

One-Shot Reachability Analysis of Neural Network Dynamical Systems

Shaoru Chen, Victor M. Preciado, Mahyar Fazlyab

Abstract—The arising application of neural networks (NN) in robotic systems has driven the development of safety verification methods for neural network dynamical systems (NNDS). Recursive techniques for reachability analysis of dynamical systems in closed-loop with a NN controller, planner or perception can over-approximate the reachable sets of the NNDS by bounding the outputs of the NN and propagating these NN output bounds forward. However, this *recursive* reachability analysis may suffer from compounding errors, rapidly becoming overly conservative over a longer horizon. In this work, we prove that an alternative *one-shot* reachability analysis framework which directly verifies the unrolled NNDS can significantly mitigate the compounding errors, enabling the use of the rolling horizon as a design parameter for verification purposes. We characterize the performance gap between the recursive and one-shot frameworks for NNDS with general computational graphs. The applicability of one-shot analysis is demonstrated through numerical examples on a cart-pole system.

I. INTRODUCTION

Robotic systems embedding learning-enabled modules such as neural network (NN) controllers, planners, or perception have achieved state-of-the-art performances in various complex tasks and are becoming increasingly popular. However, such neural network dynamical systems (NNDS) lack formal safety guarantees and are prone to failures due to the fragility of NNs to adversarial attacks or random input perturbation [1], [2]. Verifying the safety of NNDS before deployment is therefore required in safety critical applications, but handling the large scale and high complexity of NNs in verification is challenging.

Neural network verification consists in certifying that the output of a NN satisfies certain properties given a bounded set of inputs. A rich body of works has focused on developing specialized solvers [3]–[11] for NN verification. Although NN verification methods only consider NNs in isolation, they can be conveniently combined with existing reachability analysis tools to certify properties of NNDS [12]–[14] over a finite horizon. For a discrete-time NNDS, NN verification methods can readily compute an over-approximation of the one-step reachable set given a bounded input set. Then, applying such one-step over-approximation method recursively for $t = 0, 1, \dots, T$ leads to a bounding tube of the NNDS trajectories (blue boxes in Fig. 1).

We denote the above methodology as the *recursive reachability analysis* framework [15]–[17] and compare it with an alternative framework that computes over-approximations

Shaoru Chen and Victor M. Preciado are with the Department of Electrical and Systems Engineering, University of Pennsylvania. {srchen, preciado}@seas.upenn.edu. Mahyar Fazlyab is with the Mathematical Institute for Data Science, Johns Hopkins University mahyarfazlyab@jhu.edu.




Fig. 1: Reachable set over-approximations of a NN dynamical system can be obtained by applying NN verification methods recursively (blue arrows) or in one-shot (red arrows). In the one-shot analysis, unrolled NN dynamics over multiple time steps is considered. In this work, we investigate in what cases the one-shot analysis generates tighter bounds than the recursive counterpart.

of the reachable sets in one shot [18], [19]. In *one-shot reachability analysis*, the reachable set of NNDS at time t is bounded by directly applying NN verification methods on the unrolled NN dynamics for t steps (red boxes in Fig. 1). While one may observe that the one-shot analysis generates tighter bounds than the recursive counterpart due to the use of iterated dynamics [18], this property does not hold for general NN verification tools. Correspondingly, our contributions in this paper are:

- We provide a counter-example where one-shot analysis results in worse bounds than the recursive framework, highlighting that the one-shot analysis does not always lead to tighter bounds.
- We formally prove conditions under which the one-shot framework provides tighter bounds compared with the recursive framework for a general class of NN verification methods.
- Our analysis applies to NNs with general architectures and allows us to consider disturbances in the computation of the reachable set over-approximations. Numerical examples are provided to demonstrate the applicability of the one-shot framework.

A. Related works

a) NN verification: NN verification methods analyze NNs in isolation and can certify the worst-case performance of NNs under bounded input perturbations. A rich set of tools with distinct tightness/complexity trade-offs are developed, including those based on mixed-integer programming (MIP) [3], semidefinite programming (SDP) [10], [11], linear programming (LP) [5], and linear bounds propagation [6], [20], [21]. Particularly, [21] considers NNs with general architectures. Tightness comparison of different convex relaxation-based NN verifiers is discussed in [22], [23] from the dual perspective.

b) *Reachability analysis of NNDS*: Verifying a closed-loop dynamics with a NN controller often involves first bounding the set of control inputs (e.g., by polytopes [24], linear bounds [25], star sets [26], polynomial zonotopes [14], hybrid automaton [27], Taylor models [28]) in each control cycle and then estimating the reachable set of the states. Several works [13], [17], [29] process the overall closed-loop dynamics directly. The application of the recursive analysis framework to find an over-approximation of the reachable set within a finite horizon can be found in [17], [25], [26], [30]. The effectiveness of the one-shot analysis in reducing compounding errors is demonstrated and analyzed in detail in OVERT [18]¹, a safety verification method for nonlinear systems with NN controllers. In OVERT, the nonlinear dynamics is over-approximated by piecewise linear relations such that the safety of the closed-loop system can be verified through mixed-integer programming (MIP). Instead of focusing on one particular NN verification method, in this work we aim to provide guarantees on the performance of the one-shot framework for a range of NN verification methods including MIP [3], linear programming (LP) [5], and linear bounds propagation-based ones [6], [21]. Our analysis allows straightforward integration of these NN verification methods into the one-shot framework to further boost their performances in safety verification.

