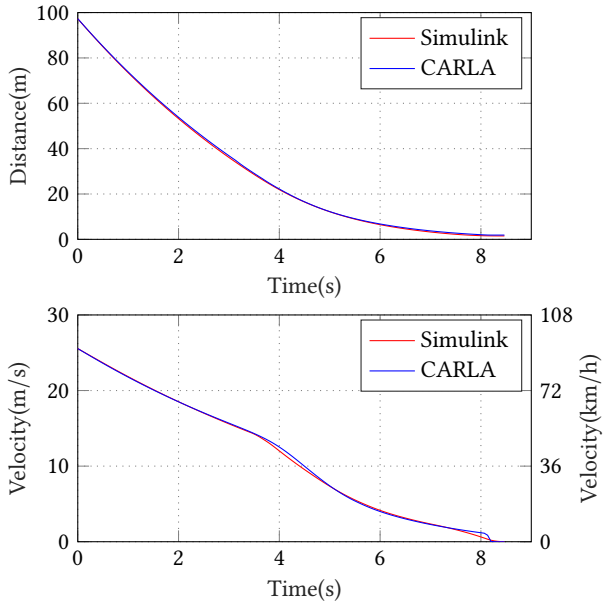

Figure 7: Validation of the Simulink model of AEBS. The Simulink model captures well the behaviors of the actual AEBS in CARLA.

5.5 Safety Verification of the AEBS

Physical constraints for safety verification. To verify the safety of the AEBS, we need to take into account some essential physical constraints of the system. First, the AEBS system uses a perception component to detect the obstacle. The operating range of the perception component is from 0 to 100 meters. Therefore, we are going to verify the safety of the AEBS for the distance (between the car to the obstacle) from 10 to 100 meters (we assume that the car is at least 10 meters far away from the obstacle). Secondly, we limit the maximum allowable velocity of the car is 35 m/s, i.e., ≈ 80 miles per hour which is a usual upper limit of the speed on highways.

Thirdly, we need to know what is a reasonable constraint between initial conditions of the car's velocity and its distance to the obstacle such that if we apply a full braking action, the car is safe. This information is important that we should know before verifying the safety of AEBS because there are cases when even if we apply the full braking action, the collision still occurs. For example, the car is too close to the obstacle and is travelling at a high speed. From Figure 5, we approximate an analytical formula for the full

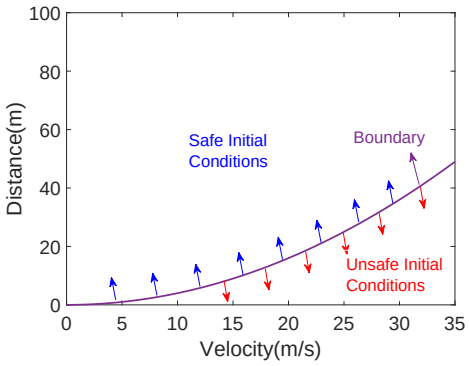

Figure 8: Safe initial conditions for full braking action.

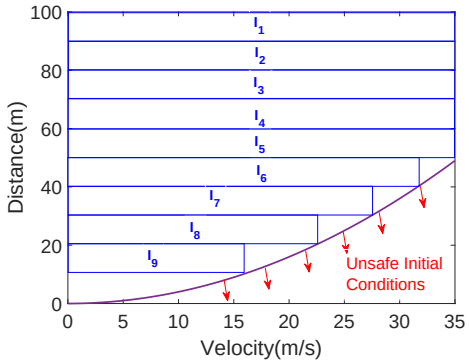

Figure 9: Set of initial conditions that needs to be verified for the AEBS with the RL controller.