II. PRELIMINARIES AND PROBLEM FORMULATION

In this section, we formalize the reachability analysis problem for NNDS and two frameworks to approach it: the recursive and the one-shot analysis. Although, intuitively, the one-shot method should outperform the recursive method in tightness at the cost of computational overhead, we provide a motivating example that falsifies this claim which calls for a formal comparison between these two frameworks.

A. Neural network dynamical systems

We denote a discrete-time NNDS with disturbances as

$$x_{t+1} = f(x_t, w_t) \quad (1)$$

where $x_t \in \mathbb{R}^{n_x}$ is the state, $w_t \in \mathbb{R}^{n_w}$ denotes the disturbance at time t , and f is a neural network with an arbitrary architecture taking (x_t, w_t) as its input. For example, Eq. (1) can represent the closed-loop dynamics of a system consisting of several NN modules:

$$\begin{cases} x_{t+1} = f_{NN}(x_t, u_t) + w_t^x \\ y_t = p_{NN}(x_t) + w_t^y \\ u_t = \pi_{NN}(y_t) \end{cases} \quad (2)$$

where f_{NN}, p_{NN}, π_{NN} denote, respectively, the NN dynamics, measurement function and control policy. The process and measurement disturbances w_t^x and w_t^y can be considered as components of a compound disturbance $w_t = [w_t^x \top \ w_t^y \top]^\top$. Then, $f(x_t, w_t)$ denotes a NN with the connection diagram shown in Fig. 2.




Fig. 2: Illustration of a neural network dynamical system with several interconnected NNs.

B. Finite-step reachability analysis

For the NNDS (1), define the one-step forward reachable set of $f(\cdot)$ from a given set $\mathcal{X} \times \mathcal{W} \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$ as

$$f(\mathcal{X} \times \mathcal{W}) := \{y \in \mathbb{R}^{n_x} \mid y = f(x, w), (x, w) \in \mathcal{X} \times \mathcal{W}\}.$$

For a given set of initial conditions $\mathcal{X}_0 \subset \mathbb{R}^{n_x}$, the forward reachable set $\mathcal{R}_t(\mathcal{X}_0)$ of the NNDS (1) at time t is defined by the recursion

$$\mathcal{R}_{t+1}(\mathcal{X}_0) = f(\mathcal{R}_t(\mathcal{X}_0) \times \mathcal{W}), \mathcal{R}_0(\mathcal{X}_0) = \mathcal{X}_0. \quad (3)$$

Due to the complex structure of the function $f(\cdot)$, exactly identifying $\mathcal{R}_t(\mathcal{X}_0)$ is computationally challenging. Instead, we aim to find over-approximations $\bar{\mathcal{R}}_t(\mathcal{X}_0) \supseteq \mathcal{R}_t(\mathcal{X}_0)$ of the reachable sets at time t . Once these reachable set over-approximations are found and shown to avoid the unsafe regions in the state space, we can certify the safety of the NNDS. When the initial set is clear from context, we drop the \mathcal{X}_0 argument from $\bar{\mathcal{R}}_t$ for notational simplicity.

C. Reachability analysis frameworks

We first define the concept of *propagator* as a basic reachability analysis module that abstracts different bounding methods of the output range of NNs.

Definition 1 (Propagator): Any method P that can bound the output of the NNDS $f(x, w)$ given a bounded input set $\mathcal{X} \times \mathcal{W}$ is defined as a *propagator*, i.e., $P(\mathcal{X} \times \mathcal{W}; f) \subset \mathbb{R}^{n_x}$ and $P(\mathcal{X} \times \mathcal{W}; f) \supseteq f(\mathcal{X} \times \mathcal{W})$.


The first argument of the propagator, $\mathcal{X} \times \mathcal{W}$, denotes the input set while the second argument denotes the NN it acts on. The propagator can be obtained by running a NN verification algorithm, and its output set $P(\mathcal{X} \times \mathcal{W})$ often follows a pre-fixed geometric template specific to each NN verification algorithm, e.g., polytope [31], ellipsoid [17], zonotope [32], or polynomial zonotope [14]. In this work, we compare two frameworks that employ propagators to compute the reachable set over-approximations $\bar{\mathcal{R}}_t$ of the NNDS over a finite horizon:

a) *Recursive reachability analysis* which applies the propagator P for one-step reachable set over-approximation recursively, i.e., by using the reachable set over-approximation $\bar{\mathcal{R}}_t$ as the input set for computing $\bar{\mathcal{R}}_{t+1}$, as follows:

$$\begin{aligned} \text{Recursive: } \bar{\mathcal{R}}_0 &= \mathcal{X}_0, \\ \bar{\mathcal{R}}_{t+1} &= P(\bar{\mathcal{R}}_t \times \mathcal{W}; f), \quad t \geq 0. \end{aligned} \quad (4)$$

b) *One-shot reachability analysis* which applies the propagator on the t -th order composition of the NNDS (1) to

¹The one-shot framework is denoted as the symbolic approach in [18].



(a) Backward bounds propagation. (b) Forward bounds propagation.

Fig. 3: Comparison of one-shot and recursive analysis of a NNDS approximating a Rayleigh-Duffing oscillator. Box reachable set over-approximations obtained by the backward (left figure)/forward (right figure) linear bounds propagation are plotted along with randomly sampled trajectories of the NNDS. In the left figure, the recursively synthesized bounds are shown for only 3 steps. In the right figure, bounds over 2 steps are plotted for both one-shot and recursive frameworks.

compute $\bar{\mathcal{R}}_t$ directly for $t > 1$, as follows:

$$\begin{aligned} \text{One-shot: } \bar{\mathcal{R}}_0 &= \mathcal{X}_0, \\ \bar{\mathcal{R}}_t &= P(\mathcal{X}_0 \times \mathcal{W}^t; f^{(t-1)}), t \geq 1, \end{aligned} \quad (5)$$

where $f^{(t)}(x_0, w_{0:t})$ denotes the t -th order composition of the NNDS, i.e., $x_{t+1} = f^{(t)}(x_0, w_{0:t})$, and \mathcal{W}^t denotes the Cartesian product of t disturbance sets. The notation $w_{0:t}$ is shorthand for the set $\{w_0, \dots, w_t\}$. Due to the cascading structure of NNs, one can easily represent the composition of the NNDS (1) as a concatenated network.

The recursive framework for reachable set over-approximations is known to be susceptible to the ‘wrapping effect’ [33], [34] which can quickly lead to a dramatic approximation error. Intuitively, the one-shot framework can mitigate this issue and generate tighter $\bar{\mathcal{R}}_t$ than the recursive one since it utilizes the exact description $f^{(t)}(x_0, w_{0:t})$ of x_{t+1} for reachable set over-approximation. However, we find a counterexample where the opposite holds.

Example 1 (Counterexample): We approximate the discretized dynamics of the 2D Rayleigh-Duffing oscillator [35] as a feed-forward ReLU NN. No disturbances are considered. In Fig. 3, the box reachable set over-approximations obtained from the one-shot and recursive frameworks are compared using the backward linear bounds propagation (left figure) and forward linear bounds propagation (right figure), respectively (see [21] for details). While the one-shot framework gives tighter bounds than the recursive one in the former case, the opposite holds in the latter.

D. Problem formulation

Motivated by Example 1, in this paper we investigate the following question:

In what cases is it beneficial to verify or optimize over the unrolled NN dynamics in one-shot?

In Section III, we characterize a general class of NN verification methods that are guaranteed to benefit from the one-shot analysis. Our proof applies to NNDS with arbitrary architecture. The implication of our analysis for verification

algorithm design is discussed in Section IV, with numerical examples validating our analysis in Section V.

III. TIGHTNESS IMPROVEMENT GUARANTEES

NNs implement mathematical functions that can be represented by a computational graph. In this section, we first describe the graph representation of NNs with arbitrary architecture, based on which the class of separable NN verification methods and propagators are defined. Then, a formal tightness improvement guarantee of the one-shot framework for separable propagators is provided in Theorem 1.

A. NN representation

Similar to [21], we can represent a general NN through its computational graph, which is defined as a directed acyclic graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$, where $\mathbf{V} = \{z_0, \dots, z_q, z_{q+1}, \dots, z_L\}$. We denote $\mathcal{I} = \{z_0, \dots, z_q\}$ as the input nodes or independent nodes of the graph and $\mathcal{D} = \{z_{q+1}, \dots, z_L\}$ as the dependent nodes. The edge set \mathbf{E} is a set of pairs (i, j) denoting that node z_i is an input argument of node z_j . Define $\text{Pre}(z_j) = \{z_{j,1}, \dots, z_{j,m(j)}\}$ as the set of input nodes to z_j with cardinality $|\text{Pre}(z_j)| = m(j)$. With this notation, we can describe the operator generating z_j as $z_j = \phi_j(\text{Pre}(z_j))$ where ϕ_j can represent any commonly used operators in deep NNs such as linear, convolutional, MaxPooling, and activation layers, etc. We have $|\text{Pre}(z_j)| = 0$ for all input nodes $z_j \in \mathcal{I}$, and $|\text{Pre}(z_j)| > 0$ for the rest. All maps from the input nodes \mathcal{I} to any dependent node $z_j \in \mathcal{D}$ are defined by the computational graph \mathbf{G} and the operators $\phi_j(\cdot)$ associated to each node. We assume z_L is the output of the NN whose property is of our concern.

Connection to NNDS: Given $T > 0$, the unrolled NNDS $x_T = f^{(T)}(x_0, w_{0:T-1})$ for T steps can be considered as a single NN with the computational graph \mathbf{G} . In this graph, the input nodes represent $\{x_0, w_0, \dots, w_{T-1}\}$ and the dependent nodes represent all intermediate variables including states $x_{1:T-1}$. The state x_T is the output node of \mathbf{G} , and our goal is to bound the reachable set of x_T for safety verification.

B. Separable propagator

In this paper, we consider NN verification methods that solve an optimization problem based on a layer-wise abstraction of $\phi_j(\cdot)$ in the NN computational graph. Specifically, these methods verify the properties of a NN by solving problems of the following form:

$$\begin{aligned} & \underset{z_{0:L}}{\text{minimize}} && J(z_L) \\ & \text{subject to} && (\text{Pre}(z_k), z_k) \in \mathcal{S}_k, k = q+1, \dots, L, \\ & && z_k \in \mathcal{X}_k, k = 0, \dots, q, \end{aligned} \quad (6)$$

where J is a real-valued function encoding the specification to be verified, \mathcal{X}_k is the bounded input set of each input node z_k for $0 \leq k \leq q$, and \mathcal{S}_k denotes the constraints describing the relationship between node z_k and its input nodes. We can view \mathcal{S}_k as a set defined on the joint space of $(\text{Pre}(z_k), z_k)$, possibly described by expressions with intermediate variables, as shown in the following example.