braking time that is $\tau(v) \approx v/12.5$. When the full braking action occurs, the car goes a distance $d_s = 0.5a\tau^2 + v\tau$ before stopping, where a is the average acceleration of the car which is equal to $a = \Delta v/\Delta t = (0-v)/\tau$. Therefore, the average travel distance of the car after applying a full brake is: $d_s = 0.5v\tau = 0.5v^2/12.5 = v^2/25$. To guarantee the safety, the initial distance of the car d_0 should be larger than this travel distance, i.e., $d_0 > d_s$. Combining the above limitation on the distance $d_{max} = 100$ m and the maximum allowable velocity $v_{max} = 35$ (m/s), a ***safe initial condition region for full braking action*** is depicted in Figure 8. By partitioning the safe initial condition region of the full braking action, we can derive ***the reasonable initial conditions that need to be verified for the safety of the AEBS with the RL controller*** as shown in Figure 9. This is because, under the safety aspect, the RL controller cannot overcome the full-braking action.

Finally, we need to find out what the minimum number of steps that we should at least give a guarantee about the safety of the system is. We should prove the safety of the system at least $\tau(v)$ seconds in the future where $\tau(v)$ is the full braking time w.r.t the velocity v. Therefore, ***the minimum number of steps that needs to prove the safety*** is: $min(k_{max}) = \tau(v)/\Delta t$. For example, if $v = 25$ m/s, we should at least prove the safety of the system until $k = k_{max} = 2/(1/15) = 30$ time steps.

Challenge and drawback of polyhedron approach [20, 24]. The main challenge in the safety verification of AEBS is how to compute a tight reachable set of the AEBS model depicted in Figure 6. One can see that the control set $U = \{u_t\}$ applied to the plant is derived from the transformation component that takes the output set $T = \{T_t\}$ from the RL controller and the velocity $V = \{v_t\}$ as the input set. Therefore, to compute the control set U, we need to compute the output set T of the RL controller and then combine with the velocity set $V = \{v_t\}$ of the plant to form the input set for the transformation neural network. The problem is how to efficiently combine these sets to form the exact input set for the transformation neural network. This problem is unsolvable if we use the polyhedron-based approach [20, 24] since the relationship between the output set T of the RL controller and the velocity set V of the plant cannot be preserved in the computation. This leads to a coarse combination which returns a coarse input set for the transformation neural network. Consequently, the over-approximation error is exploded quickly after only 3 time steps as shown in Figure 10. From the figure, we can see that the obtained reachable set is too conservative and cannot be used for safety verification.

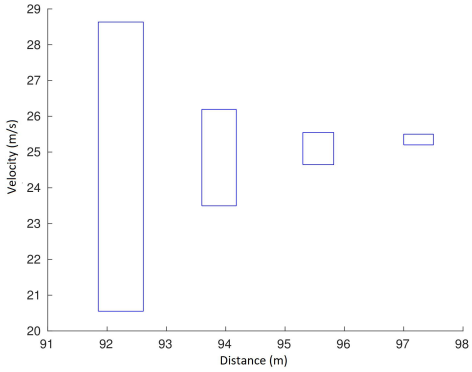

Figure 10: Reachable set of the AEBS computed using polyhedron-based approach [20, 24] with the initial conditions $d_0 \in [97, 97.5]$, $v_0 \in [25.2, 25.5]$. The Over-approximation error is exploded quickly after only 3 time steps.

Overcome the challenge with our star-based approach. Our star-based approach is an efficient technique to overcome the main challenge discussed above. We perform both the exact and over-approximate reachability analysis for the AEBS with the initial conditions $d_0 \in [97, 97.5]$, $v_0 \in [25.2, 25.5]$ for 50 time steps. The results are presented in Figure 11 and 12 with noticing that we use boxes to represent the reachable sets. Our star-based approach eliminates (in the exact method) or reduces significantly (in the over-approximation method) the over-approximation error caused by the polyhedron-based approach. The reachable sets computed from the star-based approach are tight and useful for safety verification of the AEBS. The reachability analysis times of two proposed methods are presented in Table 2. The Table shows that the over-approximation method is faster than the exact method while still produces tight reachable sets for the system. From the figures, one can see that the reachable sets computed by the two methods are almost the same. The time improvement of using the over-approximation method increases as the number of time steps grows. The reason that makes the over-approximation method

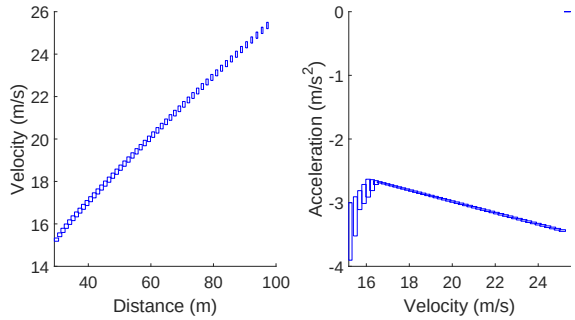

Figure 11: Reachable set of the AEBS computed using the exact star-based method with the initial conditions $d_0 \in [97, 97.5]$, $v_0 \in [25.2, 25.5]$. The Over-approximation error is eliminated.

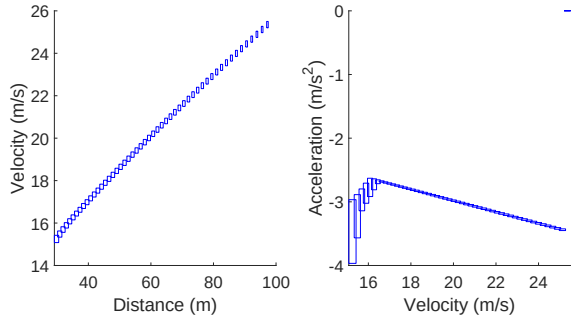

Figure 12: Reachable set of the AEBS computed using the over-approximate star-based method with the initial conditions $d_0 \in [97, 97.5]$, $v_0 \in [25.2, 25.5]$. The Over-approximation error is reduced significantly.

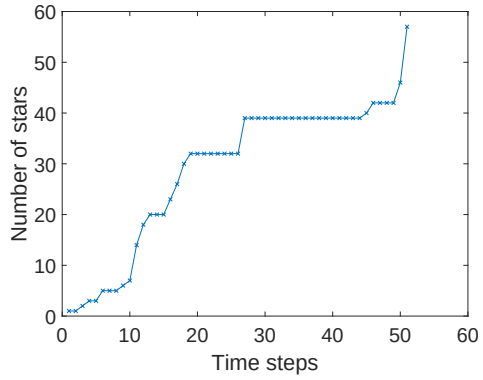

Figure 13: Number of stars in the reachable sets of the AEBS grows over time with the exact star-based method.

faster is, it produces only a single reachable set at every time step while in the exact method, the number of reachable sets may grow over time as depicted in Figure 13.

Checking safety using the computed reachable sets. To verify the safety of the AEBS, we consider the worst case, i.e., we want to verify if the following constraint is satisfied in a bounded time, $max(TTC^{-1}(d, a, v)) < \tau^{-1}(max(v))$. To do that, we estimate the

Method	N = 5	N = 10	N = 20	N = 30	N = 40	N = 50
Exact star	12.47	32.24	162.95	400.13	532.1	831
Over-approximate star	10.07	21.09	42.86	63.98	83.3	104.44
Time improvement	**1.24x**	**1.53x**	**3.8x**	**6.25x**	**6.39x**	**7.96x**

Table 2: Reachability analysis times (measured in seconds) of the exact and over-approximate star methods in which N is the number of time steps. The experiment is done on a computer with following configurations: Intel Core i7-8859H CPU @ 2.6GHz × 4 Processor, 32 GiB Memory, Window 10 Pro OS. We use 6 cores for the computation.