Example 2 (Optimization constraints): Consider a scalar ReLU function $y = \max(x, 0)$. Given lower and upper bounds on the scalar input $x \in [\ell, u]^2$, the constraint $(x, y) \in \mathcal{S}$ in different NN verification methods is formulated as:

a) Mixed-integer linear constraints [3]:

$$y \geq 0, y \geq x, y \leq x - (1 - \mu)\ell, y \leq \mu u, \mu \in \{0, 1\}. \quad (7)$$

b) Linear constraints [5]:

$$y \geq 0, y \geq x, y \leq \frac{u}{u - \ell}x - \frac{u\ell}{u - \ell}. \quad (8)$$

c) Simplified linear constraints [6] that are amenable to linear bounds propagation:

$$y \leq \frac{u}{u - \ell}x - \frac{u\ell}{u - \ell}, y \geq \alpha x. \quad (9)$$

where $\alpha \in [0, 1]$ is a user-defined parameter.

In (7), an intermediate binary variable μ is introduced to formulate \mathcal{S} as the graph of the ReLU function. In (8), (9), the set \mathcal{S} is a polyhedron in the space of (x, y) . The above abstractions of a ReLU function can be easily extended to all neurons in a NN, leading to the verification problem (6).

If a NN verifier does not consider the coupling between nodes from different layers, the constraints \mathcal{S}_k in problem (6) are separable as defined below.

Definition 2 (Separable constraints): Constraints \mathcal{S}_i and \mathcal{S}_j such that $(x_i, y_i) \in \mathcal{S}_i, (x_j, y_j) \in \mathcal{S}_j$ are separable if the intermediate variables that define \mathcal{S}_i and \mathcal{S}_j are independent.

Definition 3 (Separable NN verifier): We denote any NN verification method that solves an optimization problem of the form (6) with separable constraints \mathcal{S}_k a separable NN verifier.

Generality of separable NN verifiers: Separable NN verifiers include verification methods that either explicitly or implicitly solve a primal optimization problem (6) with separable constraints. For example, for a ReLU network, all complete verifiers [4], [8], [36] are considered separable since they are equivalent to the MIP-based method [3]. Similarly, scalable Lagrangian-based methods [37], [38] that solve the LP-based verification problem using the dual formulation are separable due to the strong duality of LP. Importantly, we observe that backward linear bounds propagation methods [6], [21], which are scalable and can handle general activation functions, are also separable since they exactly solve an LP verification problem built on linear bounds abstractions of nonlinear activation functions with constraints similar to (9). On the contrary, we cannot identify any implicit separable optimization problem for the forward linear bounds propagation method [21]. In fact, the forward method is used to construct the counterexample in Example 1.

Construction of the propagator: A propagator can be constructed by applying a NN verifier to solve problem (6) multiple times with different linear objective functions $J_i(z_L) = c_i^\top z_L, i = 1, \dots, N$ in order to obtain a polytopic

²We consider the non-trivial case where $\ell < 0, u > 0$. Otherwise, the ReLU layer is reduced to a linear one.

over-approximation of the output set of the NN. In what follows, we assume the normal vectors $\{c_i\}_{i=1}^N$ defining the polytopic over-approximation are fixed. A propagator is called separable if it applies a separable NN verifier to bound the NN outputs.

C. Tightness comparison

Recall that we can represent the T -step unrolled NNDS by a computational graph \mathbf{G} with x_T being the output node. In the one-shot framework, a propagator is directly applied to the computational graph \mathbf{G} . In the recursive framework, sub-networks of \mathbf{G} are extracted such that the ranges of the intermediate nodes $x_{1:T-1}$ are over-approximated by the propagator as $\bar{\mathcal{R}}_t$ sequentially. Once the bound $\bar{\mathcal{R}}_t$ on the range of the intermediate state x_t for $1 \leq t \leq T - 1$ is computed, x_t is treated as an input node to the downstream dependent nodes with input set $\bar{\mathcal{R}}_t$. We denote this process as *concretization* of a dependent node in \mathbf{G} .

We now show that for a separable propagator, concretizing the range of any dependent nodes in the NN computational graph \mathbf{G} increases conservatism in bounding the output node. This immediately indicates that in reachability analysis with a separable propagator, the recursive framework is more conservative than the one-shot framework. Our analysis is conducted in two steps:

a) *Recursive framework:* In the recursive framework, we first bound the range of a dependent node z_k for some $q < k < L$. This can be achieved by applying a NN verifier solving problem (6) with objectives $J_i(z_k) = c_i^\top z_k, i = 1, \dots, N$ and the truncated constraints corresponding to the sub-network with z_k as the output. To extract relevant constraints of this sub-network, we apply Breath First Search (BFS) on the graph \mathbf{G} as shown in Algorithm 1. The set of extracted constraints is denoted as $\mathcal{S}_{input \rightarrow k} = \text{EXTRACTCONSTRAINTS}(\text{InputNodes} = z_{0:q}, \text{OutputNodes} = z_k)^3$, and the NN verifier solves

$$\begin{aligned} & \text{minimize} && J_i(z_k) \\ & \text{subject to} && (z_{0:q}, z_k) \in \mathcal{S}_{input \rightarrow k}, \\ & && z_{0:q} \in \mathcal{X}, \end{aligned} \quad (10)$$

for $i = 1, \dots, N$ where the input set is given by $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_q$. Denote J_i^* the optimal values of (10). Then, the over-approximation of the range of z_k is given by ⁴

$$\bar{\mathcal{X}}_k = \{z \mid J_i(z) \geq J_i^*, i = 1, \dots, N\}. \quad (11)$$