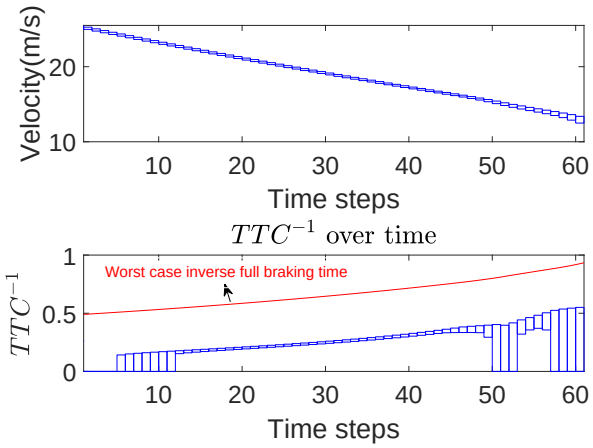

Figure 14: The inverse TTC over time is smaller than the worst case inverse full braking time $\tau^{-1}(max(v))$. The AEBS is safe (for 60 time steps) with the initial conditions $d_0 \in [97, 97.5]$, $v_0 \in [25.2, 25.5]$.

ranges of TTC^{-1} in 60 time steps (two times larger than the minimum requirement $k_{max} = 30$) using the ranges of the distance, velocity, and acceleration of the car from the computed reachable sets and check if it satisfies the requirement or not. The result is illustrated in Figure 14 which shows that the inverse TTC is smaller than the worst case inverse full braking time $\tau^{-1}(max(v))$. Therefore, the AEBS is safe for 60 time steps in the future.

5.6 Safe Initial Conditions of the AEBS

From the physical constraints of the car, we have derived the set of initial conditions that need to be verified for the AEBS with RL controller as depicted in Figure 9. It is important to determine in these initial conditions, which regions are safe for the AEBS with RL controller and which ones are risks. We perform our safety verification methods on each partition $I_i, i = 1, 2, \cdots, 9$ of the initial conditions to find the safe regions. We perform the search as follows. We partition the distance range $[10, 100]$ into 9 smaller ranges with the same width of 10, i.e., $d^i = [10i, 10(i+1)], 1 \le i \le 9$. For the i^{th} individual distance range, we search for the maximum velocity v^i_{max} such that the RL controller can guarantee the safety of the system in $k_{max} = 50$ time steps for the initial condition of $[d^i, v^i_{max}]$. The results of v_{max} are presented in Table 3. From the information of v_{max}, we visualize the safe region of the initial conditions for the AEBS as depicted in Figure 15.

$d(m)$	[0, 10]	[10, 20]	[20, 30]	[30, 40]	[40, 50]
$v_{max}(m/s)$	—	4	7	8	10
$d(m)$	[50, 60]	[60, 70]	[70, 80]	[80, 90]	[90, 100]
$v_{max}(m/s)$	15	19	21	24	26

Table 3: Safe initial conditions for the AEBS with RL controller in which d is the distance range and v_{max} is the maximum allowable velocity such that the system is still safe.

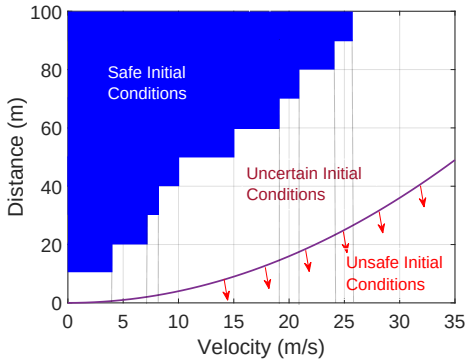

Figure 15: Safe region of the initial conditions of the AEBS with the RL controller.