Next, we treat z_k as an input node with domain $\bar{\mathcal{X}}_k$ to bound z_L . We consider the sub-network with $z_{0:q} \cup \{z_k\}$ as inputs and z_L as output. We extract the associated constraints as $\mathcal{S}_{input, k \rightarrow L} = \text{EXTRACTCONSTRAINTS}(\text{InputNodes} = z_{0:q} \cup \{z_k\}, \text{OutputNodes} = z_L)$. Bounding the range of z_L is now

³We omit the arguments of the NN computational graph \mathbf{G} and the constraint set $\{\mathcal{S}_i\}$ from (6) since they are clear from context.

⁴Here we consider arbitrary dependent nodes z_k in the computational graph of the unrolled NNDS. If the chosen node corresponds to state x_t , the over-approximation set $\bar{\mathcal{X}}_k$ from (11) coincides with $\bar{\mathcal{R}}_t$.

achieved by solving

$$\begin{aligned} & \text{minimize} && J_i(z_L) \\ & \text{subject to} && (z_{0:q} \cup z_k, z_L) \in \mathcal{S}_{input, k \rightarrow L}, \\ & && z_{0:q} \in \mathcal{X}, z_k \in \bar{\mathcal{X}}_k, \end{aligned} \quad (12)$$

for $i = 1, \dots, N$. Denote the optimal values of (12) as $J_{i, re}^*$. The over-approximation of z_L obtained by the recursive framework is given by $\bar{\mathcal{X}}_{L, re} = \{z \mid J_i(z) \geq J_{i, re}^*, 1 \leq i \leq N\}$.

b) One-shot framework: In the one-shot analysis framework, we directly solve problem (6) with objectives $J_i(z_L)$ for $i = 1, \dots, N$. Denote the optimal values of these problems as $J_{i, os}^*$. Then, the over-approximation of z_L obtained by the one-shot framework is given by $\bar{\mathcal{X}}_{L, os} = \{z \mid J_i(z) \geq J_{i, os}^*, 1 \leq i \leq N\}$.

Algorithm 1 Extract concretization constraints

Input: InputNodes $\subseteq \{z_i\}_{i=0}^L$, OutputNodes z_k , computational graph \mathbf{G} , constraints $\{\mathcal{S}_i\}_{i=q+1}^L$ from problem (6).

Output: Constraint set \mathcal{S} .

- 1: **procedure** EXTRACTCONSTRAINTS
 - 2: Let Q be a queue and $\mathcal{S} = \emptyset$.
 - 3: $Q.enqueue(z_k)$. Mark node z_k as explored.
 - 4: **while** Q is not empty **do**
 - 5: $z_i = Q.dequeue()$. ▷ Subscript i denotes the label of the node in graph \mathbf{G} .
 - 6: $\mathcal{S} = \mathcal{S} \cup \mathcal{S}_i$.
 - 7: **for** z_j in $\text{Pre}(z_i)$ **do**
 - 8: **if** $z_j \notin \text{InputNodes}$ and is not explored **then**
 - 9: $Q.enqueue(z_j)$ and label z_j as explored.
-

Theorem 1: With a separable propagator, the one-shot framework generates a tighter over-approximation of the NN output than the recursive framework, i.e., $\bar{\mathcal{X}}_{L, os} \subseteq \bar{\mathcal{X}}_{L, re}$.

Proof: To allow a straightforward comparison, we combine the two-step computation (10) and (12) of the recursive framework into one optimization problem by introducing auxiliary variables $\tilde{z}_{0:L}$ and $\tilde{z}_{0:L}^j$ for $j = 1, \dots, N$. We can equivalently rewrite (12) as

$$\begin{aligned} & \text{minimize} && J_i(\tilde{z}_L) \\ & \text{subject to} && (\tilde{z}_{0:q}^j, \tilde{z}_k^j) \in \mathcal{S}_{input \rightarrow k}, \tilde{z}_{0:q}^j \in \mathcal{X}, \\ & && J_i(\tilde{z}_k) \geq J_i(\tilde{z}_k^j), j = 1, \dots, N, \\ & && (\tilde{z}_{0:q} \cup \tilde{z}_k, \tilde{z}_L) \in \mathcal{S}_{input, k \rightarrow L}, \\ & && \tilde{z}_{0:q} \in \mathcal{X}. \end{aligned} \quad (13)$$

By the separability of constraints \mathcal{S}_k in (6), we have that any feasible solution $z_{0:L}$ of (6) constructs a feasible one of (13) by setting $\tilde{z}_\ell = \tilde{z}_\ell^j = z_\ell$ for all $0 \leq \ell \leq L$, $1 \leq j \leq N$ while achieving the same objective value. Then we have $J_{i, os}^* \geq J_{i, re}^*$ for all $1 \leq i \leq N$ and $\bar{\mathcal{X}}_{L, os} \subseteq \bar{\mathcal{X}}_{L, re}$. ■

The proof of Theorem 1 shows that the performance gap between the one-shot and recursive frameworks arises from the concretization errors, i.e., the over-approximation error of applying a finite number of hyperplanes to bound the

reachable set of a dependent node in the NN. Such errors may compound when we concretize the ranges of a sequence of dependent nodes one-by-one as in recursive reachability analysis of the NNDS (1).