6 RELATED WORK

Verification, testing, and falsification for cyber-physical systems (CPS) containing learning enable components have become an emerging research topic recently. Toward verification for CPS with learning enable components, several methods have been proposed recently to verify the safety of feedback neural network control systems [10, 16, 17, 22, 24]. The early polyhedron-based approach has been extended for safety verification of neural network controlled systems in [24] in which the plant is assumed to be *linear and discrete*. Recently, Verisig [10] proposes an approach that transforms a neural network controller with *sigmoid* activation function to an equivalent nonlinear hybrid system which is then combined with the plant dynamics before utilizing Flow* [5] to verify its safety properties. Another approach using satisfiability modulo convex (SMC) solver for formal verification of NNCS has been proposed in [17]. In this context, the closed-loop control system was encoded as monotone SMC formulas which were formally verified by SMC decision procedures. Impressively, this method is sound and complete with noticing that the plant dynamics is *linear and discrete*. Last but not least, a new abstraction method [16] has been proposed for NNCS verification in which an "local" Taylor model over-approximation of neural network controller was obtained and integrated into Flow*, a flowpipe constructor using Taylor model, to compute a tight over-approximation reachable set of NNCS. This method is impressively fast and scalable for NNCS verification and more importantly, it can reduce over-approximation errors significantly in reachable set computation process and can deal with relative large input sets. A summary of recent verification methods is given in Table 4. In the testing context, a simulation-based test generation framework for autonomous vehicles with machine learning components has been proposed in [21] to enhance the reliability of autonomous driving systems. In the falsification context, a compositional falsification framework for CPS with machine learning components has been proposed in [7]. In this framework, a temporal logic falsifier cooperates efficiently with a machine learning analyzer to find falsifying executions of the system. The effectiveness of the proposed framework was shown via Automatic Emergency Braking System (AEBS).

Approaches	Plant Dynamics	Discrete/Continuous	Activation Function	Size of Controller
Polyhedron-based [24]	Linear	Discrete	ReLU	≤ 100 neurons
Verisig [10]	Linear, Nonlinear	Discrete, Continuous	Sigmoid, Tanh	≤ 50 neurons
SMC-based [17]	Linear	Discrete	ReLU	≤ 200 neurons
Sherlock [16]	Linear, Nonlinear	Discrete, Continuous	ReLU	≤ 500 neurons
Star-based	Linear	Discrete	ReLU	≤ 200 neurons

Table 4: Approaches for neural network control systems verification in which the sizes of controllers are collected from the related experimental results.

In this paper, we focus on safety verification of linear, discrete NNCS with ReLU activation function. We mainly focus on the exact and over-approximate analysis for such a system which aims at eliminating or significantly reducing the over-approximation error in the reachable set computation. We emphasize that our approach can be combined with the zonotope-based reachability algorithms [1, 8] to deal with NNCS with both continuous and discrete nonlinear plant. However, this problem is not the main focus of the paper.

7 CONCLUSION

We have proposed two efficient, exact and over-approximate reachability schemes and an optimization-based approach for safety verification of CPS with RL controller where the safety specification is defined based on a nonlinear transformation of the system states. From thorough experiments on the practical AEBS, we have shown that our method is computationally cheaper and less conservative than the existing polyhedron approach. More important, it is applicable for real-word applications. Our future work is extending the proposed methods for nonlinear NNCS with other types of nonlinear activation functions such as Tanh or Sigmoid.

ACKNOWLEDGMENTS

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 Government purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFOSR or DARPA.

REFERENCES

[1] Matthias Althoff. 2015. An introduction to CORA 2015. In *Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems*.

[2] Stanley Bak and Parasara Sridhar Duggirala. 2017. Simulation-equivalent reachability of large linear systems with inputs. In *International Conference on Computer Aided Verification*. Springer, 401–420.

[3] Valentina E Balas and Marius M Balas. 2006. Driver assisting by inverse time to collision. In *2006 World Automation Congress*. IEEE, 1–6.

[4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).

[5] Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. 2013. Flow*: An analyzer for non-linear hybrid systems. In *International Conference on Computer*

Aided Verification. Springer, 258–263.

[6] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An open urban driving simulator. *arXiv preprint arXiv:1711.03938* (2017).

[7] Tommaso Dreossi, Alexandre Donzé, and Sanjit A Seshia. 2017. Compositional falsification of cyber-physical systems with machine learning components. In *NASA Formal Methods Symposium*. Springer, 357–372.