Corollary 1: The one-shot framework is guaranteed to generate tighter over-approximations of reachable sets than the recursive framework for the NNDS (1) when a separable propagator is applied.

Proof: We can treat the T -step unrolled NNDS as the NN with input nodes $(x_0, w_0, \dots, w_{T-1})$ and x_T as the output node. By Theorem 1, concretizing any intermediate dependent node x_t leads to more conservative over-approximation of x_T . ■

IV. PRACTICAL IMPLICATIONS

We now summarize the comparison between the one-shot and recursive frameworks: (a) The one-shot framework solves the NN verification problem on NNs $f^{(t)}(\cdot)$ with growing sizes as t increases and is therefore always more computationally costly to run than the recursive framework⁵. (b) The recursive framework generates looser bounds than the one-shot framework due to the concretization error as shown in Theorem 1 and the induced conservatism could explode as time t increases.

One-shot analysis principle: The above two points indicate that we can treat the size of NNs, or the order of composition of the NNDS $f(\cdot)$, as a tuning parameter which makes a trade-off between tightness and computational complexity. Within a given computational resource budget (either in time or in memory), our analysis suggests that we should run a separable propagator on the composition $f^{(t)}$ of the NNDS with t as large as possible.

Concretization error reduction: Problem (13) indicates that increasing N tightens the search space of \tilde{z}_k and consequently the output variable \tilde{z}_L . Therefore, the concretization error from the recursive framework can be decreased by considering a richer set of objective functions $\{c_i^\top z_L\}$ in constructing the propagator. This phenomenon is demonstrated in Fig. 4. However, the overall complexity of the recursive framework increases with N , and it remains a challenge to identify a set of bounding hyperplanes $\{c_i\}_{i=1}^N$ with low conservatism for systems with high dimensions.

V. NUMERICAL EXAMPLES

We demonstrate the improvement of tightness of the one-shot analysis on a NNDS that approximates the closed-loop dynamics of a cart-pole system. Separable propagators relying on LP and backward linear bounds propagation are considered with different NNDS structures. All experiments are run on an Intel Core i7-6700K CPU.

A. LP-based propagator

We consider the cart-pole dynamics without friction from [39] where the cart has a mass of 0.25 kg, and the pole

⁵The scaling of the computational complexity with respect to the horizon t in the one-shot framework depends on the propagator applied.

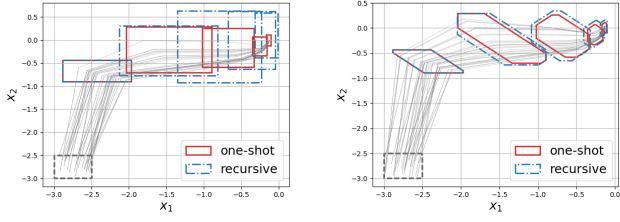


Fig. 4: The LP-based propagator using relaxation (8) is applied to bound the reachable sets of the Rayleigh-Duffing oscillator NNDS using a box template (left, 4 hyperplanes) and a more complex polyhedral template (right, 8 hyperplanes). Randomly sampled trajectories of the NNDS are shown in gray. The tightness gap between the one-shot and recursive frameworks decreases as the complexity of the bounding templates increases.




Fig. 5: Box reachable set over-approximations of the feedforward NNDS (14) obtained by the one-shot (solid red boxes, 30 steps) and recursive (dashed blue boxes, 8 steps) frameworks using the LP-based propagator. The boxes are projected onto the (x_1, x_2) and (x_3, x_4) planes for plotting. Randomly sampled trajectories of the NNDS are shown in gray.

has a mass of 0.1 kg and length of 0.4 m. The entries of the 4-dimensional state $x = [x_1 \ x_2 \ x_3 \ x_4]^T$ represent the position, the velocity of the cart, the angle, and the angular speed of the pole, respectively. We train a feedforward ReLU NN $f_{cl,NN}(x)$ with the structure 4–100–100–4 to approximate the discretized closed-loop dynamics of the cart-pole system under a nonlinear model predictive controller (MPC) [40] with $dt = 0.05s$. An LP-based verifier that bounds ReLU by (8) is applied to compute reachable sets of the NNDS

$$x_{t+1} = f_{cl,NN}(x_t). \quad (14)$$

The pre-activation bounds (ℓ, u) in (8) are computed in an inductive manner by the verifier. The initial set is given by $\mathcal{X}_0 = \{x \mid x_1 \in [2.0, 2.2] \text{ m}, x_2 \in [1.0, 1.2] \text{ m/s}, x_3 \in [-0.174, -0.104] \text{ rad}, x_4 \in [-1.0, -0.8] \text{ rad/s}\}$. In Fig. 5, the over-approximating boxes obtained by the recursive method (dashed blue) for 8 steps are plotted which quickly become overly conservative, while the one-shot method is able to provide informative bounds (solid red) on the system states for 30 steps. The LP propagators are solved by Gurobi [41], and the total solver time for running the one-shot and recursive methods for 8 steps are 161.3 and 3.1 seconds, respectively.

B. Backward propagation-based propagator

Next, we adopt a different approximation of the closed-loop dynamics of the cart-pole system which has a more




Fig. 6: Box reachable set over-approximations of the closed-loop NNDS (15) obtained by backward linear bounds propagation are plotted for 3 steps. Bounds produced by the one-shot framework (red boxes) are tighter than those from the recursive framework (blue boxes). Randomly sampled trajectories are plotted in gray.

complex structure:

$$x_{t+1} = Ax_t + Bu_t + f(x_t, u_t) + w_t \quad (15a)$$

$$u_t = \pi_{NN}(x_t) \quad (15b)$$

where the linear dynamics (A, B) are derived from the Jacobian matrices of the cart-pole dynamics at the origin. A feedforward NN $f(x, u) : \mathbb{R}^5 \mapsto \mathbb{R}^4$ is trained to learn the residual nonlinear dynamics with training data randomly sampled from the discretized cart-pole dynamics with $dt = 0.05s$. We implement a nonlinear MPC for the open-loop NNDS (15a) in MATLAB to collect samples of (x_t, u_t) and train the NN controller $\pi(\cdot) : \mathbb{R}^4 \mapsto \mathbb{R}$ in PyTorch to approximate the MPC policy. The residual NN dynamics $f(x, u)$ in (15a) is a feedforward tanh network of structure 5–50–4, and the controller $\pi(x)$ in (15b) is a feedforward ReLU network of structure 4–100–1. The disturbance w_t is assumed bounded by $\|w_t\|_\infty \leq 0.05$. We compare the one-shot and recursive frameworks on system (15) using backward linear bounds propagation [21] since it can handle general activation functions. The results are shown in Fig. 6 where the initial set is chosen as $\mathcal{X}_0 = \{x \mid x_1 \in [0.0, 0.3] \text{ m}, x_2 \in [-0.2, -0.1] \text{ m/s}, x_3 \in [0.262, 0.312] \text{ rad}, x_4 \in [-0.15, -0.05] \text{ rad/s}\}$ and the horizon is set as $T = 3$ ⁶. We observe that the one-shot framework generates tighter bounds than the recursive framework, validating Theorem 1. The total solver time is 2.22s for the one-shot framework and 1.01s for the recursive framework.

VI. CONCLUSION

We proved that for a wide range of NN verification methods with a separable structure, analyzing the unrolled dynamics of a NN dynamical system generates tighter bounds than a recursive, step-by-step analysis. Our proof holds for NN dynamical systems with general computational graphs. The applicability of one-shot reachability analysis was demonstrated numerically. Future work includes designing adaptive bounding polytopes such that the performance gap between the recursive and one-shot frameworks is minimized.

⁶Backward linear bounds propagation solves a weaker relaxation of the NN verification problem than LP-based methods. Therefore, in this example, we only consider $T = 3$ since both the one-shot and recursive frameworks become too conservative afterward.