[8] Antoine Girard. 2005. Reachability of uncertain linear systems using zonotopes. In *Hybrid Systems: Computation and Control*. Springer, 291–305.

[9] John Hertz, Anders Krogh, and Richard G Palmer. 1991. *Introduction to the theory of neural computation*. Addison-Wesley/Addison Wesley Longman.

[10] Radoslav Ivanov, James Weimer, Rajeev Alur, George J Pappas, and Insup Lee. 2019. Verisig: verifying safety properties of hybrid systems with neural network controllers. In *Hybrid Systems: Computation and Control (HSCC)*.

[11] Kyle D Julian, Mykel J Kochenderfer, and Michael P Owen. 2018. Deep Neural Network Compression for Aircraft Collision Avoidance Systems. *arXiv preprint arXiv:1810.04240* (2018).

[12] Kristofer D Kusano and Hampton Gabler. 2011. Method for estimating time to collision at braking in real-world, lead vehicle stopped rear-end crashes for use in pre-crash system design. *SAE International Journal of Passenger Cars-Mechanical Systems* 4, 2011-01-0576 (2011), 435–443.

[13] David N Lee. 1976. A theory of visual control of braking based on information about time-to-collision. *Perception* 5, 4 (1976), 437–459.

[14] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).

[15] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2574–2582.

[16] Sriram Sankaranarayanan Souradeep Dutta, Xin Chen. 2019. Reachability Analysis for Neural Feedback Systems using Regressive Polynomial Rule Inference. In *Hybrid Systems: Computation and Control (HSCC)*.

[17] Xiaowu Sun, Haitham Khedr, and Yasser Shoukry. 2019. Formal Verification of Neural Network Controlled Autonomous Systems. In *Hybrid Systems: Computation and Control (HSCC)*.

[18] Hoang-Dung Tran, Diego Manzanas Lopez, Patrick Musau, Xiaodong Yang, Luan Viet Nguyen, Weiming Xiang, and Taylor T Johnson. 2019. PHAVer: Algorithmic verification of hybrid systems past HyTech. In *23rd International Symposium on Formal Methods (FM)*. Springer.

[19] Hoang-Dung Tran, Diego Manzanas Lopez, Patrick Musau, Xiaodong Yang, Weiming Xiang, and Taylor T Johnson. 2019. nnv - Neural Network Verification Software Tool - Output Reachable Set Estimation and Verification for Multilayer Neural Networks. https://www.codeocean.com/. (June 2019). https://doi.org/10.24433/CO.1314285.v1

[20] Hoang-Dung Tran, Patrick Musau, Diego Manzanas Lopez, Xiaodong Yang, Luan Viet Nguyen, Weiming Xiang, and Taylor T. Johnson. 2019. Parallelizable Reachability Analysis Algorithms for Feed-Forward Neural Networks. In *7th International Conference on Formal Methods in Software Engineering (FormaliSE2019), Montreal, Canada*.

[21] Cumhur Erkan Tuncali, Georgios Fainekos, Hisahiro Ito, and James Kapinski. 2018. Simulation-based Adversarial Test Generation for Autonomous Vehicles with Machine Learning Components. *arXiv preprint arXiv:1804.06760* (2018).

[22] Weiming Xiang, Diego Manzanas Lopez, Patrick Musau, and Taylor T Johnson. 2019. Reachable Set Estimation and Verification for Neural Network Models of Nonlinear Dynamic Systems. In *Safe, Autonomous and Intelligent Vehicles*. Springer, 123–144.

[23] Weiming Xiang, Hoang-Dung Tran, and Taylor T Johnson. 2017. Reachable set computation and safety verification for neural networks with ReLU activations. *arXiv preprint arXiv:1712.08163* (2017).

[24] Weiming Xiang, Hoang-Dung Tran, Joel A Rosenfeld, and Taylor T Johnson. 2018. Reachable Set Estimation and Safety Verification for Piecewise Linear Systems with Neural Network Controllers. *arXiv preprint arXiv:1802.06981* (2018).