REFERENCES

- [1] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- [2] S. Huang, N. Papernot, I. Goodfellow, Y. Duan, and P. Abbeel, "Adversarial attacks on neural network policies," *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [3] V. Tjeng, K. Xiao, and R. Tedrake, "Evaluating robustness of neural networks with mixed integer programming," *arXiv preprint arXiv:1711.07356*, 2017.
- [4] R. Ehlers, "Formal verification of piece-wise linear feed-forward neural networks," in *International Symposium on Automated Technology for Verification and Analysis*, pp. 269–286, Springer, 2017.
- [5] E. Wong and Z. Kolter, "Provable defenses against adversarial examples via the convex outer adversarial polytope," in *International Conference on Machine Learning*, pp. 5286–5295, PMLR, 2018.
- [6] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, "Efficient neural network robustness certification with general activation functions," *Advances in neural information processing systems*, vol. 31, 2018.
- [7] K. Dvijotham, R. Stanforth, S. Gowal, T. A. Mann, and P. Kohli, "A dual approach to scalable verification of deep networks.," in *UAI*, vol. 1, p. 3, 2018.
- [8] R. R. Bunel, I. Turkaslan, P. Torr, P. Kohli, and P. K. Mudigonda, "A unified view of piecewise linear neural network verification," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [9] G. Singh, R. Ganvir, M. Püschel, and M. Vechev, "Beyond the single neuron convex barrier for neural network certification," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [10] A. Raghunathan, J. Steinhardt, and P. S. Liang, "Semidefinite relaxations for certifying robustness to adversarial examples," *Advances in neural information processing systems*, vol. 31, 2018.
- [11] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming," *IEEE Transactions on Automatic Control*, 2020.
- [12] M. Akintunde, A. Lomuscio, L. Maganti, and E. Pirovano, "Reachability analysis for neural agent-environment systems," in *Sixteenth international conference on principles of knowledge representation and reasoning*, 2018.
- [13] C. Huang, J. Fan, W. Li, X. Chen, and Q. Zhu, "Reachnn: Reachability analysis of neural-network controlled systems," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 18, no. 5s, pp. 1–22, 2019.
- [14] N. Kochdumper, C. Schilling, M. Althoff, and S. Bak, "Open-and closed-loop neural network verification using polynomial zonotopes," *arXiv preprint arXiv:2207.02715*, 2022.
- [15] Y. Zhang and X. Xu, "Safety verification of neural feedback systems based on constrained zonotopes," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 2737–2744, IEEE, 2022.
- [16] M. Everett, G. Habibi, and J. P. How, "Efficient reachability analysis of closed-loop systems with neural network controllers," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4384–4390, IEEE, 2021.
- [17] H. Hu, M. Fazlyab, M. Morari, and G. J. Pappas, "Reach-sdp: Reachability analysis of closed-loop systems with neural network controllers via semidefinite programming," in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 5929–5934, IEEE, 2020.
- [18] C. Sidrane, A. Maleki, A. Irfan, and M. J. Kochenderfer, "Overt: An algorithm for safety verification of neural network control policies for nonlinear systems," *Journal of Machine Learning Research*, vol. 23, no. 117, pp. 1–45, 2022.
- [19] B. Karg and S. Lucia, "Stability and feasibility of neural network-based controllers via output range analysis," in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 4947–4954, IEEE, 2020.
- [20] L. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, L. Daniel, D. Boning, and I. Dhillon, "Towards fast computation of certified robustness for relu networks," in *International Conference on Machine Learning*, pp. 5276–5285, PMLR, 2018.
- [21] K. Xu, Z. Shi, H. Zhang, Y. Wang, K.-W. Chang, M. Huang, B. Kailkhura, X. Lin, and C.-J. Hsieh, "Automatic perturbation analysis for scalable certified robustness and beyond," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [22] H. Salman, G. Yang, H. Zhang, C.-J. Hsieh, and P. Zhang, "A convex relaxation barrier to tight robustness verification of neural networks," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [23] A. De Palma, R. Bunel, A. Desmaison, K. Dvijotham, P. Kohli, P. H. Torr, and M. P. Kumar, "Improved branch and bound for neural network verification via lagrangian decomposition," *arXiv preprint arXiv:2104.06718*, 2021.
- [24] W. Xiang, H.-D. Tran, J. A. Rosenfeld, and T. T. Johnson, "Reachable set estimation and safety verification for piecewise linear systems with neural network controllers," in *2018 Annual American Control Conference (ACC)*, pp. 1574–1579, IEEE, 2018.
- [25] M. Everett, G. Habibi, C. Sun, and J. P. How, "Reachability analysis of neural feedback loops," *IEEE Access*, vol. 9, pp. 163938–163953, 2021.
- [26] H.-D. Tran, F. Cai, M. L. Diego, P. Musau, T. T. Johnson, and X. Koutsoukos, "Safety verification of cyber-physical systems with reinforcement learning control," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 18, no. 5s, pp. 1–22, 2019.
- [27] R. Ivanov, J. Weimer, R. Alur, G. J. Pappas, and I. Lee, "Verisig: verifying safety properties of hybrid systems with neural network controllers," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pp. 169–178, 2019.
- [28] C. Schilling, M. Forets, and S. Guadalupe, "Verification of neural-network control systems by integrating taylor models and zonotopes," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 8169–8177, 2022.
- [29] S. Dutta, X. Chen, and S. Sankaranarayanan, "Reachability analysis for neural feedback systems using regressive polynomial rule inference," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pp. 157–168, 2019.
- [30] S. Dutta, S. Jha, S. Sankaranarayanan, and A. Tiwari, "Learning and verification of feedback control systems using feedforward neural networks," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 151–156, 2018.
- [31] G. Singh, T. Gehr, M. Püschel, and M. Vechev, "An abstract domain for certifying neural networks," *Proceedings of the ACM on Programming Languages*, vol. 3, no. POPL, pp. 1–30, 2019.
- [32] G. Singh, T. Gehr, M. Mirman, M. Püschel, and M. Vechev, "Fast and effective robustness certification," *Advances in neural information processing systems*, vol. 31, 2018.
- [33] W. Kühn, "Rigorously computed orbits of dynamical systems without the wrapping effect," *Computing*, vol. 61, pp. 47–67, 1998.
- [34] A. Neumaier, "The wrapping effect, ellipsoid arithmetic, stability and confidence regions," in *Validation numerics*, pp. 175–190, Springer, 1993.
- [35] J. Giné and C. Valls, "On the dynamics of the rayleigh–duffing oscillator," *Nonlinear Analysis: Real World Applications*, vol. 45, pp. 309–319, 2019.
- [36] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient smt solver for verifying deep neural networks," in *International conference on computer aided verification*, pp. 97–117, Springer, 2017.
- [37] R. Bunel, A. De Palma, A. Desmaison, K. Dvijotham, P. Kohli, P. Torr, and M. P. Kumar, "Lagrangian decomposition for neural network verification," in *Conference on Uncertainty in Artificial Intelligence*, pp. 370–379, PMLR, 2020.
- [38] S. Chen, E. Wong, J. Z. Kolter, and M. Fazlyab, "Deepsplit: Scalable verification of deep neural networks via operator splitting," *IEEE Open Journal of Control Systems*, vol. 1, pp. 126–140, 2022.
- [39] R. V. Florian, "Correct equations for the dynamics of the cart-pole system," *Center for Cognitive and Neural Studies (Coneural), Romania*, 2007.
- [40] S. Lucia, A. Tăulea-Codrean, C. Schoppmeyer, and S. Engell, "Rapid development of modular and sustainable nonlinear model predictive control solutions," *Control Engineering Practice*, vol. 60, pp. 51–62, 2017.
- [41] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2022.